



Wind farm value, revenue-triggered curtailment and support-scheme design under wind- and solar-driven negative prices

Moritz Gräfe¹, Samuel Kainz², Vasilis Pettas¹, Anastasia Ioannou¹, Athanasios Kolios¹, Tuhfe Göçmen¹, Iñaki Sandua-Fernández³, Irene Eguinoa³, and Nikolay Dimitrov¹

¹DTU Wind and Energy Systems, Technical University of Denmark, Frederiksborgvej 399, 4000 Roskilde, Denmark

²Wind Energy Institute, Technical University of Munich, Boltzmannstr. 15, 85748 Garching b. München, Germany

³Wind Energy Department, National Renewable Energy Centre (CENER), Sarriguren, Spain

Correspondence: Moritz Gräfe (mograf@dtu.dk)

Abstract. Negative electricity prices are becoming more frequent in Europe, and whether they fall in windy hours depends on whether wind or solar feed-in causes the surplus. We ask how this driver affects the operation, the economics and the environmental value of a wind farm that stops its turbines on a revenue signal. We simulate the Lillgrund offshore wind farm hour by hour in one model of control, fatigue loads, contract settlement, financing and carbon accounting, on three constructed price futures with the same share of negative-price hours, driven by wind, by wind and solar, and by solar. One farm-wide rule stops all turbines in hours in which the projected farm revenue falls below a threshold, which we vary from zero to low positive values. The controlled shutdown removes the hours of lowest revenue, and the driver decides what else these hours carry. Under wind-driven prices they are windy hours with high fatigue loads, and the fatigue budget left unused permits a lifetime extension of about ten years. Under solar-driven prices they are hours of average wind, and the extension is about one year. The years gained are worth a few million euros in every future. The net displacement of emissions rises by about 17 % under wind-driven prices and falls slightly under solar-driven prices. Three contracts for difference require the same support in present value, but the strike of the production-based contract, which pays nothing on energy delivered at negative prices, rises by about 40 % under wind-driven prices, while the strikes of two capability-based contracts, which pay on the energy the farm could have delivered, stay near their historical level. Raising the threshold above zero adds little value here, and a contract settled on delivered energy blocks any load-reducing control whose fatigue value lies below the difference payment. Only a contract that pays on available energy at every price leaves the decision to the operator.

1 Introduction

The share of hours with negative day-ahead electricity prices in northwestern European bidding zones has grown persistently with the deployment of variable renewables (ACER, 2024; Prokhorov and Dreisbach, 2022). Prices fall below zero when wind and solar feed-in exceeds demand and part of the generation does not reduce output (ACER, 2025). System studies expect such surplus hours to become more frequent, and they differ widely in how many of them clear at negative prices (ENTSO-E and ENTSOG, 2025; Danish Energy Agency, 2024; Seyyedeh-Barhagh et al., 2026). Which hours they are depends on the generation mix. Where wind causes the surplus, the negative-price hours are windy hours. Where solar causes it, they are



midday hours. For a wind farm this driver of the negative prices decides what a given share of negative-price hours costs.
 25 Under wind-driven negative prices the farm loses revenue in the hours in which it produces most. Under solar-driven negative prices it produces at or below its average in those hours and loses less.

The operator can respond by curtailing. Stopping production in hours of negative revenue gives up energy and avoids the loss. It also changes the fatigue accumulation of the turbines, by an amount that again depends on which hours are curtailed. Whether an operator takes this response is decided by the support contract and not by the market price alone. Where payments
 30 are withheld in negative-price hours, production drops (Hirth, 2026; Van Steenberghe and Ovaere, 2026). Where payments continue, generators have been found to keep producing through negative prices (Van Steenberghe and Ovaere, 2026). The settlement function of the contract determines whether a curtailed hour forfeits support, and who carries the consequences when the driver changes during the life of the contract.

Two research domains each address part of this situation, and each holds fixed what the other varies. Wind farm control
 35 research quantifies the physical side. Curtailment and derating change fatigue accumulation and can extend structural lifetime (Meyers et al., 2022). In this domain the market enters as a price level or an exogenous series, the support contract is absent, and loads are coupled to money in a separate model chain or by scaling after the fact (Réthoré et al., 2014). By contrast, support-scheme research addresses the contract side and treats the settlement function as an instrument of risk allocation and operational incentives (Schlecht et al., 2024; Ioannou et al., 2026). In this domain operation has no physical consequence. A
 40 curtailed hour changes cash flows but never loads or lifetime. The curtailment decision sits at the intersection that the two domains leave open. Evaluating it requires five elements together: a controlled price-shape axis, loads carried in the same simulation as revenue, a comparison of settlement functions, strikes solved endogenously, and a financing layer. Each element has precedent in its own domain: parametric price axes (Mehta et al., 2024), joint revenue–load operation (Zheng et al., 2026; Anand et al., 2025), endogenous design-dependent strikes (Johanndeiter et al., 2025a) and contract design against debt sizing
 45 (Đukan et al., 2025). To our knowledge, no prior study combines them. Section 2 discusses the closest studies.

Beyond the economic consequences, the resulting control strategies also have environmental implications. On the one hand, changes in fatigue loading may affect component reliability and, in turn, the emissions associated with plant maintenance over its lifetime (Gräfe et al., 2025). On the other hand, changes in the timing of electricity production affect the greenhouse-gas emissions displaced in the connected electricity grid (Wheatley, 2013; Thomson et al., 2017; Kainz et al., 2025). While these
 50 effects have been studied individually, to our knowledge, no study has assessed how varying operational strategies affect net environmental performance over the plant lifetime, considering both plant- and grid-related consequences.

This paper addresses four research questions (RQs), each asked as a function of the driver of the negative prices at a given negative-price share:

- **RQ1** (market value): How does the driver of negative prices change the market remuneration of wind production, and
 55 how much of the loss does revenue-triggered curtailment recover for the operator?
- **RQ2** (physical value): What structural fatigue relief and lifetime extension does revenue-triggered curtailment provide?



- **RQ3** (contract value): How does the settlement function change the required strike, the composition of public payments and the operator’s curtailment incentive?
- **RQ4** (environmental value): How do the settlement function and the price futures affect the net environmental performance of the wind farm, considering both plant-related emissions and displaced grid emissions?

The study answers these questions in hourly simulations of the Lillgrund offshore wind farm, in which fatigue accumulates from the same operation that produces the energy and settles the contracts. Prices follow a historical path and three futures constructed from observed prices. The futures are not forecasts. They hold the negative-price share fixed and vary the driver between wind and solar feed-in. Five settlement functions are compared, each priced at the operation its own rule implies. Section 2 reviews the literature, Sects. 3 and 4 present the methodology and the case study, Sect. 5 gives the results in the order of the research questions, and Sect. 6 discusses them with their limitations.

2 Background and related work

2.1 Negative prices and the market value of wind

Wind revenue depends on the shape of the price distribution, not only on its level. The value factor is the ratio of the capture price, the production-weighted average price, to the time-weighted average price. It declines as the share of variable renewables grows, because their output lowers the price in the hours in which they produce (Hirth, 2013; López Prol et al., 2020; Stiewe et al., 2025). In model results, the decline is steep where subsidies or quotas drive the expansion of wind and solar. It is much smaller where a carbon price drives it (Brown and Reichenberg, 2021). The occurrence of negative prices in European bidding zones has grown alongside the decline (ACER, 2024, 2025). Negative-price hours are the extreme of the same mechanism. In a model with a high wind share, the price falls to zero in the hours in which wind produces most (Hirth, 2013). Solar output lowers the price around midday instead (Hirth, 2013; López Prol et al., 2020). In European zones, solar output has so far had little effect on the value factor of wind (Stiewe et al., 2025). A surplus caused by solar therefore costs a wind farm less than a surplus caused by wind. The growth of solar capacity has widened the midday surplus (ACER, 2025). In several European zones, the negative-price hours have moved from windy hours to sunny hours with average wind (Sect. 4.1).

How often prices turn negative in the future is uncertain. Energy system models usually leave out support payments and the operating limits of power plants, so their prices do not fall below zero (Ruhnau, 2022; Antweiler and Müsgens, 2025; Geis et al., 2026). The number of hours with a price of zero or with curtailment varies widely with the assumed storage capacity, flexible demand, and interconnection (Koivisto et al., 2020; ENTSO-E and ENTSG, 2025). The mechanism behind that spread is the absorption of the surplus. In a system study for Germany, batteries set the price in the surplus hours first, and electrolyzers join them from 2035 on (Geis et al., 2026). Where flexible electrolysis grows with the renewable fleet, it absorbs the surplus at its own willingness to pay, and the market value of wind stabilises. The European network scenarios vary these assumptions together with the generation mix, but they publish no prices (ENTSO-E and ENTSG, 2025). Which technology causes the surplus is therefore as uncertain as how often it occurs.



Stylised or parametric price scenarios have been used in turbine design (Mehta et al., 2024) and in the comparison of support instruments (Kitzing, 2014). In both, the farm produces whenever the wind allows. Mehta et al. (2024) vary the price level and its correlation with wind, and they test a higher price variability as a sensitivity. Their prices are drawn from a normal distribution and include negative hours, and the farm produces in these hours as well. Kitzing (2014) work with monthly price statistics. Neither represents a farm that stops.

2.2 Wind farm control, fatigue and lifetime

Wind farm control research quantifies how operation changes structural loading. The review of Meyers et al. (2022) covers wake steering, induction control and dynamic individual pitch control as instruments that redistribute loading across a farm. It notes that the value of load mitigation is harder to assess than a change in power, because the link from loading to component reliability, remaining lifetime and O&M cost is difficult to establish. Kanev et al. (2018) simulate wake control on several existing commercial farms. They fold the fatigue loads into a lifetime extension factor and maximise the lifetime energy production. No price enters, so the hourly price shape plays no role. Pettas and Cheng (2018) show down-regulation and individual blade control as levers that slow fatigue accumulation and enable lifetime extension, at the turbine level and without a market. Braunbehrens et al. (2024) test wake steering that maximises profit subject to a lifetime constraint in a wind tunnel. Fatigue damage enters through a cost model and the lifetime constraint, and no support contract is represented.

Two companion studies supply the control axis of the present work. Pettas et al. (2026) present a farm-response framework with surrogate-based turbine loads and condensed health metrics. Gräfe et al. (2026a) study revenue-triggered shutdown with lifetime extension under a feed-in tariff, a two-sided contract for difference and market revenue. The shutdown responds to the market revenue alone, and the study uses one historical price series.

The economic value of load-reducing operation depends on the joint hourly timing of price, wind and fatigue accumulation. Studies that weigh energy or revenue against loads commonly approximate this coupling. In Réthoré et al. (2014), damage-equivalent loads from a separate aeroelastic chain enter the financial balance through component degradation cost functions. Kölle et al. (2022) evaluate farm flow control against simulated day-ahead prices. For a single turbine, they also switch hour by hour between shutdown and two pitch-control modes on wind and price, and they report the load changes without monetising them. Pettas (2024) optimises the operation of a single turbine on time series of wind and price, with surrogates for its revenue and for the fatigue damage of several load channels. The gain from managing both together depends on the wind–price correlation of the scenario. Hulsman et al. (2020) combine independent power and load surrogates through a user-set weighting factor. Simley et al. (2024) value wake steering against nodal prices with structural loads outside the model. Requate et al. (2023) optimise energy against damage budgets with the economic valuation deferred to a separate downstream step. Recent work integrates market revenue and priced component damage in one operational framework (Zheng et al., 2026; Anand et al., 2025). Anand et al. (2025) also count the revenue gained or lost through a longer or shorter life. Zheng et al. (2026) shut the farm down at negative prices, with a minimum duration that limits start-stop cycles. Neither study includes a support contract or project financing. No chain carries loads, contract settlement and project financing together. Where a price enters, it is a



given series or scenario. None of the studies varies the price formation as a controlled axis, so none can state under which price formation load-reducing operation has market value.

Lifetime extension adds an assessment question on top of the accumulation question. Lifetime-extension standards assess the remaining fatigue life of each load-carrying component, and the IEC specification does so for each of its failure modes (DNV, 2021; IEC, 2025). The component with the shortest remaining life limits the extension (Ziegler et al., 2018). Techno-economic assessments take the extension period as given and weigh operating costs against revenue (Ziegler et al., 2018). Component-level assessments with an economic case exist (Rubert et al., 2019; Nielsen and Sørensen, 2021) but are the exception. A fatigue measure averaged over components or load channels can diverge from the weakest one when the channels respond heterogeneously to the same operation.

2.3 Environmental performance

Wind farm control also has environmental consequences beyond its economic and structural effects, affecting both the emissions associated with the wind farm itself and the emissions displaced in the connected electricity system. The former is commonly assessed using life-cycle assessment (LCA), covering impacts from manufacturing and installation through operation and maintenance (O&M) to end-of-life (Arvesen and Hertwich, 2012; Bonou et al., 2016). In the operational phase, electricity production is typically represented through estimated capacity factors (Arvesen and Hertwich, 2012), without explicitly modeling the operational strategy of the wind farm. O&M emissions are often either neglected or based on specific component-replacement scenarios (Arvesen and Hertwich, 2012; Garcia-Teruel et al., 2022). More recent work has incorporated detailed failure rates into offshore wind LCA (Garcia-Teruel et al., 2022) and studied the impacts of varying failure rates on the lifetime O&M emissions (Gräfe et al., 2025). However, these approaches prescribe failure rates independently of wind farm operation and therefore do not capture how control-induced changes in fatigue loading propagate to component reliability and O&M emissions.

Beyond emissions associated with the wind farm itself, its operational strategy also affects the emissions displaced in the connected electricity system. The emissions displaced by an additional unit of wind generation can be represented by the marginal displacement factor (MDF), which varies over time as it depends on factors such as market rules, dispatch decisions, cross-border electricity flows, and real-time system conditions (Kainz et al., 2025). To estimate the MDF of wind energy, data-driven approaches have mainly been used, ranging from regression-based models (Wheatley, 2013; Thomson et al., 2017) to more sophisticated machine-learning approaches that capture nonlinear relationships and interactions between multiple system variables. Kainz et al. (2025) combined an artificial neural network (ANN) with a system-balance constraint to estimate technology-specific MDFs. They also examined the effect of wind farm control on displaced emissions over a two-year period, considering additional generation from wake steering. However, the environmental consequences of revenue-based curtailment and other operational strategies, particularly over the wind farm lifetime, remain unexplored.



2.4 Support-scheme design and curtailment incentives

Support-scheme research addresses the contract side. Its focus has moved from the cost of support to how contracts allocate risk and how they shape operating and investment decisions. Beiter et al. (2024) argue that contracts for difference remain the central instrument for allocating price risk and for financing new capacity. The EU electricity market reform takes the same position. From July 2027 it requires direct price support for new wind and solar generation, among other technologies, to take the form of two-way contracts for difference (CfDs) or equivalent schemes, which may include capability-based contracts (European Union, 2024). The Commission followed in December 2025 with design guidance on reference prices, reference volumes and dispatch incentives (European Commission, 2025b). Neuhoﬀ et al. (2022) show that once technology costs approach wholesale prices, the comparison between instruments is dominated by how they allocate risk and hence by the cost of capital they imply. Schlecht et al. (2024) show that the conventional form of the CfD, in which payments follow metered production, distorts short-term operation because the difference payment is earned only by generating. They propose a financial CfD whose payments follow a reference generator modelled from weather data instead of the plant's own output, so that the operator faces the market price at the margin. Newbery (2023) reaches the same end by paying on the plant's capacity times a reference capacity factor for its location, whether the plant generates or not. The capability contracts of this study settle on the farm's own available energy and are closest to this second design (Sect. 3.3). The gated variant does so only in negative-price hours and is otherwise production-based.

The evidence that operators respond to these rules is direct. Where support is withheld in negative-price hours, stopping pays. Germany withholds the market premium for runs of negative prices that reach a set length, and for plants commissioned since February 2025 in every negative-price period. Which plants still have a reason to keep producing therefore depends on their remuneration class and commissioning year (Kern et al., 2024; Hirth, 2026). In 2025, almost all German wind generators had a reason to stop at deeply negative prices, and about half of them did (Hirth, 2026). British contracts from the fourth allocation round onward pay no difference amount in any negative-price hour (Low Carbon Contracts Company and Department for Energy Security and Net Zero (UK), 2024). Where payments continue, the opposite holds. British offshore wind farms under the first CfDs, which carry no negative-price rule, kept producing close to full output in negative-price hours. Comparable merchant farms cut their output by about three quarters in the same hours. The difference payment oﬀsets the negative market price and so removes the reason to stop. Farms whose payments are withheld in runs of six or more negative-price hours cut their output in those runs about as much as the merchant farms (Van Steenberghe and Ovaere, 2026). The same physical action is therefore worth taking under one settlement rule and not under another, whether or not the rule was designed with that response in mind. Scheme design is an operational incentive instrument, not only a risk-allocation instrument.

Three properties of this literature limit its reach for the present question. First, operation has no physical consequence in these models, so the fatigue value of the incentivised behaviour cannot appear. Second, the evaluations diﬀer in their strike convention, which sets what a comparison of designs can show. Where strikes are scenario parameters or administrative prices (Veenstra and Mulder, 2024; Fabra and Imelda, 2023), a comparison of designs is also a comparison of support levels. Where strikes are solved for each design (Johanndeiter et al., 2025a, b; Đukan et al., 2025), the designs are compared at their own



strikes. Kitzing et al. (2017) give each scheme the same expected support and solve the margin at which investment starts. The operational response to the contract is either not represented or evaluated together with the design, and not separated from it. Third, Johanndeiter et al. (2025a) fix the strike at award and report how widely investor profits and consumer prices spread across market scenarios. Which side carries a shift of price formation in one direction, such as from wind-driven to solar-driven negative prices, is not examined there or in the other evaluations. Sections 3.4 and 4.3 describe how this study separates the settlement function from the operation it induces.

3 Methodology

All results derive from WINPACT, an integrated techno-economic assessment framework for wind farm projects (Gräfe et al., 2026a), motivated by the observation that the dominant risks in wind farm development interact across technical and financial domains and resist evaluation in separate models (Gräfe et al., 2026b). The framework couples its modules over one hourly timeline spanning the project life (Fig. 1). This study relies on six modules. The meteorological and market environments supply the resampled wind series, the price futures and the displacement factors (Sects. 3.1, 3.7). The farm-response module of Pettas et al. (2026) converts the hourly inflow of each turbine under the active controller into farm power and fatigue damage per load channel, with an engineering wake model and load surrogates trained on aeroelastic simulations (Sects. 3.2, 3.5). The damage-coupled O&M model carries its own cost and emission accounting (Sect. 3.6). The lifetime-extension module (Sect. 3.5) and the settlement functions of the support schemes (Sect. 3.3) feed the valuation, which hosts the strike solver (Sect. 3.4) and the carbon balance (Sect. 3.7).

Energy, loads, settlement and emissions thus derive from one hourly simulation. Sub-hourly shutdown and start-up transients are represented by composing each event hour from one transient load window plus trailing-state windows, so that lifetime damage and energy stay consistent with the hourly resolution. Relative to Gräfe et al. (2026a), the present study adds the capability contracts with the control-gated availability basis, the strike convention of anchoring at the reference dispatch, and the two-driver merit-order transform.




Figure 1. Integrated simulation pipeline.

3.1 Price futures: observed prices, merit-order transform, bootstrap

210 Price futures are constructed from observed data rather than from a stochastic price model. The observed day-ahead price p_t and the observed system wind and solar feed-in W_t and S_t , each normalised to installed capacity, share the calendar of the meteorological data, so the correlation between local wind and prices is measured, not modelled. Deeper negative-price futures are generated by a merit-order transform of the observed series,

$$p'_t = \max(p_t - \lambda D_t, p_{\text{floor}}), \quad D_t = \rho \frac{W_t}{\bar{W}} + (1 - \rho) \frac{S_t}{\bar{S}} - 1, \quad (1)$$

215 where $\bar{W}$ and $\bar{S}$ are the mean feed-ins over the path, D_t is the renewable surplus driver, ρ is the wind share of the variable renewable energy in the driver, λ is the additional price depression per unit of driver above its mean, and p_{floor} is the exchange price floor. Hours with mean driver at its mean keep their observed price. Hours of surplus are depressed, hours of scarcity lifted. Because D_t averages to zero, the transform preserves the mean price and changes only the shape of the price distribution: it shears the price–driver cloud about the point of mean driver and mean price. This is a convention that separates the shape from the level: the hours of surplus lose what the other hours gain (Sect. 6.4). For each target negative-price share, λ is solved by bisection. With $\rho = 1$ the driver is wind alone and the transform reduces to $p'_t = \max(p_t - \lambda_W(W_t - \bar{W}), p_{\text{floor}})$ with $\lambda_W = \lambda/\bar{W}$, the price depression per unit of wind feed-in above its mean. The driver mix ρ decides which hours become negative: with wind alone they are windy hours, with solar alone they are sunny hours. The futures thus impose most of the



price response to the drivers, while the observed path contributes its weather, its calendar and its residual price variability. The
 225 driver mix, the share target and the calibration of the futures are case settings (Sect. 4.1). All financial quantities are in real
 terms, and prices and strikes are not indexed to inflation.

Twenty-year hourly series are assembled by block-bootstrapping whole calendar years. Each drawn year brings its local
 wind, its system feed-in and its prices together, so the measured correlation between them is kept. The hourly prices of every
 drawn year are then shifted, so that the mean price of the year equals a common level. The transform of Eq. (1) is applied
 230 afterwards and preserves this level. Different seeds give replicate weather–price paths. The pool of years, the common price
 level and the number of replicates are case settings (Sect. 4.1).

3.2 Control: revenue-triggered shutdown

The control axis is a single farm-wide rule: all turbines shut down in hour t if the projected farm revenue falls below a threshold
 θ ,

$$235 \quad \hat{P}(v_t) p_t < \theta, \quad (2)$$

where $\hat{P}(v_t)$ is the farm power projected from the ambient wind speed via the turbine power curve, without wake correction.
 This is the information available to an operator at gate closure. Restart follows when the condition clears. Start-up and shutdown
 transients contribute their own load windows (Sect. 3).

The threshold $\theta = 0$ is the negative-revenue stop, the reference dispatch of the price-exposed schemes (Sect. 3.4). Larger
 240 thresholds curtail hours of positive but low revenue. Whether the fatigue relief and the lifetime they buy outweigh the forgone
 income is the question of Sects. 5.2–5.3. The rule is the farm’s operating rule for its whole life, the extension years included
 (Sect. 3.5). The sweep over thresholds is an experiment on that rule, not an optimisation of dispatch. The threshold values,
 their derivation and their calibration are case-specific and given in Sect. 4.2. A derating control arm is outside the scope of this
 study.

245 3.3 Support schemes as settlement functions

Support schemes are implemented as pure settlement functions. They map the price series, the delivered energy and an avail-
 ability basis to cash flows, and they do not alter dispatch. All physical curtailment is decided by the controller (Sect. 3.2). This
 keeps the scheme and control axes of the experiment apart. Any behavioural difference between schemes enters only through
 the operating point at which each scheme is evaluated (Sect. 3.4).

250 Let p_t denote the day-ahead price in settlement interval t (1 h), E_t^{del} the delivered farm energy, K the strike price, and $\bar{p}_{m(t)}$
 the time-weighted average price of the month containing t , which serves as the reference price for the CfD family. Revenue in
 interval t is the market leg plus a support leg on a scheme-specific basis B_t :

$$R_t = p_t E_t^{\text{del}} + (K - \bar{p}_{m(t)}) B_t, \quad (3)$$

with the following five configurations:



- 255 1. **Market:** no support leg ($B_t = 0$), the farm earns the market leg only.
2. **FiT:** $R_t = K E_t^{\text{del}}$. The tariff, the FiT's strike, replaces the market leg instead of adding a support leg to it. FiT support totals are therefore gross payments, not directly comparable with the schemes of Eq. (3).
3. **CfD:** two-sided generation-based contract, $B_t = E_t^{\text{del}} \cdot \mathbb{1}\{p_t \geq 0\}$. The indicator implements the negative-price rule of current European practice: no support is paid on energy delivered in hours with negative day-ahead price. The support leg pays the difference payment $K - \bar{p}_{m(t)}$ per unit of basis and is signed, so settlement periods with $\bar{p}_{m(t)} > K$ produce paybacks.
- 260 4. **Capability CfD:** $B_t = \hat{E}_t^{\text{av}}$, the effective available energy defined below. Support is paid on what the farm could have delivered, so operator curtailment does not forfeit support.
5. **Capability CfD, delivered-gated:** $B_t = E_t^{\text{del}} + (\hat{E}_t^{\text{av}} - E_t^{\text{del}}) \cdot \mathbb{1}\{p_t < 0\}$. The availability basis applies in negative-price hours only. Curtailment at non-negative prices is settled on delivered energy. The two capability variants differ in this one clause and are otherwise identical. We denote them as the ungated and the gated capability contract in the following.
- 265

The available energy E_t^{av} is the production of the same farm under the same weather without price-driven control, taken from the uncontrolled baseline simulation including wake losses. The effective basis credits availability deficits only when they are caused by the controller:

270

$$\hat{E}_t^{\text{av}} = E_t^{\text{del}} + (E_t^{\text{av}} - E_t^{\text{del}})^+ \cdot \mathbb{1}\{\text{shutdown commanded in } t\}. \quad (4)$$

This control gate mirrors settlement practice, in which availability deficits count only during logged downregulation commands. Deficits caused by technical unavailability, such as storm cut-outs and startup transients, are never compensated, and numerical disagreement between the baseline and scenario simulations cannot enter the settlement. A relative dead band of 2 % of the interval's available energy is retained as a secondary numerical guard. The gated quantity is hereafter simply the available energy. Throughout the paper the support leg of the capability contracts is split by its basis into support on delivered energy and support on curtailed availability, at negative and at non-negative prices. The delivered-energy gate of the gated variant removes the part at non-negative prices. The lifetime totals of the two legs are the market revenue and the support.

275

The capability contracts take from Schlecht et al. (2024) the principle that support does not depend on the operator's own output. Their basis is the farm's own available energy, which Schlecht et al. (2024) avoid because it has to be estimated at the site. They follow the companion study of contract design (Ioannou et al., 2026), which examines the settlement basis and the relation between support and strike for the same contracts. Both variants are stylised. Curtailment instructed by the system operator is not modelled, because settling it requires an estimate of the production foregone, as the curtailment compensation of the British contract does (Low Carbon Contracts Company and Department for Energy Security and Net Zero (UK), 2024).

280 That estimate lies outside this study.

285



Support is paid for the nominal life only, as in contracts written today. The support window closes at the nominal end, and the extension years of Sect. 3.5 earn market revenue under the same operating rule. Support that follows the extension years is discussed as a sensitivity (Sect. 6.2).

3.4 Strike determination: anchoring at the reference dispatch

290 Comparing support schemes requires a strike price for each. We solve a separate required strike $K^*(s, r)$ for every scheme s and price future r : the strike at which the project's net present value is zero on the reference weather–price path,

$$\text{NPV}(s, r, K^*; \theta_s^{\text{ref}}) = 0, \quad (5)$$

evaluated at the scheme's *reference dispatch* θ_s^{ref} . The reference dispatch is the negative-revenue stop of Sect. 3.2 applied to the revenue the scheme itself settles. Under the CfD, support is suspended in negative-price hours. Under both capability contracts, 295 support continues on available energy and does not depend on whether the operator delivers. In all three cases producing in these hours adds only the negative market revenue, so all three schemes induce the negative-revenue stop ($\theta = 0$). Under the FiT the settled revenue is never negative, the stop never triggers, and the reference dispatch is continuous production, the uncontrolled baseline. Market carries no strike and is evaluated at the stop. A scheme evaluated at its reference dispatch with the strike so solved is its *anchor*, and value results are reported relative to it. The NPV is computed in real terms from unlevered 300 free cash flow with a loss-carry-forward tax ledger at a corporate tax rate of 22 %, discounted at a fixed real WACC of 5.4 % for all schemes. Decommissioning is paid at the end of operation and moves with the extension. The fixed discount rate closes the financing-risk channel, under which lower cash-flow variance would earn cheaper capital. The strike is found by bisection over the full settlement and valuation chain on fixed physics, to a tolerance of 10^4 € in NPV.

The strike absorbs the cash flows of the dispatch it is solved under. A capability contract priced under continuous production, 305 for example, would carry the losses of producing through negative prices in its strike, although an operator under that contract stops. Solving every strike at the scheme's reference dispatch ties each strike to the operation its rule implies. Because that dispatch is the same negative-revenue stop for all price-exposed schemes, their strikes are solved on identical physics and differ through the settlement functions alone. First, the present value of support at the anchor is the same for all price-exposed schemes. With identical physics and $\text{NPV} = 0$ it equals total costs minus market revenue, the support requirement. Scheme 310 choice therefore changes the composition, timing and risk of the support, not its amount. Second, evaluating each scheme under the complementary dispatch, continuous production for the price-exposed schemes and the stop for the FiT, with its strike frozen, compares the two dispatches under one contract. It shows whether the reference dispatch is worth more to the operator than the complementary one, and it prices what a scheme that induced the complementary dispatch would have to cover.

315 Strikes are solved once, on the reference path, and frozen across the replicate paths, consistent with contracts priced ex ante. The replicate paths show whether the results depend on the drawn path. All results are reported on the reference path (Sect. 5).



3.5 Fatigue accumulation and lifetime: the aggregation chain

Fatigue information exists at four levels: load channel, component, turbine and farm. The O&M model reads it as one hazard modifier per component and the lifetime decision as one scalar per farm, and each averaging step on the way can hide a
 320 channel that responds differently from its neighbours. The chain from the channel series to these two consumers follows the normalisation and aggregation conventions of the farm-response framework of Pettas et al. (2026). We state them first, then the lifetime decision and the worst-channel check that bypasses the averages.

3.5.1 Channel level: the load-surrogate pipeline.

The primary fatigue data are produced per turbine $i \in \{1, \dots, 48\}$ and hour by the load surrogates of the farm-response frame-
 325 work, from the turbine's waked inflow and power setpoint. Each operating mode is served by a surrogate bundle trained on aeroelastic simulations of the corresponding IEC design load case. The modes are power production and derated operation, start-up, shutdown and idling. Transient events are therefore represented by their own trained models and not by scaling. The training method is documented by Vad et al. (2026). The output per interval is a damage-equivalent load for each of the nine channels of Table 1. Channel-specific Wöhler exponents and reference cycle counts convert it into a damage increment,
 330 and Palmgren–Miner accumulation yields the cumulative damage $D_{i,ch}(t)$ of turbine i in channel ch . Everything below is a functional of these series.

Table 1. Exercised load channels, their component mapping and the weights of the component index (Eq. 7, applied to baseline-normalised channel ratios, with the weights of Pettas et al. (2026)). All channels are covered by all four operating-mode bundles.

Component	Channel	Physical quantity	w
Rotor blades	wrot_BllRad0FlpMnt	blade-root flapwise moment	0.5
	wrot_BllRad0EdgMnt	blade-root edgewise moment	0.3
	wrot_BllRad0TorMnt	blade-root pitching moment	0.2
Drive train	wtrm_LoSdpShftRotXMnt	shaft torsion	0.7
	wtrm_LoSdpShftStaProjYZMnt	shaft combined bending moment	0.3
Tower	wtow_H0FAMnt	tower-base fore–aft moment	0.5
	wtow_H0SSMnt	tower-base side–side moment	0.3
	wtow_H0TorMnt	tower-base torsional moment	0.2
Yaw system	wtow_H100ProjFASSMnt	tower-top combined moment	1.0

3.5.2 From damage-equivalent loads to normalised damage

Damage values need a common reference before they can be aggregated. The Palmgren–Miner increment of an interval scales
 as $d_{i,ch} \propto n \cdot \text{DEL}_{i,ch}^{m_{ch}}$, with channel-specific Wöhler exponents and reference cycle counts that are not calibrated to the as-built
 335 structures. Absolute damage values are therefore arbitrary, and across Wöhler classes they differ by orders of magnitude. Only



ratios of like quantities are meaningful. The reference is always the same quantity under the uncontrolled baseline, so every reported fatigue result is a ratio that equals 1 under no control.

Ratios are formed at the channel level first, following the two-step aggregation of Pettas et al. (2026). Each channel series is divided by its own baseline end-of-life value, for the same turbine and channel,

$$340 \quad R_{i,ch}(t) = \frac{D_{i,ch}(t)}{D_{i,ch}^{\text{base}}(T)}. \quad (6)$$

The component index is the weighted sum of these ratios over the channels of the component,

$$R_{i,c}(t) = \sum_{ch \in c} w_{c,ch} R_{i,ch}(t), \quad \sum_{ch \in c} w_{c,ch} = 1, \quad (7)$$

with the weights of Table 1. A raw-weighted index would track a single channel, because of the difference in magnitude between Wöhler classes. The component index is a modelling convention that provides a scalar per-component state. It is not
 345 the damage of any physical location, since the channels of a component load different hot spots. Raw damages are added only over turbines within one channel, for the worst-channel check of Sect. 3.5.3.

3.5.3 Lifetime extension

The farm health at nominal end-of-life T is the equal-weight mean of the normalised consumptions over the N_i turbines and the N_c structural components per turbine, rotor blades, tower and yaw system,

$$350 \quad H_{\text{farm}}(T) = \frac{1}{N_i N_c} \sum_i \sum_c R_{i,c}(T), \quad (8)$$

equal to 1 for the baseline by construction. The drive train is replaceable machinery and enters the O&M route only. The fatigue relief of an operation is $1 - H_{\text{farm}}(T)$, the share of the nominal-life fatigue budget, the baseline's end-of-life damage, that it leaves unused. Per component and channel the relief is $1 - R$. The extension continues operation at the same damage rate until the budget is used,

$$355 \quad L = \frac{1}{H_{\text{farm}}(T)}, \quad T_{\text{ext}} = T(L - 1), \quad (9)$$

the farm-mean rule, under which a farm that consumed 80 % of its budget at T operates $0.25T$ longer. Extension operation is simulated by replaying the operating pattern deterministically, so the budget is exactly used at the end of the extension years. No cost is charged for the assessment and refurbishment a real extension requires, so the value of the extension years (Sect. 5.3) is the margin available for them. The extension so defined is the fatigue-permitted one: it isolates the limit state that
 360 the operating rule changes, while corrosion, erosion, ageing of the electrical system, spare-part supply and the terms of permits and grid connection limit a real extension as well and do not respond to curtailment, so the fatigue-permitted extension bounds the realisable one from above.

The chain feeding Eq. (8) averages twice: the weighted channel mean inside each component (Eq. 7) and the equal-weight component mean inside the farm health. Averages do not resolve heterogeneous responses within a junction: a channel with



365 little relief, or with damage growth, is concealed by a strongly relieved partner, and a highly consumed component by relieved neighbours. Lifetime-extension standards, by contrast, assess each failure mode of each component (Sect. 2.2). As a convention of this study, not a certification assessment, we therefore carry a worst-channel check alongside the farm health,

$$L_{\text{worst}} = \frac{1}{\max_{ch} R_{ch}}, \quad R_{ch} = \frac{\sum_i D_{i,ch}(T)}{\sum_i D_{i,ch}^{\text{base}}(T)}, \quad (10)$$

over the channels of the structural components, and read every average-based extension as an upper bound on what an assess-
 370 ment of each failure mode would allow.

3.6 Operation and maintenance: the load-to-cost bridge

The O&M model converts failure rates into expected costs, analytically. Each component carries a set of failure modes (minor repair, major repair, replacement) specified by published offshore reliability data: Weibull-derived base occurrence rates, repair times, spares costs, labour, and the service vessel each mode requires.

375 The bridge from loads to costs is the wear-out hazard coupling. The model consumes fatigue per component, at farm level, as a time series: the farm-mean consumption ratio of Sect. 3.5.2, $\tilde{D}_c(t) = \frac{1}{N_i} \sum_i R_{i,c}(t)$, which equals 1 at the baseline's end of life. For each maintenance window W , the replacement (wear-out) failure rate of component c is scaled by the accumulated consumption,

$$\lambda_{\text{eff}}(c, W) = \lambda_{\text{base}}(c) \left(1 + \alpha \langle \tilde{D}_c \rangle_W \right), \quad (11)$$

380 with $\langle \cdot \rangle_W$ the time average over the window, α the coupling sensitivity, and the multiplier clamped to a plausible band. Minor- and major-repair rates do not depend on damage. The linear coupling is a stylised representation of how wear-out rates depend on accumulated fatigue. The sensitivity α is not calibrated (Sect. 6.4). The fatigue relief of the control lowers the repair intensity more and more over the life. The end-of-life ratios $R_c = \tilde{D}_c(T)$ are reported per component with the results.

Each component and failure mode is a small continuous-time Markov chain over operation, waiting for logistics, waiting
 385 for a weather window and repair, solved for its steady-state occupation per maintenance window. Its outputs are expected failure counts, downtime and component availability, expectations without sampled events and deterministic given the inputs. They are priced with spares, labour and vessel costs into a monthly payment series, together with the preventive-maintenance schedule and fixed O&M. The technical availability of the farm (approximately 95 %) scales the settled energy (Sect. 3.3). The extension years use the same failure rates, repair durations and access conditions as the nominal life. Ageing enters only
 390 through the accumulated-damage hazard of Eq. (11), which makes the O&M of the extension years slightly more expensive per year.

The same expected event counts carry the O&M emissions of Sect. 3.7. They are summed into a monthly series next to the cost series, separately for the nominal and the extension years.



3.7 Environmental assessment

395 In addition to the economic assessment, the impact of the market scenarios, support schemes and revenue-triggered shutdowns is evaluated from an environmental perspective. Following the metric definitions of Canet et al. (2023), three carbon-related indicators are used: environmental cost, environmental value and net environmental value. The environmental cost, also known as the carbon footprint, reflects the life-cycle greenhouse gas (GHG) emissions attributable to the wind farm, including emissions from its manufacturing, construction, maintenance, and end-of-life treatment. The environmental value reflects the GHG emissions displaced in the connected electricity system by the energy the wind farm generates. The net environmental value is the difference of the two, with a positive value indicating that the wind farm saves emissions in net terms. Each indicator is evaluated both as a total over the realised lifetime (ktCO₂e) and per unit of electricity delivered (kgCO₂e/MWh). The totals are referred to as total life-cycle emissions, total displacement and net displacement, while the energy-normalised quantities retain the terms environmental cost, environmental value and net environmental value.

405 The life-cycle emissions that do not depend on the operating strategy are quantified with the DETECT framework (Kainz et al., 2024), a parameterised toolchain that derives the life-cycle inventory and the resulting GHG emissions of an offshore wind farm configuration. These emissions cover the life cycle from raw material extraction through manufacturing and installation to decommissioning and end-of-life treatment. O&M emissions are assessed separately because the shutdown strategies change them. They follow the failure-rate-based method of Gräfe et al. (2025): component failure rates are translated into O&M emissions through the vessel fuel consumed in transit and on site, the production and maintenance emissions of the vessels allocated to their lifetime usage, and the embodied emissions of the spare parts replaced. The embodied emissions encompass all upstream emissions from raw material extraction through processing, manufacturing, and transportation to the installation port. A limitation of this study is that emissions arising from decommissioning and end-of-life treatment related to spare parts are disregarded, as their calculation would require a full software coupling with the DETECT framework, which is left for future work.

An operating strategy enters the environmental cost in two ways: through the lifetime electricity production that normalises it, and through the component loads, whose failure rates set the O&M activities and their emissions (Sect. 3.6). For each strategy, the O&M emissions are added to the strategy-independent life-cycle emissions and divided by the electricity delivered over the realised life.

420 The environmental value is quantified with MDFs, the GHG emissions displaced in the electricity system by a marginal increase of generation from a given technology. The MDF model of Kainz et al. (2025) combines an ANN that predicts the emissions of the electricity system with a mathematical model that derives technology-specific MDFs from the trained network. Adapted to the electricity system of this study (Sect. 4.1.2), it gives an hourly MDF series for wind energy. The displaced emissions of each hour are the wind farm's generation multiplied by the MDF of that hour, integrated over the operational life. Because the MDF varies strongly over time, the environmental value depends not only on how much electricity is produced but also on when it is produced. Strategies with different temporal generation profiles differ in environmental value even at similar total production.



Market scenarios and support schemes affect the environmental assessment through the operating strategies they incentivise, which change the amount and timing of generation and the O&M activities. The assessment therefore records how economic rules translate into operating behaviour and, through it, into environmental performance. For the support schemes, the environmental benefit is also related to the public support required, giving the benefit per unit of support.

4 Case study and design of experiments

4.1 Site, data and assumptions

The case is the Lillgrund offshore wind farm in the Øresund, with 48 Siemens SWT-2.3-93 turbines and 110.4 MW. Its layout is one of the densest in operation, with a minimum spacing of 3.3 rotor diameters, so wake interaction is strong and results on wake-related quantities are at the upper end of the plausible range.

4.1.1 Meteorological and market data.

Meteorological input is an ERA5-based hourly series of wind speed, wind direction and turbulence intensity at the site for 2015–2025. Market input is the ENTSO-E day-ahead price and the wind and solar feed-in of the DK2 bidding zone for the same hours. Lillgrund settles in the Swedish SE4 zone. We use DK2 as the market environment because of its data coverage and its strong interconnection with SE4 across the Øresund.

4.1.2 Emission data.

The life-cycle emissions of the plant, excluding O&M and therefore independent of the operating strategy, are derived from the DETECT inventory adapted to the Lillgrund farm (Kainz et al., 2024), amounting to approximately 140 kt CO₂e. O&M emissions are computed from the event counts of the O&M model with the vessel and spare-part factors of Gräfe et al. (2025) (Sect. 3.7). The marginal displacement factors come from the MDF model of Kainz et al. (2025) adapted to the Danish electricity system of 2017–2020 via hourly generation, demand and exchange data (Energinet, 2026), with technology-specific emission factors from Ecoinvent (Wernet et al., 2016) and emission factors for electricity imports from Electricity Maps (2024). The mean derived from an ensemble of ANNs trained on different data splits gives the hourly series. The factor averages 252 kgCO₂e/MWh over the four years and varies strongly with the generation mix, with values ranging from –588.4 to 1336.1 kgCO₂e/MWh. Negative values, where additional wind generation raises the system emissions, occur in less than 2% of the hours. Figure 3a shows the hourly factor over ten days in June 2017, where it falls during the wind passages and peaks in the calm spell. The production-weighted factor of the farm is about 230 kgCO₂e/MWh.

4.1.3 Path construction.

The twenty-year paths of Sect. 3.1 are drawn from the pre-crisis years 2015–2020. The years 2021–2023 are excluded because the gas-price crisis is a structural break in price formation that the merit-order transform is not intended to represent. The years



2024–2025 are not drawn either, so that all drawn years share one price formation. Each drawn year is shifted to the common real mean of 60 € MWh^{-1} , so paths differ in weather and price shape but not in price level. Five seeds give five replicate paths. One of them is the reference path, on which the strikes are solved and the figures are drawn (Sect. 4.3).

460 The hourly MDFs cover 2017–2020 and follow the same year draw. Drawn years outside this coverage, 2015 and 2016, take a factor conditioned on the feed-in decile and the month, calibrated on the 2019–2021 overlap of the displacement-factor and feed-in series. These filled years carry the conditional mean without its hourly variation. The factors are held fixed across the price futures. The futures move the hours of surplus in the prices, not the marginal mix of the electricity system.

4.1.4 Price futures.

465 Figure 2 shows, for six bidding zones, the wind and solar feed-in in negative-price hours relative to their annual means. In 2019 and 2020 the negative hours of the Danish and German zones were windy hours with little sun. By 2025 they are sunny hours in every zone, with wind near or below its annual mean. The futures of this study span that range. We construct three futures that share one negative-price share of 10 % and differ in the driver of Eq. (1): wind alone ($\rho = 1$), a mix ($\rho = 0.6$) or solar alone ($\rho = 0$). With the historical path they form the future axis of the paper, and the three transformed futures are the
 470 stressed futures. On the wind axis the zone-years of 2019–2020 lie between the solar-driven and the wind-driven future, and those of 2024–2025 near or below the solar-driven future.




Figure 2. Wind and solar feed-in in negative-price hours as a multiple of the annual mean, six bidding zones 2019–2025 (2021–2023 excluded), and the three futures of the study at the 10 % negative-price share. Zone-years with fewer than 0.5 % negative hours are omitted; dotted: no relation.

The share of 10 % lies above the 3 to 7 % observed in 2025 and below the deep-penetration projections. It is the level at which the three drivers are compared at an equal number of negative hours. The historical path keeps the observed prices. The price floor is the exchange minimum of -500 € MWh^{-1} . Figure 3 shows the transform on the observed 2015–2020 series: a



ten-day window with a calm spell, a storm passage and clear days (a), the price–feed-in cloud with the fitted slope of every future (b), and the calibration of λ (c). An hour without the driver gains the full slope, up to a third of the mean at night under the solar-driven future and up to two thirds of it in calm hours under the wind-driven one. The fitted slope is steepest for the wind-driven future and flat for the solar-driven one, whose price no longer responds to wind. The solar driver saturates near a quarter of the hours, because only daylight hours can be pushed below zero. Table 2 calibrates the four futures on the reference path. The solar-driven future needs the smallest slope because its driver is concentrated in daylight hours. Its negative hours are daylight hours with wind near its annual mean, whereas half of the negative hours of the wind-driven future fall at night.




Figure 3. The transform on the observed 2015–2020 series, each year re-levelled to 60 € MWh^{-1} , and the displacement factor of the same hours. (a) Wind and solar feed-in (top), hourly prices (middle) and the hourly displacement factor (bottom), 12–21 June 2017. Prices: observed, and the three futures; solid black: zero; dashed: the common mean. Displacement factor: the series of the Danish system, the same in every future. (b) Price vs system wind feed-in over the full pool (observed: grey, wind-driven future: blue; thinned) with the fitted slope of each future. (c) Share of negative-price hours as a function of λ per driver; dotted: the 10 % target; markers: the solved λ .



Table 2. Calibration of the future axis on the reference path: transform slope λ per unit of driver (for the wind-driven future λ_W per unit of wind feed-in, Eq. 1), achieved negative-price share, mean price, correlation between price and wind feed-in, wind and solar feed-in in negative-price hours relative to their annual means, and the share of negative hours in daylight.

	hist.	wind-driven	mixed	solar-driven
driver ρ	–	1	0.6	0
λ [€MWh ⁻¹]	–	$\lambda_W = 129.0$	51.0	21.4
achieved share [%]	0.04	10.00	10.00	10.00
mean price [€MWh ⁻¹]	60.00	60.00	60.00	60.00
corr(price, wind feed-in)	–0.35	–0.93	–0.59	–0.03
wind feed-in in negative hours [$\times$ mean]	1.55	2.49	1.53	0.86
solar feed-in in negative hours [$\times$ mean]	–	0.68	4.04	4.89
daylight share of negative hours	–	0.47	0.94	1.00

4.1.5 Engineering and cost data.

The load surrogates of Sect. 3.5 are trained for the SWT-2.3-93. Failure rates of the O&M model follow the offshore reliability data set of Gräfe et al. (2025). Repair times, material costs and logistics parameters are scaled from published offshore O&M data to the 2.3 MW class. Capital and operating costs are representative of the nearshore 2–3 MW class and not calibrated to the asset, so absolute financial levels are stylised. Proportional cost errors leave the orderings, signs and mechanisms of the reported differences intact but scale their magnitudes. Every required strike fills the financing gap that the cost level sets, so the strike spreads of Sect. 5.3 are specific to this cost basis.

4.2 Threshold grid

The revenue thresholds are constructed, not optimised. The reference scale is the farm’s rated-power revenue at the mean price,

$$\Theta = P_{\text{rated}} \bar{p} = 110.4 \text{ MW} \times 60 \text{ €MWh}^{-1} = 6624 \text{ €h}^{-1}, \quad (12)$$

and the grid comprises $\theta = 0$ (the negative-revenue stop) and $\theta = f \Theta$ with $f \in \{3, 6, 9, 12, 15, 18\}\%$, i.e.

$$\{199, 397, 596, 795, 994, 1192\} \text{ €h}^{-1}.$$

The dimensionless form makes the grid portable to other farms and price levels. The fractions were calibrated on the historical path so that the induced farm idle time rises densely from the stop, about 18 % of hours with weather idling included, to a deliberate overshoot of plausible operation, about 35 % on the historical path and more in the stressed futures. Two design choices matter for interpretation. First, the threshold is defined on projected farm revenue (€h⁻¹), not on price (€MWh⁻¹). The rule curtails hours by the value at stake, so a low-wind hour is curtailed at a higher price than a high-wind hour. Second, the



500 grid is held fixed in absolute terms across futures. Because all futures share the same mean price and the same farm, a given θ represents the same value at stake everywhere, and the fraction of hours it triggers becomes a future-dependent outcome rather than a controlled input.

4.3 Design of experiments

505 The experiment varies four dimensions: the support scheme (five settlement functions, Sect. 3.3), the future (historical, wind-driven, mixed and solar-driven, Sect. 4.1), the shutdown threshold (Sect. 4.2) and the weather–price path (five replicates, Sect. 4.1). Not every combination is run. Scenarios fall into three classes, defined by the dispatch they evaluate (Table 3).

Anchor scenarios evaluate every scheme at its reference dispatch in every future (Sect. 3.4). The required strikes are solved here, and all value results are reported relative to these scenarios. *Complementary-dispatch scenarios* evaluate the same contracts, strikes frozen, under the other dispatch. The uncontrolled run of this class is also the physical reference of the whole experiment. The *threshold sweep* evaluates all five schemes on the threshold grid beyond the negative-revenue stop in the three stressed futures.

515 The three classes comprise 130 scenarios per path, 650 in total. Two sensitivities run on the reference path only. The first repeats the wind-driven future at negative-price shares of 5 and 20 %, with all three classes, and is reported in Sect. 5.4. The second is an anchor run without lifetime extension, whose re-solved strikes are reported with the strike results (Sect. 5.3). The worst-channel check of Sect. 3.5.3 is evaluated on the same simulations as the farm-mean rule, for comparison, and shown with the lifetime results (Sect. 5.2). Three invariants of the design are checked on every campaign and reported with the results. The FiT strike is identical across futures, because its settlement does not depend on prices. The two capability variants coincide at the negative-revenue stop, where all curtailment occurs at negative prices and the delivered-energy gate is inactive. Energy, fatigue and extension are identical across schemes for a given dispatch, so settlement does not enter the physics.



Table 3. Scenario classes. Anchor scenarios solve the required strike K^* (Eq. 5) on the reference weather–price path. All other scenarios evaluate the frozen strikes. All classes are repeated on five replicate paths.

Class	Dispatch	Role	Futures	Scenarios/path
Anchors (reference dispatch)	negative-revenue stop ($\theta = 0$) for Market, CfD and both capability variants; continuous production for FiT	defines the contracts ($K^* \rightarrow \text{NPV} = 0$); origin of the value axis	4	20
Complementary dispatch	continuous production for the price-exposed schemes; negative-revenue stop for FiT	value of the reference dispatch under a frozen contract; the uncontrolled run is also the physical reference (fatigue normalisation, availability basis)	4	20
Threshold sweep	$\theta \in \{199, \dots, 1192\} \text{ € h}^{-1}$, all five schemes	value and physics of curtailment beyond the negative-revenue stop	3	90

Market carries no strike and participates in the anchor class as the unsupported reference. The historical path has no sweep. The share sensitivity and the run without lifetime extension use the reference path only.

520 5 Results

5.1 Market revenue: delivered energy and capture price

Figure 4 relates the energy delivered in the nominal 20 years to its capture price, the production-weighted average day-ahead price of that energy over the same period, for every threshold and future. Their product is the undiscounted market revenue of the nominal 20 years, so the grey contours, lines of constant product, are lines of equal market revenue.

525 The uncontrolled farm delivers the same energy in every future, because the physics does not see prices, so its points stack vertically. Its capture price is set by the driver of the future, not by the number of negative hours, which is the same in the three stressed futures. The transform moves value out of the hours in which the driver is in surplus. Under wind-driven prices these are the farm’s own producing hours, and the farm captures about half of the mean price. Under solar-driven prices they are hours in which the farm sells little, while the lifted night hours are the ones it produces in. The farm then captures the mean
 530 price, at the fixed mean price of the futures slightly more than on the historical path. The mixed future lies between.

The negative-revenue stop moves the farm across the contours, to less energy and more revenue, in every future. The energy given up and the revenue gained depend on the future. Under wind-driven prices the stop removes the windiest hours, about a quarter of the energy, and the remaining production still sits in the cheaper positive hours, so the capture price stays below the mean. Under solar-driven prices the stop removes midday hours in which the farm produces little, and the remaining production
 535 is priced above the mean, because it falls in the hours solar does not serve. Producing less earns more, but the gain shrinks along the axis.



Beyond the negative-revenue stop each further step removes hours of positive but low revenue. Revenue falls, little at first and faster with the threshold, because each step removes hours of higher revenue than the last. The capture price rises, because the removed hours carry far less revenue per megawatt-hour than the farm captures on average. Under wind-driven prices the sweep from the stop to the deepest threshold removes about a quarter of the remaining energy and about 10 % of the revenue. The three futures form separate curves, ordered by the wind coupling of their negative prices.

The higher capture price does not recover the revenue. Under wind-driven and mixed prices the farm earns less in the market over the nominal life than on the historical path, under solar-driven prices more. The negative-revenue stop recovers only the negative revenue. No dispatch can earn back what the historical path earned in the hours now priced at or below zero, and the remaining hours sell for less than before. After the stop the farm delivers less energy at a higher capture price, and its market revenue still falls short of the historical path. The financing gap, the costs that the market revenue does not cover, is therefore wider than on the historical path under wind-driven and mixed prices and narrower under solar-driven prices.




Figure 4. Capture price against delivered energy per future, both over the nominal 20 years (energy settled, reference path). Grey: lines of equal undiscounted market revenue over the same period, in M€, the product of the two axes. Open squares: uncontrolled baseline; dashed: activation of the negative-revenue stop; labels: threshold in € h^{-1} .

5.2 Physical impact of curtailment: fatigue relief and lifetime

Figure 5 shows the relation between prices and fatigue on which this section builds. For five channels of Table 1, one per response family, it plots the cumulative share of the channel's damage held by the hours below a price cut, next to the share of delivered energy, per future. Under wind-driven prices the channels separate. The blade-flap curve lies far above the energy curve, the tower-base side-side curve well above it, the blade-edge and tower-base fore-aft curves a little above it, and the tower-top curve below it. The hours priced below zero hold most of the blade-flap damage but only about a quarter of the energy, so the negative-revenue stop removes more damage than energy for four channels and less for the tower-top channel. This order follows from the wind speed of the removed hours, which are the windiest hours of this future, and from how each channel's damage grows with wind speed, a property of the load surrogates. Blade-flap damage grows steeply with wind speed up to cut-out, and tower-base side-side damage keeps growing to cut-out as well. Blade-edge and tower-base fore-aft damage



grow up to rated wind speed and little beyond it, where the controller limits thrust. The tower-top moment is largest in partial load, so its damage sits in hours of moderate wind. Under mixed prices the curves keep this order with smaller gaps, because
 560 only part of the negative-price hours are windy hours. Under solar-driven prices the negative-price hours are midday hours with average wind. Every channel's curve lies on the energy curve, the stop removes damage in proportion to energy, and the price carries no information about loads.




Figure 5. Cumulative share of a channel's damage and of delivered energy (grey) in the hours below a price cut, per future. Channels: blade-root flapwise and edgewise, tower-base fore–aft and side–side, tower-top combined moment. Damage: farm mean of the hourly damage-equivalent load over the 48 turbines, raised to the channel's Wöhler exponent; price: hourly farm revenue per megawatt-hour; producing hours of the uncontrolled baseline, reference path. Vertical line: the negative-revenue stop.

Figure 6a aggregates this to the component indices used by the O&M and lifetime models (Sect. 3.5). Every component is relieved at every threshold in every stressed future, and the relief grows along the sweep, but its size and its order are set by the
 565 future. Under wind-driven prices the blades are relieved most, then the drive train, the tower and the yaw system. Under solar-driven prices the relief is small and the blades are relieved least, because their damage sits in the windiest hours, which a stop at midday does not touch. Three mechanisms limit the relief. Channels that accumulate damage while parked, the tower-base side–side channel above all, keep a large share of it, because curtailment converts producing hours into idle hours rather than removing them. Start-up transients keep a part of the damage the negative-revenue stop removes from the producing hours,
 570 about 20 % of the blade-flap baseline under wind-driven prices, independent of the threshold. Channels whose damage sits in partial-load hours, the tower-top channel of the yaw system, are little affected by a stop that removes windy hours. A farm-level average hides all three. The worst-channel check reports them (Appendix A).

Figure 6b converts farm health into lifetime. The farm-mean rule yields about ten years at the negative-revenue stop under wind-driven prices, three under mixed and one under solar-driven prices, and more along the sweep. The worst-channel check
 575 confirms an extension wherever the least-relieved channel is relieved. Under wind-driven and mixed prices it yields roughly half of the farm-mean years. At the stop under solar-driven prices it yields nothing, because the start-up transients return what the stop removes from the blade-flap channel, which ends at its baseline budget (Appendix A).



The damage-coupled O&M model (Sect. 3.6) turns the fatigue relief into a small reduction of the nominal-life O&M cost and a larger O&M cost during the extension years, which Sect. 5.3 nets against the extension revenue. The O&M response is small against the market and settlement responses.

Figure 7 compares curtailment strategies by their revenue cost and their fatigue benefit. It plots the cumulative revenue given up, ρ , against the cumulative fatigue consumption avoided, δ . Each producing hour of the uncontrolled farm consumes an increment of the farm health of Eq. (8), and the hours of the nominal life sum to one. δ counts the baseline increments of the removed hours, so it exceeds the realised relief of Fig. 6, because the idle hours add damage. The solid curves remove hours in the order of a falling revenue threshold. The dotted curves are the upper bound for any selection of hours beyond the negative-revenue stop.




Figure 6. Component damage and lifetime along the threshold sweep (reference path). (a) End-of-life damage of the component index (Eq. 7) relative to the uncontrolled baseline, one panel per stressed future, one line per component (marker and line style); $x = 0$ is the negative-revenue stop; shared axes. (b) Farm lifetime ratio L from the farm mean (solid) and from the worst channel, $1 / \max_{ch} R_{ch}$ (dotted); right axis: extension years. Channel detail in Appendix A.

Up to the negative-revenue stop no load information is needed. Every negative-revenue hour raises revenue and avoids fatigue, so the solid curves run left and upward from the origin, and the dotted curves coincide with them. The leftmost point of each curve is the stop, and every selection that removes all negative-revenue hours ends there, whatever its order. Beyond the stop the curves turn right, because every further hour gives up revenue, and the order of removal matters. An hour of low positive revenue is either a calm hour, which carries little fatigue at any price, or a windy hour at a price just above zero, which carries much. Revenue does not tell them apart. The solid curves, which remove hours in revenue order, therefore avoid fatigue only in proportion to revenue and straighten towards the unselective diagonal in every future. Under solar-driven prices, as on the historical path, they run along the diagonal from the start. The dotted curves, which remove the windy low-price hours first, rise steeply from the stop, and the gap between the two is what load information would add. The revenue signal is an exact proxy for a fatigue-efficient selection up to the stop and a proxy of declining quality beyond it. An optimal curtailment beyond the stop requires load information.



The overlaid points show the realised value of each simulated threshold under Market settlement. On the historical path the negative-revenue stop and the uncontrolled baseline lie within a million euros of each other, because few hours are priced negative. Under mixed prices the optimum lies a few steps beyond the stop, where the value of the additional extension exceeds the small revenue loss of curtailing further. Under wind-driven and solar-driven prices it lies at the stop or one step beyond it, within a million euros of the stop's value. This is within the variation between the replicate paths.

Figure 8 collects the structural response of every operating rule against the wind coupling of negative prices with the observed bidding zones on the same axis. Relief and extension rise with the coupling at every threshold, and the gain from a deeper threshold rises with it as well. A deeper threshold does not make up for a weak coupling. The structural value of price-responsive operation is a property of wind-driven price formation.




Figure 7. Cumulative revenue given up ρ (relative to gross positive revenue) vs cumulative fatigue consumption avoided δ (farm-health increments, Eq. 8), per future. Solid: hours removed in revenue order; dotted: upper bound beyond the negative-revenue stop (all negative-revenue hours, then the remaining hours in order of fatigue avoided per unit revenue); grey: unselective. Points: simulated thresholds coloured by realised ΔNPV (Market); star: best realised threshold. $\rho < 0$: curtailment increases revenue.




Figure 8. Relief at the nominal end (a) and lifetime extension (b) against the wind feed-in in negative-price hours of the three stressed futures, one line per threshold (light to dark), reference path. Below (a): the observed bidding zones on the same statistic, from their first observed year (2019 or 2020, ES 2024; hollow marker) to 2025 (filled marker).

5.3 Contracts decide what flexibility is worth

5.3.1 Required strikes.

Figure 9 shows the required strikes $K^*(s, r)$ of Eq. (5), one per scheme and future. On the historical path the four schemes
 610 with a strike require nearly the same strike, between 95 and 100 € MWh⁻¹. Under recent price shapes the choice of scheme
 has little effect on the strike. The futures separate the schemes, and the separation follows the wind coupling of the future,
 not its negative-price share, which is the same in all three. The CfD pays on delivered energy in non-negative-price hours.
 Under wind-driven prices the negative-revenue stop removes a quarter of the energy and the rest captures half the mean price
 (Sect. 5.1). A large financing gap is then paid on a small basis, and the strike rises far above its historical level. Under solar-
 615 driven prices the stop removes little energy and the farm captures more than the mean, so the gap is small and the strike stays at
 its historical level. The capability contracts pay on available energy at the stop, where the two variants coincide, and this basis
 does not shrink with the stop. Their strike stays near its historical level under wind-driven prices and falls below the FiT under
 solar-driven prices, where the farm earns more in the market than on the historical path at the fixed mean price of the futures.



The FiT strike does not move, because its settlement does not depend on prices. The spread between the CfD and the capability
 620 strike thus arises between contracts that deliver the same zero-NPV project on the same physics, so it is the settlement function
 alone. The two bases differ by the energy in the negative-price hours. The wind coupling sets how much of the farm's energy
 falls into these hours and how large the gap paid on the bases is, so the spread grows with the coupling.

The hollow markers are the strikes re-solved for a farm decommissioned at the nominal end. The extension years earn
 market revenue and cost O&M, and their net income closes part of the financing gap, so the strike of the base case can be
 625 lower. Removing the extension raises every price-exposed strike by a few euros per megawatt-hour, most under wind-driven
 prices, where the extension is longest, and more for the CfD, whose basis of delivered energy at non-negative prices is the
 smaller one. The extension therefore has positive market value. An operator without support in the extension years extends.
 The FiT strike is unchanged, because its anchor carries no extension, and the ordering of the schemes is the same with and
 without the extension. On the replicate paths the magnitudes change slightly, and no ordering of the schemes and no sign of a
 630 result changes.




Figure 9. Required strike K^* per scheme and future ($NPV = 0$ at each scheme's reference dispatch, reference path). Filled: base case, support over the nominal life and the extension on market revenue; hollow: re-solved without the extension; stem: the change of the strike when the extension is removed. The capability variants overlap; dashed grey: mean price.

5.3.2 Composition of the support.

Figure 10a shows the lifetime support at the negative-revenue stop. The CfD and the capability contracts pay the same total
 at each future, as the strike determination implies (Sect. 3.4), and that total is the financing gap of the future, highest under
 wind-driven prices (Sect. 5.1). The schemes differ in what the support is paid for. The CfD pays only for delivered energy. The
 635 capability contracts pay support on delivered energy and on curtailed availability alike, and the support on availability curtailed
 at negative prices is about 50, 30 and 20 % of the total under wind-driven, mixed and solar-driven prices. The total does not
 increase by this. The support moves from hours in which the farm also earns the market leg into hours in which the market leg
 is zero, which is why the capability strike can fall while the CfD strike rises (Fig. 9). The FiT is shown as a line rather than a



bar, because it pays the tariff on all delivered energy and is the same in every future. A farm that stops under a FiT forfeits part
 640 of it.

Figure 10b shows, per contract, how the lifetime support changes if the operator curtails to the deepest threshold instead of
 at the negative-revenue stop. The stop has already removed every negative-price hour. A deeper threshold removes only hours
 with non-negative prices, and the contracts differ in how they settle these hours. The CfD pays on delivered energy. Every
 curtailed megawatt-hour forfeits the difference payment, the strike minus the reference price, and the support falls. The loss
 645 is largest under wind-driven prices, where both factors are largest. The strike is highest there. The energy removed is largest
 there too, because the hours priced just above zero are windy hours. The same threshold in euros per hour therefore removes
 about twice the energy it removes under solar-driven prices, where those hours are calm midday hours. The gated capability
 contract also settles curtailment at non-negative prices on delivered energy. Its support falls in the same way, at the smaller
 difference payment of the lower capability strike. The ungated capability contract pays for curtailed availability at every price.
 650 Its support on delivered energy falls by the same amount as under the gated contract, and its support on availability curtailed at
 non-negative prices rises by that amount, so the total does not change in any future. The small remainder is the higher technical
 availability that the fatigue relief produces in the O&M model. The public therefore pays the same support for less delivered
 energy.




Figure 10. Lifetime support (reference path). (a) At the negative-revenue stop, per future: CfD (left bar) and capability contracts (right bar; both variants coincide at the stop), split into support on delivered energy and on availability curtailed at $p < 0$; dashed: FiT (gross basis). (b) Change of the lifetime support when the operator curtails to the deepest threshold instead of at the stop, per contract and future; three bars per future: CfD, gated capability, ungated capability. The distance between the two capability bars is the support on availability curtailed at $p \geq 0$.



5.3.3 Private value of curtailment.

Figure 11 shows the NPV of threshold operation relative to each scheme's anchor. Because all schemes settle the same physical dispatch, every curve is the Market curve plus the present value of the change in that scheme's support, decomposed in Fig. 10b. The Market curve is the physical value of the action. The hours beyond the negative-revenue stop carry little revenue, and the revenue given up per step is small at first and grows with the threshold (Fig. 4). Over the first steps the extension value and the O&M cost reduction balance it, and the optimum lies at the stop or a few steps beyond it. At the deepest threshold a merchant operator loses up to eight million euros under wind-driven prices, small against the contract effects below. Against continuous production the stop itself is worth 20, 18 and 10 million euros to a merchant under wind-driven, mixed and solar-driven prices. Two thirds of that is the avoided negative revenue, one third the extension years. The three contracts solved for the stop pay no support on delivery at negative prices, so producing through these hours costs a CfD or capability operator the merchant's loss. Under the CfD the forfeited difference payments make the same curtailment the most costly, deepening with the wind coupling. Under the gated contract the forfeit at the lower strike places its curve between the CfD and the market, and its optimum returns to the stop. Under the ungated contract the support does not change with curtailment, so its curve coincides with the Market curve up to the small technical-availability effect, as long as support ends at the nominal life. In total, the same physical action changes the operator's NPV from a gain of one or two million euros to a loss of tens of millions among the price-exposed schemes, depending only on the settlement function. The same comparison can be made per megawatt-hour. Every megawatt-hour removed beyond the stop buys the operator a little more extension and a little less O&M cost within the nominal life. Together these gains, the fatigue value of the removed energy, amount to between one and seven euros per megawatt-hour, least under wind-driven prices, where the extension years are the cheapest. Against them stand two losses. The first is the market revenue of the removed energy, which is small just beyond the stop. The second arises under the contracts that pay on delivered energy, the CfD and the gated capability contract. Under these, every removed megawatt-hour forfeits its difference payment, the strike minus the reference price. That payment averages between thirty and eighty euros per megawatt-hour across the contracts and futures. The forfeit exceeds the fatigue value several times over, and this ratio, not settlement on delivered energy as such, keeps a CfD operator and a gated operator at the stop. A control whose fatigue value exceeded the difference payment would be worth taking under these contracts too. The forfeited payments are support that the public no longer pays, so the operator's loss is not a loss to the electricity system. The system loses only the fatigue value that stays unused when the operator remains at the stop, at most one or two million euros in any future. The FiT is solved for continuous production, and every stopped hour forfeits the gross tariff. The stop costs a FiT operator tens of millions of euros in every stressed future. Continuous production is worth more to it than the stop, and the FiT lies outside the figure.




Figure 11. NPV of threshold operation relative to each scheme’s anchor, per future (stars: best realised thresholds). All schemes settle the same physical dispatch. FiT omitted for scale (−32 to −112 M€ at the deepest threshold, relative to continuous production).

5.3.4 Carbon effects of curtailment

Support is paid per megawatt-hour of availability or delivery. One rationale for such support is the reduction of carbon emissions in the connected electricity system through wind generation. This subsection explores that rationale in three steps: First, the effect of curtailment on the amount of displaced emissions is examined. Second, the balance between plant-related emissions and displaced emissions is established, and the impact of lifetime extension on this balance is assessed. Third, the support amount per displaced tonne is determined. A key limitation of the following analysis is that, although the results are evaluated under different futures, the MDF time series used to calculate displaced emissions is held constant across all scenarios and corresponds to the historical case.

Curtailment removes production within the nominal life, and the resulting reduction in displacement depends on the hours in which production is curtailed. Across all price scenarios, the reduction in displacement is smaller than the reduction in energy. Under wind-driven prices, this difference is most pronounced. The displacement factor during curtailed hours is, on average, about 10 % below the overall average, resulting in a correspondingly smaller reduction in displacement than in energy. Under mixed prices, a similar pattern is observed, but the difference between the reductions in displacement and energy is less pronounced. Under solar-driven prices, the reductions are almost identical, although the reduction in displacement remains slightly smaller than the energy reduction. Further details are provided in Appendix B.

Figure 12 shows the balance between total displacement in the grid (environmental value of energy) and total life-cycle emissions attributed to the farm (environmental cost of energy) at the negative-revenue stop of each future, expressed as absolute GHG emissions (Fig. 12a) and per unit of energy delivered to the electricity grid (Fig. 12b). O&M emissions are separated from other life-cycle emissions, and displaced and O&M emissions associated with lifetime extension are distinguished from those occurring during the nominal lifetime. At the uncontrolled baseline, the total life-cycle emissions amount to about a seventh of the total displaced emissions. The nominal-life O&M emissions account for about one third of the total life-cycle emissions, mainly caused by fuel consumption due to on-site vessel activities (see Appendix B).



705 The negative-revenue stop removes displacement in the nominal years and adds displacement in the extension years. The sign of the overall change in displacement is set by the extension length, which, in turn, is determined by the fatigue relief of the future (Sect. 5.2). In the wind-driven price scenario, the stop removes a quarter of the nominal displacement. A lifetime extension of 10 years compensates for this loss, resulting in a net increase in displacement of about 17 %. This is because the curtailed hours have below-average displacement factors, while the extension years displace above the average factor. Under
 710 mixed prices, the three extension years return what the stop removed, and the net displacement stays within two percent of the baseline. Under solar-driven prices, the extension is little more than a year and the net displacement falls by a few percent.

Similar to displaced emissions, curtailment and lifetime extension affect the life-cycle O&M emissions. Fatigue relief reduces O&M emissions by a few kilotonnes during the nominal lifetime through fewer replacements, while the increase associated with lifetime extension is approximately an order of magnitude larger, mainly due to vessel fuel consumption. Nev-
 715 ertheless, these changes remain small compared with the changes in displaced emissions and do not alter the sign of the net displacement. Along the studied threshold sweep, a net gain persists at every threshold under wind-driven and mixed prices, and a net loss at every threshold under solar-driven prices (see Appendix B).

Per megawatt-hour delivered, environmental cost, value, and net value respond less to curtailment than the respective total emission quantities, as they are expressed relative to the lifetime energy yield. Curtailment reduces energy production, displace-
 720 ment, and O&M emissions, while lifetime extension increases all three. Differences in the energy-specific metrics therefore only reflect differences in the magnitude of these changes. As a result, environmental cost, value, and net value remain within a few percent of the baseline at every stop. The trends observed for the overall emission quantities remain unchanged, but the changes are smaller in magnitude.

The carbon payback time, defined as the total life-cycle emissions divided by the displacement per operating year, rises
 725 with curtailment. Production per nominal year decreases with curtailment, reducing the associated displacement and O&M emissions. At the same time, the associated lifetime extension substantially increases life-cycle O&M emissions. Although the extension also adds energy production, this additional production is not accounted for in the payback calculation. Thus, carbon payback performance is mainly determined by the amount of curtailment. With a payback time of about three years, the baseline shows the shortest payback, followed by the solar-driven and mixed-price futures, while the wind-driven future has
 730 the longest payback, exceeding four years at the negative-revenue stop and reaching about six years at the deepest threshold (Appendix B). However, a longer payback time does not indicate poorer environmental performance, because the extended lifetime also provides additional years of displacement that are not captured by the payback metric.

Table 4 relates the financial support received by the wind farm to the associated emissions displacement. The support per displaced tonne is calculated as the lifetime support divided by the total displacement over the realised life and can therefore be
 735 interpreted as the implicit abatement cost of the support. At the reference dispatch, it has one value per future, because every price-exposed scheme pays the same support there. Similarly, the table lists the support per net tonne, which also accounts for the life-cycle emissions linked to the wind farm, resulting in about 17 % higher values. The displacement per megawatt-hour delivered is nearly the same in every future, about a quarter of a tonne, because it is set by the generation the farm displaces in the Danish electricity system based on historical data only. However, it rises slightly with curtailment. The support




Figure 12. Lifetime carbon balance at the uncontrolled baseline and at the stopping point of each future. In (a), the absolute quantities are shown over the realised lifetime, with CO₂ displacement above zero and life-cycle emissions below zero; the markers indicate the net environmental value in absolute terms, the net displacement of Sect. 3.7. In (b), the same quantities are expressed per megawatt-hour delivered, with environmental value above zero and environmental cost below zero; the markers indicate the net environmental value per megawatt-hour delivered. Displacement is separated into the nominal 20-year period and the extension years. For life-cycle emissions, O&M emissions during the nominal lifetime and during the extension years are shown separately. The uncontrolled baseline is shown only once because its environmental performance is identical across all price futures: weather, displacement factors, and maintenance are unchanged, while only prices differ between futures.

740 per megawatt-hour delivered is the financing gap per unit of energy and follows the market revenue of the future: highest under wind-driven prices, and under solar-driven prices below the historical path, because the farm earns more in the market (Sect. 5.3). The support per displaced tonne therefore varies between futures with the financing gap, while the displacement per megawatt-hour hardly differs. The extension period introduces two additional effects: the wind-driven and the mixed future result in similar values, although the wind-driven gap is larger. This is because the ten extension years deliver energy without

745 support and dilute the support per megawatt-hour more than during the three extension years in the mixed scenario.



Table 4. Implicit abatement cost of the support at the anchors per future: support per megawatt-hour delivered, displacement per megawatt-hour, and their ratio per displaced and per net tonne.

	hist.	wind-driven	mixed	solar-driven
support per displaced tonne [€ t^{-1}]	168	217	208	137
support per net tonne [€ t^{-1}]	196	253	244	161
displacement [t MWh^{-1}]	0.233	0.244	0.238	0.234
support per MWh delivered [€ MWh^{-1}]	39	53	49	32

5.4 Share sensitivity under the wind driver

The wind-driven future was also solved to negative-price shares of 5 and 20 % (Sect. 4.3), with all three scenario classes on the reference path. Table 5 lists the results at the negative-revenue stop. The wind feed-in in negative-price hours stays a property of the driver. It drifts down mildly with the share, because the hours that a deeper share adds are less windy than the first ones.

750 A deeper share removes more of the same windy hours, so the relief and the extension grow with it. The net displacement stays above the baseline at every share, because a deeper share removes more nominal displacement and the longer extension returns it. The support total rises by about a tenth over the range, because the transform preserves the mean price. The strikes therefore move through the basis and the reference price rather than through the amount. The CfD strike rises steadily with the share, to four fifths above its historical level at 20 %, because its basis shrinks. The capability strike lies above its historical level at 5 %

755 and below it at 20 %. Two effects work against each other. The financing gap grows with the share, and the monthly reference price falls in the windy months, in which the available energy is largest. At 5 % the first dominates. At 10 % the two balance. At 20 % the second dominates. Over the whole range the capability strike stays within about 10 % of its historical level.

Table 5. Share sensitivity under the wind driver: the wind-driven future solved to 5, 10 and 20 % negative-price hours, reference path, at the negative-revenue stop. Coupling: wind feed-in in negative-price hours relative to its annual mean; relief: $1 - H$ at the nominal end; extension from the farm mean; strikes at the reference dispatch; support: mean of the contracts for difference, undiscounted over the nominal life; net displacement over the realised life relative to the uncontrolled baseline.

negative-price share	5 %	10 %	20 %
wind coupling [$\times$ mean]	2.55	2.49	2.30
relief at the nominal end	0.22	0.34	0.53
extension [y]	5.5	10.5	22.2
CfD strike [€ MWh^{-1}]	123	140	181
capability strike [€ MWh^{-1}]	106	101	90
support at the negative-revenue stop [M€]	339	358	372
net displacement vs baseline [%]	+15	+17	+14



6 Discussion

Control, support contract and price future were co-varied in one hourly simulation, so that loads, settlement and finance derive from the same operation. The market and settlement effects are the largest, at tens to hundreds of millions of euros over the project life. The lifetime extension is worth a few million euros in every future. The damage-coupled O&M effect is small at the assumed coupling strength (Sect. 5.2). The integration is needed where these interact, in the position of the value optimum, which balances an immediate revenue loss against discounted extension value after the nominal end.

6.1 Price driver and negative-price share

At one negative-price share the futures differ in market revenue, fatigue relief and extension, required strikes and carbon balance. The share alone therefore does not fix these results. At a given share they are set by the wind feed-in in negative-price hours relative to its annual mean, which is observable from published market data. The negative-price share, the stress variable of the scenario literature, needs this second quantity beside it. Under the wind driver the share scales the magnitudes without changing the ordering (Sect. 5.4). Which value of the coupling a given system reaches, and when, is the output of an electricity system model. Linking such a model to the project-level valuation of this study would replace the constructed futures by projected ones and turn the ordering reported here into a projection of the actual value. This link does not exist yet.

Where the coupling is strong, evaluations by energy yield or levelised cost rank operating policies inversely to their value, because the negative-revenue stop removes energy and raises project value at once (Sect. 5.1). A ranking has to carry revenue. The fatigue value of price-responsive operation is likewise a property of price formation. A price signal reduces fatigue only to the extent that low prices coincide with high loads. Under wind-driven price formation they do, and price-responsive operation is load-reducing operation. Under solar-driven price formation they do not, and fatigue relief requires load information in the control (Sect. 5.2). The observed zones moved from wind-driven to solar-driven price formation between 2019 and 2025 (Fig. 2), so this value declined in the observed markets over that period.

The money value of the relief declines less than the relief itself, because the years it buys become more valuable as price formation turns solar-driven. Three effects set the value of an extension year. First, the extension replays the future under the negative-revenue stop's own dispatch, so an extension year sells the hours the stop leaves, at the prices of that future. Under wind-driven prices the stop has removed the windiest hours, and the remaining energy sells at the depressed capture price of a wind-driven system. Under solar-driven prices the stop has removed little, and the remaining energy sells above the mean. Second, the O&M cost of an aged farm does not depend on the future. It is close to the revenue of a wind-driven extension year, so such a year yields little net income, while the revenue of a solar-driven year is well above it. Third, the extension years follow the nominal life, and the longer the extension, the further its last years lie from the investment and the less they count after discounting. The price formation that relieves the most fatigue therefore pays the least for the years it adds. Wind-driven prices give many years of small net income. Solar-driven prices give fewer years of larger net income. The steps beyond the stop show the same pattern. The threshold is set in euros per hour, so a step removes hours of the same revenue in every future, and the revenue it gives up is of similar size. The step removes different hours in each future. Under wind-driven prices it



removes two to three times the energy it removes under solar-driven prices, because the hours in the same revenue band are windy hours at low prices in the one future and calm hours at moderate prices in the other. The years a step buys follow the energy removed, most under wind-driven prices and fewest under solar-driven. Their value runs the other way, little under wind-driven prices and much under solar-driven prices. The product is largest under mixed prices, where a step still buys most of a year and the year still earns near the mean price. This is why the mixed curve lies above the wind-driven one in Fig. 11.

Monetised through an unsupported extension, fatigue is therefore worth a few euros per megawatt-hour of production. This is the marginal value of the fatigue budget under the conventions of this study, and it values any control that saves fatigue. For a control that gives up energy, such as the negative-revenue stop, it is set against the revenue given up. For a control that saves fatigue without giving up energy, it is the value of the relief itself. A real extension is an option exercised at the nominal end, with the prices and the condition known then.

6.2 Contract terms: basis, strike and support window

The support requirement under each price future is paid by each scheme at the reference dispatch. A scheme that induces production through negative prices, as the tariffs of the past did, carries the cost of that production in its requirement, between ten and twenty million euros in the stressed futures. Beyond that, the contract decides the basis on which it pays, the strike and the support window. These three terms decide who carries which risk, which operation the contract pays for, and where the fatigue value ends up.

6.2.1 Settlement basis.

Every price-exposed scheme covers the same gap, so a strike exceeds the reference price by the gap per unit of the scheme's settlement basis, the B_t of Eq. (3). Strikes on different bases are therefore not comparable: a comparison of schemes by strike, or an auction in which bids on different bases compete in one price per megawatt-hour, misranks them, and the comparable quantity is the present value of support. Because the CfD strike rises with the wind coupling at unchanged costs successive auction rounds would show rising CfD strikes where price formation turns wind-driven and none where it turns solar-driven.

The two variants of the capability contract separate two properties of the basis. Available energy does not depend on the operator's decisions, so under the ungated contract the operator is paid the same whether it delivers or curtails and faces the market price at the margin. It curtails only where the fatigue value of an hour exceeds its market value, so the ungated contract is the efficient one in this model and its optimum is the merchant's. The delivered-energy gate removes this property at non-negative prices. Under the gated contract an hour curtailed at a non-negative price forfeits its difference payment, although the energy is worth only its market price to the electricity system, a few euros per megawatt-hour at the thresholds of the sweep. This is the distortion of the production-based CfD. The gated operator therefore returns to the negative-revenue stop and forgoes one or two million euros of extension value, while the public pays the same support either way. The gate buys verification. The available energy is an estimate, and the gate limits the volume settled on it to the negative-price hours. It also removes the gain from declaring a commanded shutdown in an hour of technical unavailability. The choice between the variants is a trade between a small forgone surplus and a smaller, less contestable volume settled on available energy. A settlement on



the output of a reference plant that the operator does not control (Newbery, 2023) needs no estimate of the farm's available energy and keeps the efficiency of the ungated contract without its contestable volume. The ungated contract is anchored at the stop although its rule induces a dispatch at the stop or a few steps beyond it. A strike solved at the induced dispatch would be lower by less than one euro per megawatt-hour, because the extension value the operator gains beyond the stop would then count towards the financing gap. Solved at the stop, the strike leaves that value with the operator.

6.2.2 Strike.

What the capability contracts change is who carries the risk that price formation develops differently from what was assumed at award. The CfD's exposure to a wind-driven development is large and one-directional. Its basis shrinks, and the required strike rises by up to about 80 % with the depth of the development. A strike fixed at award then falls short of the gap, which the operator carries. The capability contract's exposure is small and of either sign. Its basis, the available energy, does not shrink with the development, and its strike moves through two terms that work against each other. A wind-driven development raises the financing gap, which raises the strike. It also lowers the reference prices of the windy months, in which most of the available energy lies, so every megawatt-hour of basis earns a larger difference payment at the same strike, which lowers it. The two nearly cancel. The required strike ends above the awarded one under a mild development and below it under a deep one, within about 10 % either way. Under a solar-driven development the reference price of the farm's productive months hardly changes. The gap shrinks, because the farm earns more at the fixed mean price of the futures, and the strike falls below its historical level. A fall of the price level itself is not represented here. It would lower the reference price and raise the gap per unit of basis by nearly the same amount, so the strike would move little. The stability is a property of the reference price. With an hourly reference the support leg would be the strike times the basis minus the market value of the basis. The capability strike would then be the costs per unit of available energy, independent of the price shape altogether, and the CfD strike would move through its basis alone. The monthly reference leaves the deviation of the capture price from the reference price within the month with the operator. That deviation is why the capability strike moves at all. For the same expected support the operator therefore carries a lower risk under a capability contract, because its strike moves less with the price formation. Under risk-sensitive financing a lower risk means a lower cost of capital and a smaller gap. Among the schemes that induce the negative-revenue stop, this is the one route by which the choice of scheme lowers what the public pays. It lies outside the fixed cost of capital of this study, so the equal-support result is a bound.

Which side should carry the remaining exposure is a question of allocation. Price formation in surplus hours is set by the fleet mix, flexibility and interconnection, which the public steers, and by siting, technology and the curtailment behaviour of the fleet, which developers set under the rules they are given. The observed shift to solar-driven formation came largely from solar investment outside tenders, so neither side controls the driver alone. The systemic part belongs with the public and the plant-specific part with the developer. A strike indexed to the realised coupling of the fleet, the wind feed-in in negative-price hours relative to its annual mean, performs that split. It removes from the contract a forecast that both sides would find wrong for most of a contract life that outlasts the observed drift (Fig. 2). The public carries the systemic development as it occurs. The operator keeps capture price and operation, because the index is a system statistic and not the plant's own revenue.



6.2.3 Support window.

Under a support window that ends at the nominal life the extension lies outside the contract. The operator extends on market
860 revenue, the public pays nothing for the extension years, and the carbon gain of those years accrues without support (Sects. 5.3,
5.3.4). The contract's only part in the fatigue value is then to leave it intact, which the CfD beyond the negative-revenue stop
does not. The basis decides whether relief beyond the stop is produced at all. The window decides whether the relief that is
produced remains a matter between operator and market or becomes a public payment. A window that follows the extension
years turns them into supported years. Under an ungated capability contract this is additional support for the curtailment that
865 produces them. The size of that payment follows the extension, which the aggregation level of the assessment and every limit
outside the fatigue budget set (Sect. 3.5). A window defined in years, as in current practice, ends at the nominal life and leaves
the extension to the market.

6.3 Carbon balance and carbon metrics

The carbon results of Sect. 5.3.4 show that the settlement function and the price futures maintain or improve the overall
870 environmental performance of the wind farm in most cases, with a substantial improvement in the wind-driven case and a
slight reduction only in the solar-driven case. Although curtailment reduces electricity delivery and consequently the emissions
displaced during the nominal lifetime, the associated reduction in fatigue loading enables a lifetime extension that largely
offsets, and in most cases exceeds, the displacement lost during the nominal lifetime. The additional O&M emissions associated
with the extension remain approximately one order of magnitude smaller than the corresponding displaced emissions. Likewise,
875 the life-cycle emissions of the wind farm are substantially smaller than the emissions displaced over its lifetime. These results
are obtained, however, from operating strategies selected on economic grounds, without environmental performance being
included as a decision criterion. Future work should investigate co-decision criteria that consider economic and environmental
implications jointly and move towards optimization-based operational strategies that seek to maximize both economic and
environmental performance.

880 The support paid for providing net environmental value reflects two complementary objectives of renewable-energy support.
First, it enables additional wind generation that displaces electricity from other generators and thereby reduces system-level
carbon emissions. Second, it ensures the delivery of low-carbon electricity to the system, even when its market value alone is in-
sufficient to cover the cost of generation. The support per displaced net tonne therefore captures only one part of the purpose of
the subsidy. At the reference dispatch, the support amounts to approximately 161–253 € per net tonne of CO₂ displaced across
885 the considered futures. For comparison, the average European Union Emissions Trading System allowance price was around
65 €/tCO₂ in 2024 (European Commission, 2025a). This difference should not be interpreted as the support being an alterna-
tive price for CO₂: an emissions price primarily provides an incentive to avoid emissions, whereas renewable-energy support
additionally enables the production and delivery of low-carbon electricity. This distinction becomes particularly relevant as the
electricity system becomes cleaner. In the present study, the lifetime emissions displaced by the wind farm are approximately
890 seven times larger than its life-cycle emissions, with the relatively modest ratio reflecting Denmark's already high share of



renewable generation during the 2017–2020 period used to derive the displacement factors. In more carbon-intensive electricity systems, such as Germany, higher marginal displacement factors result in larger displacements and, consequently, higher ratios of displaced to life-cycle emissions (Kainz et al., 2025). Conversely, as the electricity system decarbonises, less carbon-intensive generation remains available to be displaced, reducing the CO₂ displaced per additional MWh and consequently increasing the support per displaced tonne. This should not be interpreted as a deterioration in the environmental performance of wind energy, but rather as a shift in the role of renewable support: while displacement of carbon-intensive generation is particularly important in a carbon-intensive system, support in a highly decarbonised system increasingly facilitates the continued supply of low-carbon electricity to meet growing demand.

6.4 Limitations

Links to the electricity system. The price series is an input. The farm’s curtailment does not feed back into prices. At fleet scale it would. If every wind farm stopped at negative revenue, the surplus in those hours would shrink and with it the negative-price hours that drive the results. The futures inherit the shape of the observed prices and are not projections of the kind system studies provide (ENTSO-E and ENTSG, 2025; Danish Energy Agency, 2024) (Sect. 6.1). The cost of capital is fixed for all schemes (Sect. 6.2).

Conventions the results depend on. The futures hold the mean price, so the hours without the driver are lifted by what the surplus hours lose. The results that follow from the negative-revenue stop, the energy removed, the relief, the extension and the carbon balance, do not depend on this convention, and neither does the rise of the CfD strike under wind-driven prices, which is a basis effect. The farm’s revenue gain under solar-driven prices, the lower support and support per tonne there, and the fall of the capability strike do depend on it, and a development in which the price level falls is not represented. The extension is the fatigue-permitted one and therefore an upper bound (Sect. 3.5.3). The re-solved strikes without extension give the lower bound of its value. Its value counts the deferral of decommissioning and no cost for the assessment and refurbishment (Sects. 3.4, 3.5), and the extension years are merchant years on a stationary future, so it is an ex-ante value under these conventions rather than the value of a decision taken at the nominal end. The damage–O&M coupling sensitivity α is not calibrated. Its contribution is linear in α and small at $\alpha = 1$ (Sect. 5.2), so no conclusion depends on it.

Quantities taken as known and terms not simulated. The capability contracts are evaluated with the available energy of the uncontrolled baseline, a perfectly observed counterfactual. In practice available energy must be estimated from operating data, and the estimate can be inaccurate or contested, so the results are the best case for settlement on it. Support that follows the extension years and the indexed strike are proposed and discussed (Sect. 6.2), not simulated.

Fidelity of the physical model. The channel-level checks are per-axis in the nacelle-following frame while the structure is fixed, so the channel directions sweep the tower circumference as the nacelle tracks the wind rose. This mixing spreads single-channel damage over many locations, while the discarded phase correlation between channels can concentrate it at single locations. The physically mixed case lies between the farm-mean and the worst-channel extension. The parked-mode surrogates are wind-only, without wave excitation (Appendix A).



Marginal displacement factors. A key limitation of the carbon analysis is that the hourly MDFs are held constant across all future scenarios and are based on the historical 2017–2020 period. The resulting carbon metrics therefore do not capture changes in the electricity system, generation mix, or marginal generation associated with the different future scenarios. A fully consistent assessment would require scenario-specific MDF time series derived from an explicit model of the future electricity system. Developing such a coupled electricity-system model is beyond the scope of this study.

Transferability. The case is a single site with unusually strong wake interaction, representative costs and a greenfield project. Orderings, signs and mechanisms carry over. Magnitudes depend on the site, the wake interaction and the cost basis (Sect. 4.1).

7 Conclusions

We evaluated wind farm value, revenue-triggered curtailment and five support designs on the historical path and three price futures with the same negative-price share, driven by wind, by wind and solar, and by solar feed-in. Loads, settlement and finance were computed from one hourly simulation, and every contract was priced at its own reference dispatch. Four results answer the research questions.

First, the share of negative-price hours alone does not measure a wind farm's exposure to negative prices. The exposure depends on whether the farm produces in these hours. When wind drives the negative prices, they are the farm's most productive hours. The uncontrolled farm then captures a price far below the mean, and the negative-revenue stop removes a large part of its energy. When solar drives them, they are midday hours which are statistically decoupled from the wind speed. The stop then removes little energy, and the farm captures a price close to or above the mean. At a given share, the wind feed-in in negative-price hours relative to its annual mean therefore orders every result of this study, and the share scales them. The observed bidding zones moved across the range of this coupling between 2019 and 2025.

Second, revenue-triggered shutdown is load-reducing control only under wind-driven price formation. There the negative-price hours are hours of high wind and high fatigue damage, so the negative-revenue stop removes more fatigue than energy. The lifetime can then be extended by several years. The market revenue of these years exceeds their operating cost by a small margin. This margin is what remains to pay for the lifetime assessment and the refurbishment. Under solar-driven price formation the stop removes fatigue only in proportion to the energy it removes. Fatigue relief then requires a control that uses load information. The worst-channel check confirms a shorter extension than the farm-mean rule under wind-driven prices and none under solar-driven prices.

Third, the contracts differ little where few hours are priced negative, and they separate where negative prices coincide with wind. The CfD pays no support on energy delivered at negative prices. Under wind-driven prices a large part of the farm's energy falls into these hours, so the same financing gap has to be covered on less energy, and the CfD strike rises by about 40 %. At a negative-price share of 20 % it rises by about four fifths. The capability contracts pay support on the available energy in these hours. Their strikes stay within about 10 % of the historical level and fall under solar-driven prices. The present value of support at the reference dispatch is nevertheless the same for every price-exposed scheme. The scheme decides on which energy and at what time the public pays, and who carries the risk. It does not decide how much the public pays, but it decides



what curtailment beyond the negative-revenue stop is worth to the operator. The same physical action yields a gain of one or two million euros under the ungated capability contract and a loss of tens of millions under the CfD. Most of this loss is support that the operator no longer receives and the public no longer pays.

960 Fourth, the results show that curtailment based on economic criteria generally maintains the net environmental benefit of the wind farm, with a significant improvement in the wind-driven future and a slight reduction in the solar-driven future. Although curtailment reduces electricity production and emissions displacement during the nominal lifetime, the resulting fatigue relief enables lifetime extension that largely offsets, and in most cases exceeds, the displacement lost. Additional O&M emissions remain small compared with the displaced emissions, such that the net environmental benefit remains positive
965 across all considered cases and improves in several cases. While support per displaced tonne varies between futures with the financing gap, displacement per megawatt-hour changes only slightly. These results describe the environmental consequences of economically derived operating strategies. Future work should consider environmental performance alongside economic criteria in operational decision-making.

These findings on contracts apply to any control that gives up delivered energy to reduce fatigue loads. The negative-revenue
970 stop studied here is one example, derating and partial curtailment are others. For every megawatt-hour such a control gives up, the operator weighs three terms. It gains the fatigue value of that energy, which the control sets through the fatigue it saves per megawatt-hour and the price formation sets through what an extension year earns. It loses the market value of the energy, which the price formation sets through the hours in which the control acts. Under a contract settled on delivered energy it also loses the support paid on that energy. This third term is a hurdle that the control cannot move. It equals the difference payment under
975 settlement on delivered energy and zero under settlement on available energy. A contract settled on delivered energy therefore blocks every control whose fatigue value, net of the market value of the energy, lies between zero and the difference payment. It saves the public nothing by doing so, because the operator simply does not act. The width of this band is a property of the contract and the price formation, tens of euros per megawatt-hour here against a fatigue value of a few euros. It is widest under wind-driven prices, where curtailment relieves the most fatigue and the CfD strike is highest. Better load information raises
980 the fatigue value and lets a control clear the hurdle, but it does not lower the hurdle. The negative-revenue stop clears it under every scheme, because delivery at negative prices earns no support under any of them and the energy given up has negative market value. Beyond the stop, the settlement basis decides whether a control is blocked, and the support window decides who pays for the fatigue value it produces.

Four recommendations follow. First, the rule that pays no support on energy delivered at negative prices should stay. It makes
985 delivery at negative prices earn nothing under the CfD, as settlement on available energy does under the capability contracts, so every price-exposed scheme induces the negative-revenue stop, and the stop needs no load information. A scheme that induces production at negative prices has to cover the cost of that production, which adds between ten and twenty million euros to the support requirement in the stressed futures. Second, at non-negative prices the hurdle is removed by settling support on a basis that the operator's control does not change, either the plant's own available energy or the output of a reference plant.
990 The first has to be estimated, as the output forgone under instructed curtailment is estimated under existing contracts, and the estimate can be disputed. The second needs no estimate but leaves the plant with the difference between its own output and the



reference. Under either basis the operator curtails where the fatigue value exceeds the market value of the energy. The choice between them is a choice of verification and risk allocation. Third, the strike can be indexed to the coupling of the fleet, so that it follows a change in price formation after award. This index is proposed here and was not simulated. Fourth, the support window should end at the nominal life. The operator then receives the fatigue value of its control through the extension years it sells on the market, and the public pays nothing for the extension. A window that follows the extension would pay the operator more than the fatigue value is worth and induce curtailment beyond the optimum.

Code and data availability. The WINPACT simulation framework, the Lillgrund input set, the experiment driver and the postprocessing package that generates every figure and table from the raw campaign outputs will be archived at [DOI TODO] upon acceptance. Market and meteorological input data are available from ENTSO-E and the ERA5 archive respectively.

Appendix A: Load response of revenue-triggered shutdown: from operating modes to component damage

This appendix carries the channel-level evidence behind Sect. 5.2 in two steps: in which operating mode the damage accrues, and what the component aggregation does to the channel responses. Both figures use the reference weather–price path.

A1 Attribution: damage by operating mode

Assigning each hour's damage increment to its operating mode (Fig. A1) locates each channel's response and the two mechanisms that keep damage after the negative-revenue stop. Tower-base fore–aft damage accrues in producing hours and is removed with them. Its idle and transient shares are small. Tower-base side–side damage also accrues while the turbines are parked. Curtailment converts producing hours into idle hours instead of removing the exposure, and the side–side channel, which has no aerodynamic damping, keeps accumulating wind-induced damage in those hours. Under wind-driven prices its idle contribution is about a third of the channel's baseline total and nearly independent of the threshold. This floor limits the side–side relief and makes the channel the binding one for the worst-channel check under wind-driven and mixed prices. Blade-flap damage accrues almost entirely in producing hours, and the stop removes most of it, but the start-up transients that follow every stop keep about 20 % of the channel's baseline total, again nearly independent of the threshold. Under solar-driven prices the stop removes little producing-hour damage, because the stopped midday hours carry average wind, and the start-up transients return what it removes. The flap channel ends at its baseline budget and binds the worst-channel check, which therefore yields no extension at that stop (Sect. 5.2). Shutdown contributions are negligible for all channels. The remaining uncertainties are properties of the surrogate library. Parked-mode loads are wind-only, without wave excitation, and the transient windows are represented by the dedicated start-up and shutdown surrogate bundles.




Figure A1. End-of-life damage by operating mode for three channels (columns), normalised by the channel’s baseline total, for the wind-driven and the solar-driven future (rows) and per threshold (bars; the first bar is the uncontrolled baseline). Bar height: channel ratio R_{ch} ; colours: causing mode.

A2 Aggregation: channels to components

1020 The channel responses meet the weighted normalised component index of Eq. (7) (Fig. A2, wind-driven future). The raw
 baseline shares annotated in the figure show why the index is formed on ratios rather than on raw damages. Because Wöhler
 exponents differ between material classes, one channel can carry almost all of a component’s raw damage, fore–aft for the
 tower base and edgewise for the blades. A raw-weighted index would track those channels alone. The weighted mean of ratios
 gives every channel the influence of its weight. The index remains an average. Where the channels of a component are relieved
 1025 unequally, it reports a value between them, and a channel relieved less than its partners, such as tower-base side–side, enters
 only at its weight. Integrity statements therefore use the worst-channel check of Eq. (10), which involves no weights.




Figure A2. Channels and the weighted normalised component index, per component (wind-driven future), relative to own baselines. Annotated: the first channel's share of the raw baseline sum.

Appendix B: Carbon balance mechanics

This appendix documents the mechanics behind the carbon balance of Sect. 5.3.4: the displacement factor of the hours that curtailment removes, the response of the balance terms along the threshold sweep, and the composition of the life-cycle emissions by category.

Figure B1 shows the displacement lost within the nominal life against the energy removed, per future and threshold. The slope of each line is the displacement factor of the removed hours relative to the farm's average: below the proportional line under wind-driven and mixed prices, on it under solar-driven prices (Sect. 5.3.4).

Figure B2 follows the terms of the balance along the threshold for the three stressed futures, relative to the uncontrolled baseline. A deeper threshold removes more hours and lengthens the extension, and the terms respond to the two differently. The displacement (a) rises along the sweep under mixed prices, where each added extension year returns more than the step removes, stays about level under wind-driven prices and falls under solar-driven prices, where the extension is short. The O&M emissions (b) rise with the threshold in every future, because each additional operating year adds nearly a full year of O&M




Figure B1. GHG displacement lost against energy removed within the nominal life, relative to uncontrolled operation, per future (markers: the seven thresholds). The dashed line represents proportional losses, where the relative loss in GHG displacement equals the relative loss in energy.

emissions. In contrast, O&M emissions during the nominal lifetime decrease slightly due to the reduced repair intensity from
 1040 fatigue relief, but this reduction remains small compared with the additional emissions from the extension. The net displacement
 (c) is the difference of the two: it stays above the baseline at every threshold under wind-driven and mixed prices and below
 it at every threshold under solar-driven prices, so its sign is set by the future, not by the threshold. The carbon payback (d)
 rises with the threshold in every future since curtailment reduces the displacement per operating year. However, fatigue relief
 extends the operational lifetime and adds further displacement, so a longer carbon payback time does not necessarily indicate
 1045 a worse overall environmental performance.

Figure B3 separates O&M emissions from the fixed emissions associated with all other life-cycle phases. At the uncontrolled
 baseline, O&M emissions account for about one third of total life-cycle emissions. Vessel fuel consumption alone accounts for
 about one quarter of total life-cycle emissions, split between the crew transfer vessel and the field service vessel. Spare-part
 production contributes about 5 %, while transport fuel and vessel construction together account for less than 3 %. Curtailment
 1050 leaves the fixed emissions unchanged but increases O&M emissions with each extension year. Consequently, the O&M share
 increases with the extension, most under wind-driven prices, almost entirely due to additional vessel fuel consumption.




Figure B2. Carbon balance terms along the shutdown threshold for the three stressed futures, relative to the uncontrolled baseline: (a) change in CO₂ displaced over the realised life; (b) change in O&M emissions, total (solid) and nominal years only (dotted); (c) change in net displacement; (d) carbon payback time, baseline value dashed. A threshold of zero refers to the negative-revenue stop.




Figure B3. Life-cycle emissions by category over the realised life. The bars show the uncontrolled baseline, the negative-revenue stop, and the deepest sweep threshold of each future. (a) shows the totals, while (b) reports relative shares. Transport fuel covers all vessel types; on-site fuel is split by vessel class.

Author contributions. MG: ideation, concept, simulations, writing. SK: environmental value methodology, writing. VP: response framework, review. AI: support-scheme representation, review. AK, TG, ISF, IE: review and discussions. ND: ideation, concept, review.



Competing interests. At least one of the (co-)authors is a member of the editorial board of Wind Energy Science.

1055 *Acknowledgements.* This work has been supported by the TWAIN project, which received funding from the European Union's Horizon Europe Programme under grant agreement No. 101122194. The authors acknowledge Annamaria Scherzl and Carlo L. Bottasso, with the former preparing the training data for the MDF model and the latter supervising the environmental performance methodology and contributing to its conceptualization. The authors acknowledge the use of large language model tools to support language editing and improve clarity and coherence of the manuscript.



1060 References

- ACER: Key developments in electricity wholesale markets – 2024 Market Monitoring Report, European Union Agency for the Cooperation of Energy Regulators, https://www.acer.europa.eu/monitoring/MMR/electricity_key_developments_2024, 2024.
- ACER: Key developments in European electricity and gas markets – 2025 Market Monitoring Report, European Union Agency for the Cooperation of Energy Regulators, https://www.acer.europa.eu/monitoring/MMR/electricity_gas_key_developments_2025, 2025.
- 1065 Anand, A., Braunbehrens, R., Guilloiré, A., and Bottasso, C. L.: Economic lifetime-aware wind farm control, Wind Energy Science Discussions, preprint wes-2025-67, accepted for publication, <https://doi.org/10.5194/wes-2025-67>, 2025.
- Antweiler, W. and Müsgens, F.: The new merit order: The viability of energy-only electricity markets with only intermittent renewable energy sources and grid-scale storage, Energy Economics, 145, 108 439, <https://doi.org/10.1016/j.eneco.2025.108439>, 2025.
- Arvesen, A. and Hertwich, E. G.: Assessing the life cycle environmental impacts of wind power: A review of present knowledge and research needs, Renewable and Sustainable Energy Reviews, 16, 5994–6006, <https://doi.org/10.1016/j.rser.2012.06.023>, 2012.
- 1070 Beiter, P., Guillet, J., Jansen, M., Wilson, E., and Kitzing, L.: The enduring role of contracts for difference in risk management and market creation for renewables, Nature Energy, 9, 20–26, <https://doi.org/10.1038/s41560-023-01401-w>, 2024.
- Bonou, A., Laurent, A., and Olsen, S. I.: Life cycle assessment of onshore and offshore wind energy – from theory to application, Applied Energy, 180, 327–337, <https://doi.org/10.1016/j.apenergy.2016.07.058>, 2016.
- 1075 Braunbehrens, R., Anand, A., Campagnolo, F., and Bottasso, C. L.: First experimental results on lifetime-aware wind farm control, Journal of Physics: Conference Series, 2767, 032 042, <https://doi.org/10.1088/1742-6596/2767/3/032042>, 2024.
- Brown, T. and Reichenberg, L.: Decreasing market value of variable renewables can be avoided by policy action, Energy Economics, 100, 105 354, <https://doi.org/10.1016/j.eneco.2021.105354>, 2021.
- Canet, H., Guilloiré, A., and Bottasso, C. L.: The eco-conscious wind turbine: design beyond purely economic metrics, Wind Energy Science, 8, 1029–1047, <https://doi.org/10.5194/wes-8-1029-2023>, 2023.
- 1080 Danish Energy Agency: Analyseforudsætninger til Energinet 2024, Energistyrelsen, <https://ens.dk/analyser-og-statistik/analyseforudsætninger-til-energinet>, tODO: switch to the AF25 edition if published before submission, 2024.
- DNV: DNV-ST-0262: Lifetime extension of wind turbines, Standard, Edition 2016-03, amended 2021-11, DNV AS, first edition 2016 as DNVGL-ST-0262; companion certification procedure DNV-SE-0263, 2021.
- 1085 Electricity Maps: Electricity Maps emission factors methodology: comparison and validation of the Electricity Maps emission factors dataset against the IEA emission factors, White paper, Electricity Maps, <https://portal.electricitymaps.com/docs>, accessed 19 August 2026, 2024.
- Energinet: Generation per Production Type and Exchange, Energi Data Service, <https://www.energidataservice.dk/tso-electricity/GenerationProdTypeExchange>, accessed 19 August 2026, 2026.
- ENTSO-E and ENTSG: TYNDP 2024 Scenarios Report, Final version, January 2025, <https://2024.entsos-tyndp-scenarios.eu/>, 2025.
- 1090 European Commission: Report on the functioning of the European carbon market in 2024, Report from the Commission COM(2025) 735 final, European Commission, Brussels, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52025DC0735>, 2025a.
- European Commission: Commission guidance on the design of two-way contracts for difference, Official Journal of the European Union, C/2025/6701, 19 December 2025, <https://eur-lex.europa.eu/eli/C/2025/6701/oj>, cELEX: 52025XC06701, 2025b.
- European Union: Regulation (EU) 2024/1747 of the European Parliament and of the Council of 13 June 2024 amending Regulations (EU) 2019/942 and (EU) 2019/943 as regards improving the Union’s electricity market design, Official Journal of the European Union, <https://eur-lex.europa.eu/eli/reg/2024/1747/oj>, eLI: eli/reg/2024/1747/oj, 2024.
- 1095



- Fabra, N. and Imelda: Market power and price exposure: Learning from changes in renewable energy regulation, *American Economic Journal: Economic Policy*, 15, 323–358, <https://doi.org/10.1257/pol.20210221>, 2023.
- 1100 Garcia-Teruel, A., Rinaldi, G., Thies, P. R., Johanning, L., and Jeffrey, H.: Life cycle assessment of floating offshore wind farms: An evaluation of operation and maintenance, *Applied Energy*, 307, 118 067, <https://doi.org/10.1016/j.apenergy.2021.118067>, 2022.
- Geis, J., Neumann, F., Lindner, M., Härtel, P., and Brown, T.: Price formation in a highly-renewable, sector-coupled energy system, *Energy Economics*, 157, 109 213, <https://doi.org/10.1016/j.eneco.2026.109213>, 2026.
- Gräfe, M., Kainz, S., Ludot, A., Pettas, V., Anand, A., and Bottasso, C. L.: Impact of reliability parameters on O&M cost and greenhouse gas emissions of offshore wind farms, *Journal of Physics: Conference Series*, 3131, 012 034, [https://doi.org/10.1088/1742-](https://doi.org/10.1088/1742-6596/3131/1/012034)
 1105 6596/3131/1/012034, 2025.
- Gräfe, M., Pettas, V., Ioannou, A., Kolios, A., and Dimitrov, N.: Wind farm project valuation under revenue-driven farm control considering life time extension, *Journal of Physics: Conference Series*, 3224, 062 063, <https://doi.org/10.1088/1742-6596/3224/6/062063>, proc. The Science of Making Torque from Wind (TORQUE 2026), 2026a.
- Gräfe, M. J., Ludot, A., Shields, M., Kolios, A., Pulikollu, R., and Dimitrov, N.: Review and quantification of major risks in wind farm
 1110 development and operation, *Wind Energy Science*, 11, 3475–3507, <https://doi.org/10.5194/wes-11-3475-2026>, 2026b.
- Hirth, L.: The market value of variable renewables: The effect of solar wind power variability on their relative price, *Energy Economics*, 38, 218–236, <https://doi.org/10.1016/j.eneco.2013.02.004>, 2013.
- Hirth, L.: Do wind and solar curtail at negative electricity prices? Incentives and evidence across two decades of German renewable support schemes, *arXiv preprint 2609.10053*, <https://doi.org/10.48550/arXiv.2609.10053>, 2026.
- 1115 Hulsman, P., Andersen, S. J., and Göçmen, T.: Optimizing wind farm control through wake steering using surrogate models based on high-fidelity simulations, *Wind Energy Science*, 5, 309–329, <https://doi.org/10.5194/wes-5-309-2020>, 2020.
- IEC: IEC TS 61400-28:2025: Wind energy generation systems – Part 28: Through-life management and life extension of wind power assets, Technical specification, Edition 1.0, International Electrotechnical Commission, 2025.
- Ioannou, A., Gräfe, M., Beiter, P., and Kitzing, L.: Comparing contract-for-difference designs under market price uncertainty, *SSRN preprint*
 1120 7269013, <https://doi.org/10.2139/ssrn.7269013>, 2026.
- Johanndeiter, S., Finke, J., and Heuer, J.: Most certainly certain? The impact of contract for difference design on renewables’ strike prices and electricity market risks, *arXiv preprint 2512.17508*, <https://doi.org/10.48550/arXiv.2512.17508>, 2025a.
- Johanndeiter, S., Helistö, N., and Bertsch, V.: Does the difference make a difference? Evaluating contracts for difference design in a fully decarbonised European electricity market, *Resource and Energy Economics*, 83, 101 495, <https://doi.org/10.1016/j.reseneeco.2025.101495>,
 1125 2025b.
- Kainz, S., Guilloré, A., and Bottasso, C. L.: How do technological choices affect the economic and environmental performance of offshore wind farms?, *Journal of Physics: Conference Series*, 2767, 082 005, <https://doi.org/10.1088/1742-6596/2767/8/082005>, 2024.
- Kainz, S., Scherzl, A., Guilloré, A., Anand, A., and Bottasso, C. L.: An initial study on the environmental value of wind farm control, *Journal of Physics: Conference Series*, 3131, 012 044, <https://doi.org/10.1088/1742-6596/3131/1/012044>, 2025.
- 1130 Kanev, S. K., Savenije, F. J., and Engels, W. P.: Active wake control: An approach to optimize the lifetime operation of wind farms, *Wind Energy*, 21, 488–501, <https://doi.org/10.1002/we.2173>, 2018.
- Kern, T., Wasmeier, L., Bantelmann, P., and Wendlinger, C.: Negative electricity prices – how many renewable energy systems will run through?, *Forschungsstelle für Energiewirtschaft (FfE)*, Munich, 21 October 2024, <https://www.ffe.de/en/publications/negative-electricity-prices-how-many-renewable-energy-systems-will-run-through/>, 2024.



- 1135 Kitzing, L.: Risk implications of renewable support instruments: Comparative analysis of feed-in tariffs and premiums using a mean-variance approach, *Energy*, 64, 495–505, <https://doi.org/10.1016/j.energy.2013.10.008>, 2014.
- Kitzing, L., Juul, N., Drud, M., and Boomsma, T. K.: A real options approach to analyse wind energy investments under different support schemes, *Applied Energy*, 188, 83–96, <https://doi.org/10.1016/j.apenergy.2016.11.104>, 2017.
- Koivisto, M., Gea-Bermúdez, J., Kanellas, P., Das, K., and Sørensen, P.: North Sea region energy system towards 2050: integrated offshore grid and sector coupling drive offshore wind power installations, *Wind Energy Science*, 5, 1705–1712, <https://doi.org/10.5194/wes-5-1705-2020>, 2020.
- 1140 Kölle, K., Göçmen, T., Eguinoa, I., Alcayaga Román, L. A., Aparicio-Sanchez, M., Feng, J., Meyers, J., Pettas, V., and Sood, I.: FarmConnors market showcase results: wind farm flow control considering electricity prices, *Wind Energy Science*, 7, 2181–2200, <https://doi.org/10.5194/wes-7-2181-2022>, 2022.
- 1145 López Prol, J., Steininger, K. W., and Zilberman, D.: The cannibalization effect of wind and solar in the California wholesale electricity market, *Energy Economics*, 85, 104 552, <https://doi.org/10.1016/j.eneco.2019.104552>, 2020.
- Low Carbon Contracts Company and Department for Energy Security and Net Zero (UK): FiT Contract for Difference Standard Terms and Conditions, Version 6, GOV.UK, 13 March 2024, <https://assets.publishing.service.gov.uk/media/65f1ab91133c221271cd3861/cfd-ar6-standard-terms-and-conditions.pdf>, condition 1.1: “Curtailement”, “Defined Curtailement Compensation” (revenues foregone including Difference Amounts); Part 11 (Curtailement), 2024.
- 1150 Mehta, M. K., Zaaier, M., and von Terzi, D.: Designing wind turbines for profitability in the day-ahead market, *Wind Energy Science*, 9, 2283–2300, <https://doi.org/10.5194/wes-9-2283-2024>, 2024.
- Meyers, J., Bottasso, C., Dykes, K., Fleming, P., Gebraad, P., Giebel, G., Göçmen, T., and van Wingerden, J.-W.: Wind farm flow control: prospects and challenges, *Wind Energy Science*, 7, 2271–2306, <https://doi.org/10.5194/wes-7-2271-2022>, 2022.
- 1155 Neuhoﬀ, K., May, N., and Richstein, J. C.: Financing renewables in the age of falling technology costs, *Resource and Energy Economics*, 70, 101 330, <https://doi.org/10.1016/j.reseneeco.2022.101330>, 2022.
- Newbery, D.: Eﬃcient renewable electricity support: Designing an incentive-compatible support scheme, *The Energy Journal*, 44, 1–22, <https://doi.org/10.5547/01956574.44.3.dnew>, 2023.
- Nielsen, J. S. and Sørensen, J. D.: Risk-based derivation of target reliability levels for life extension of wind turbine structural components, *Wind Energy*, 24, 939–956, <https://doi.org/10.1002/we.2610>, 2021.
- 1160 Pettas, V.: Wind turbine operational optimization considering revenue and fatigue damage objectives, Ph.D. thesis, University of Stuttgart, <https://doi.org/10.18419/opus-13959>, 2024.
- Pettas, V. and Cheng, P. W.: Down-regulation and individual blade control as lifetime extension enablers, *Journal of Physics: Conference Series*, 1102, 012 026, <https://doi.org/10.1088/1742-6596/1102/1/012026>, 2018.
- 1165 Pettas, V., Gräfe, M., Dirik, D. G., Ramaswamy, H. R., Riva, R., Quick, J., and Réthoré, P.-E.: A framework to evaluate structural fatigue accumulation and health metrics of wind farms under diﬀerent operational strategies, *Journal of Physics: Conference Series*, 3224, 032 079, <https://doi.org/10.1088/1742-6596/3224/3/032079>, proc. The Science of Making Torque from Wind (TORQUE 2026), 2026.
- Prokhorov, O. and Dreisbach, D.: The impact of renewables on the incidents of negative prices in the energy spot markets, *Energy Policy*, 167, 113 073, <https://doi.org/10.1016/j.enpol.2022.113073>, 2022.
- 1170 Requate, N., Meyer, T., and Hofmann, R.: From wind conditions to operational strategy: optimal planning of wind turbine damage progression over its lifetime, *Wind Energy Science*, 8, 1727–1753, <https://doi.org/10.5194/wes-8-1727-2023>, 2023.



- Réthoré, P.-E., Fuglsang, P., Larsen, G. C., Buhl, T., Larsen, T. J., and Madsen, H. A.: TOPFARM: Multi-fidelity optimization of wind farms, *Wind Energy*, 17, 1797–1816, <https://doi.org/10.1002/we.1667>, 2014.
- Rubert, T., Zorzi, G., Fusiek, G., Niewczas, P., McMillan, D., McAlorum, J., and Perry, M.: Wind turbine lifetime extension decision-making based on structural health monitoring, *Renewable Energy*, 143, 611–621, <https://doi.org/10.1016/j.renene.2019.05.034>, 2019.
- 1175 Ruhnau, O.: How flexible electricity demand stabilizes wind and solar market values: The case of hydrogen electrolyzers, *Applied Energy*, 307, 118 194, <https://doi.org/10.1016/j.apenergy.2021.118194>, 2022.
- Schlecht, I., Maurer, C., and Hirth, L.: Financial contracts for differences: The problems with conventional CfDs in electricity markets and how forward contracts can help solve them, *Energy Policy*, 186, 113 981, <https://doi.org/10.1016/j.enpol.2024.113981>, 2024.
- 1180 Seyyedeh-Barhagh, S., Khajeh, H., Laaksonen, H., Mohammadi-Ivatloo, B., and Hatzigiorgiou, N.: A comprehensive review of sector coupling as a strategy for power system flexibility, *Renewable and Sustainable Energy Reviews*, 239, 117 102, <https://doi.org/10.1016/j.rser.2026.117102>, 2026.
- Simley, E., Millstein, D., Jeong, S., and Fleming, P.: The value of wake steering wind farm flow control in US energy markets, *Wind Energy Science*, 9, 219–234, <https://doi.org/10.5194/wes-9-219-2024>, 2024.
- 1185 Stiewe, C., Xu, A. L., Eicke, A., and Hirth, L.: Cross-border cannibalization: Spillover effects of wind and solar energy on interconnected European electricity markets, *Energy Economics*, 143, 108 251, <https://doi.org/10.1016/j.eneco.2025.108251>, 2025.
- Thomson, R. C., Harrison, G. P., and Chick, J. P.: Marginal greenhouse gas emissions displacement of wind power in Great Britain, *Energy Policy*, 101, 201–210, <https://doi.org/10.1016/j.enpol.2016.11.012>, 2017.
- Đukan, M., Keles, D., and Kitzing, L.: The impact of two-sided contracts for difference on debt sizing for offshore wind farms, *The Energy Journal*, 46, 145–188, <https://doi.org/10.1177/01956574251331942>, 2025.
- 1190 Vad, A., Guilloré, A., Anand, A., Pettas, V., Shah, A. H., Lizarraga-Saenz, I., Aparicio-Sanchez, M., Eguinoa, I., Conti Gost, N., Tsaklis, I., Frère, A., Hermans, K. W., Kamau, J. K., Dimitrov, N., Göçmen, T., and Bottasso, C. L.: Modeling wind farm response: a modular, integrated, and multi-stakeholder approach, *Wind Energy Science Discussions*, pp. 1–49, <https://doi.org/10.5194/wes-2026-45>, preprint wes-2026-45, in review, 2026.
- 1195 Van Steenberghe, M. and Ovaere, M.: Market inefficiencies in renewable support policies: Evidence from offshore wind contracts for difference in Great Britain, Ghent University FEB Working Paper 2025/1124, revised version, June 2026, http://wps-feb.ugent.be/Papers/wp_25_1124_rev.pdf, 2026.
- Veenstra, A. and Mulder, M.: Impact of contracts for differences for non-carbon electricity generation on efficiency of electricity market, *Energy Economics*, 136, 107 674, <https://doi.org/10.1016/j.eneco.2024.107674>, 2024.
- 1200 Wernet, G., Bauer, C., Steubing, B., Reinhard, J., Moreno-Ruiz, E., and Weidema, B.: The ecoinvent database version 3 (part I): overview and methodology, *The International Journal of Life Cycle Assessment*, 21, 1218–1230, <https://doi.org/10.1007/s11367-016-1087-8>, 2016.
- Wheatley, J.: Quantifying CO₂ savings from wind power, *Energy Policy*, 63, 89–96, <https://doi.org/10.1016/j.enpol.2013.07.123>, 2013.
- Zheng, F., Vukobrat, A., Anand, A., Tamaro, S., Göçmen, T., and Bottasso, C. L.: Market-aware operation of grid-integrated wind power plants: beyond revenue, towards structural loads and grid services, *Journal of Physics: Conference Series*, 3224, 032 049, <https://doi.org/10.1088/1742-6596/3224/3/032049>, 2026.
- 1205 Ziegler, L., Gonzalez, E., Rubert, T., Smolka, U., and Melero, J. J.: Lifetime extension of onshore wind turbines: A review covering Germany, Spain, Denmark, and the UK, *Renewable and Sustainable Energy Reviews*, 82, 1261–1271, <https://doi.org/10.1016/j.rser.2017.09.100>, 2018.