

Appendix A. Methods and data

This appendix details the contract specifications, the measure of cash-flow risk, the project-finance model, and the data used.

A.1 Contract designs

We assess ten different contract specifications in total, corresponding to three groups: two-sided CfDs, one-sided CfDs, and financial CfDs. As a benchmark, we use the merchant case (no long-term contract).

Merchant case

As “merchant case” M we describe the situation that a generator has no long-term contract and receives revenue only from selling output in the day-ahead market. The revenue R of a generator i for an hour t is defined as the product of the generation $q_{i,t}$ and the electricity spot price p_t . For simplicity, we assume that for each hour there is a unique spot price, which can be interpreted as generators selling their production only on the day-ahead auction of the wholesale market. In other words, we ignore intraday markets and imbalance settlement, which in practice do affect a project's cash flow; since none of the contracts we compare alters the obligation to balance, including them would probably raise measured risk under all designs. The revenue over the course of year y is given by

$$R_{i,y}^M = \sum_{t=1}^{t=T} q_{i,t} p_t, \quad (1)$$

where the set $t = \{1, T\}$ refers to all hours in year y .

Two-sided CfD

We denote the two-sided CfD as $2CfD$ in equations. We decompose hourly revenue into market income and CfD payments. For a given hour t , the total revenue of a generator i is defined as the product of generation $q_{i,t}$ and the sum of the realized spot price (p_t) and the CfD payment, which depends on the difference between the strike price (S) and the reference price (p_t^{ref}). The total annual revenue is

$$R_{i,y}^{2CfD} = \sum_{t=1}^{t=T} q_{i,t} p_t + q_{i,t} (S - p_t^{ref}). \quad (2)$$

Under a classical “UK-style” two-sided CfD, the reference price equals the hourly spot price at which generators sell their electricity in the spot market ($p_t^{ref} = p_t$), such that the CfD payment simply offsets

deviations of p_t from the strike. This case constitutes our first specification. In this case there is only one source of variation, the *volume risk* of production $q_{i,t}$. The price risk disappears since the strike price remains the same for the duration of the contract. This contract leads to the so-called “produce-and-forget” incentives: the investor has only an incentive to maximize the volume of production $q_{i,t}$ regardless of the economic value of the energy produced (which depends on the time when the electricity was produced).

Because this variant of the contract is rarely considered by policymakers nowadays, we model a second specification using an annual reference price. The reference price is the market value of onshore wind, defined as the volume-weighted average price of onshore-wind-generated electricity across all wind parks in the country:

$$\bar{p}_y^w = \frac{\sum_1^T (Q_t^w p_t)}{\sum_1^T Q_t^w}, \quad (3)$$

where Q_t^w refers to the market wide generation of onshore wind power. When a reference price other than the hourly spot market is used, different incentives emerge. Because the reference is exogenous to a single plant but correlated with its output, it provides an imperfect but good hedge against price risk while keeping some of the wholesale price incentives alive. In hours when the spot price falls significantly below zero and the strike is not high enough to justify production ($S + p_t - p_t^{ref} \leq 0$), we assume that generators curtail all production ($q_{i,t} = 0$) to avoid losses.

Our third and last specification of the two-sided CfD suspends payments from governments to generators whenever spot prices turn negative. This incentivizes generators to always curtail their production at negative prices ($p_t \leq 0$). The European Commission's current guidance points in this direction. It prefers a two-sided CfD, that is, with some payments from the generator to the government, but it does not endorse the standard version. Instead, it requires that generators retain their incentives to bid into the day-ahead, intraday, balancing and ancillary-services markets, and recommends quarterly or annual rather than settlement-period reference prices (European Commission, 2025).

In sum, we assess three variants of the two-sided CfD: an hourly reference price, an annual reference price, and annual reference in addition to suspension during negative prices. In this order, incentives improve but hedging quality is suspected to deteriorate.

One-sided CfD

For the one-sided CfD (1CfD), we use the same decomposition into market income and CfD payments as above. The overall annual revenue a generator producing electricity under a one-sided CfD is

$$R_{i,y}^{1CfD} = \sum_{t=1}^{t=T} q_{i,t} p_t + q_{i,t} \cdot \max(S - p_t^{ref}, 0). \quad (4)$$

Under this contract, the price risk is only partially mitigated. In the extreme case where $p_t^{ref} \geq S$ for all t , the contract is identical to that of having no support, whereas in the opposite extreme case where $p_t^{ref} \leq S$ for all t , the contract is identical to that of a two-sided CfD.

As with the two-sided CfD, we consider three specifications:

- (i) using the hourly spot price p_t as the reference price p_t^{ref} ;
- (ii) using the annual market value of wind $\bar{V}_y^W$ (defined above in equation 3) as the reference price p_t^{ref} ; and
- (iii) using the annual market value of wind $\bar{V}_y^W$ while suspending payments when spot prices are negative.

As above, we assume generators always dispatch according to economic incentives, i.e. curtail if total revenue is negative.

Financial CfD

The third major design we consider is the financial CfD, as proposed by Schlecht et al. (2024). This contract consists of two separate hourly payments between generator and government: a fixed (lump-sum) payment from the government to the generator F in exchange for a variable (floating) payment from the generator to the government that reflects market conditions V_t . Effectively, the contract is a fixed-for-floating swap.

Unlike conventional CfDs, this contract is not tied to asset outputs, hence it has been described as “injection-independent” (as opposed to “production-based”). The Commission also contemplates this as a direction for future contract design (European Commission, 2025). For a generator with installed capacity K_i (in MW), the contracted capacity k_i can be defined as:

$$k_i = \alpha \times K_i, \quad (5)$$

Where $\alpha > 0$ is the contract factor.

At each hour t , the three revenue components are:

- The market income of the generator:

$$M_{i,t} = q_{i,t} p_t, \quad (6)$$

- The lump-sum payment from government to generator:

$$F = k_i A, \quad (7)$$

where A is the payment in €/MW agreed in the contract. It is somewhat analogous to the strike price S in production-based CfDs.

- The payment from generator to government:

$$V_t = k_i g_t^{ref} p_t, \quad (8)$$

where g_t^{ref} is the reference generation profile (a normalized capacity factor), in MWh/MW. It is zero when the price is negative, hence V_t payments are suspended.

The annual revenue of a generator under a financial CfD is

$$R_{i,y}^{f-CfD} = \sum_{t=1}^{t=T} (M_{i,t} + F - V_t) = \sum_{t=1}^{t=T} (q_{i,t} p_t + k_i A - k_i g_t^{ref} p_t). \quad (9)$$

As contract factor α and contracted capacity k approach 0, such that $F = V_t = 0$, the revenue and risk profile of this contract converge to the merchant case. Conversely, for a generator whose production follows exactly the adjusted reference generation $q_{i,t} = k_i g_t^{ref}$, we have $M_{i,t} = V_t$ and the hourly income equals the fixed payment F , so revenues become completely independent of price and volume fluctuations. This highlights that the remaining source of risk under a financial CfD stems from deviations of actual production from the reference profile ("basis risk").¹

Two parameters need to be specified to define the contract: the contract factor α and the reference profile g_t^{ref} (A is, like S , determined by an auction). We use the national hourly capacity factor of onshore wind as the reference profile:

$$g_t^{ref} = \frac{Q_t^w}{K_t^w} \quad (10)$$

where K_t^w is the total installed capacity for onshore wind in Germany at time t .

By varying the way in which we set the contract factor, we model four specifications of the financial CfD. In contrast to our specifications for two- and one-sided CfDs, which affect both dispatch incentives and cash-flow risk, the financial CfD specifications only change the level of basis risk and do not affect

¹ As with all contracts that entail a payment from generator to government, this design does not imply the necessity that it is a subsidy. for a given value of $A = \frac{1}{T} \sum_{t=1}^T (g_t^{ref} P_t)$ the generator would perceive, on average, only its own market income $M_{i,t}$, making this contract revenue neutral while still smoothing (or not) the cash flow.

dispatch: under this contract, generators always have an incentive to stop producing at negative prices. We therefore focus on how policymakers might choose the contract factor α_i to shape risk and support.

In our results, we consider the following four specifications of the financial CfD:

- (i) Baseline ($\alpha_i = 1$): All parks use their full installed capacity as contracted capacity.
- (ii) Fleet-optimal ($\alpha_i = \alpha^* = 0.525$): All parks share a common contract factor chosen to minimize average investment cost (equivalently, average downside risk) across our sample; the calculation of α^* is described in Section A.2.
- (iii) Technology-adjusted: Contract factors are park-specific and proportional to expected production levels. We approximate technology differences by each park's average annual capacity factor over the first three years after installation and scale these relative to the national capacity factor, obtaining a park-specific factor α_i^{tech} . This specification aims to align fixed payments and claw-backs with expected production, thereby mitigating adverse selection concerns.
- (iv) Park optimal ($\alpha_i = \alpha_i^*$): Finally, we allow for park-specific optimal contracts, where α_i is chosen specifically to minimize the LCOE for each project. The optimization procedure is detailed in Section A.2.

A.2 Method

Measuring cash-flow risk

We assess the cash-flow risk that investors face under each contract design. Cash-flow risk is a standard concept in the corporate finance and financial distress literature. As discussed in Section 2.1, many studies measure it using the standard deviation of earnings normalized by average sales, that is, the coefficient of variation (CoV) of cash flow. We follow this literature and use the CoV of annual revenues across the sample period as a normalized measure of cash-flow risk for each wind park and contract.

The coefficient of variation (CV) is defined as

$$CV_i = \frac{\sigma(R_{i,y}^c)}{\mu(R_{i,y}^c)}, \quad (11)$$

where $R_{i,y}^c$ is the annual revenue in year y under contract c expressed in €/MW of installed capacity, $\sigma(\cdot)$ denotes the standard deviation, and $\mu(\cdot)$ is the mean of all years. Normalizing by average revenue allows us to compare risk across parks with different sizes and production profiles.

We simulate hourly revenues and aggregate them to the annual level before computing CoVs. Aggregation to annual revenues removes intra-year seasonality and focuses on the time horizon relevant for meeting debt service and other yearly obligations (Alcorta et al., 2024). We interpret each observed year as a distinct and equally likely state of the world, so the CoV summarizes the dispersion of revenues across these states. This provides a simple, comparable measure of overall cash-flow risk, which we complement with a project finance model focused on downside risk in the next subsection.

Financial model for downside risk

While the CoV is a measure of risk, policy makers and investors ultimately care about how such risk affects debt-to-equity ratios, the cost of capital, and LCOE. In project finance practice, lenders size debt based on low-percentile cash flows or minimum debt service coverage ratios (DSCRs), rather than on average volatility alone. In this spirit, we develop a stylized financial model that maps the simulated annual revenue distribution under each contract into feasible capital structures and implied financing costs (WACC). The purpose is to formalize a simple and transparent mechanism through which cash-flow risk affects those indicators.

Consider a representative wind project with a lifetime of $Y = 30$ years. Investment costs C^{fix} are incurred upfront, while annual operating and maintenance costs C^{om} are constant over time. Annual revenues $R_{i,y}$ are uncertain and depend on the spot market and the long-term contract, as described in Appendix A.1, with each observed year interpreted as an equally likely and independent state of the world within the context of this project finance model.

Project cash flows in each year are allocated in the following order:

1. Operating and maintenance costs, assumed to be fixed and paid with certainty;
2. debt service, consisting of a constant annual payment; and
3. equity payouts, the residual cash flow.

Projects are financed through debt and equity with fixed and exogenous returns of 1.15% and 10%, respectively (see Table 2). We assume no reserve accounts and no refinancing. Debt is modeled as a fixed annuity loan with maturity equal to the project lifetime. Let D denote total debt, and r_d the cost of debt. Annual debt service is given by:

$$DS = D \cdot \frac{r_d(1 + r_d)^Y}{(1 + r_d)^Y - 1}. \quad (12)$$

The solvency condition we impose is that debt service must be fully covered in every simulated year, i.e.,

$$R_{i,y}^c - C^{om} \geq DS \quad \forall y. \quad (13)$$

This corresponds to a DSCR of at least one in all states. While this assumption is stricter than some standard project finance benchmarks, it allows us to directly link revenue volatility to maximum feasible leverage in a transparent way. For a given contract design c and wind park i , this condition defines the maximum sustainable debt level $D_{i,c}^{max}$. Intuitively, contracts that raise the minimum annual cash flow permit higher leverage, which reduces WACC.

The remaining investment cost is financed with equity. Let total investment cost be $I_i = C^{fix} * K_i$. Equity is then $E_{i,c} = I_i - D_{i,c}^{max}$. Equity holders receive the residual cash flow after debt service and O&M. Their participation constraint requires the expected present value of dividends to cover the equity investment,

$$\sum_{y=1}^Y \frac{\mathbb{E}[\Pi_{i,c,y}]}{(1+r_e)^y} \geq E_{i,c}. \quad (14)$$

where $\Pi_{i,c,y}$ denotes annual dividends and r_e is the required return on equity. We assume equity holders bear the residual risk, so we evaluate their participation in expectation, in contrast to debt which is constrained by the worst realization.

In equilibrium, competitive auctions for support contracts imply that projects bid such that they just break even. For each wind park and contract design, we jointly determine:

- the maximum feasible debt level $D_{i,c}^{max}$,
- the implied equity requirement $E_{i,c}$, and
- the support level as defined by contract parameters (strike price or fixed payment)

such that both the debt solvency condition and the equity participation constraint are just binding. In other words, competition for contracts ensures that equity investors do not earn excessive profits. This yields an implied LCOE for each contract and project, reflecting both the distribution of annual cash flows from the contract and the endogenous financing structure. Contracts that reduce downside cash-flow risk allow projects to carry more debt, substituting cheaper debt for expensive equity, thus lower the weighted average cost of capital, and hence have a lower LCOE.

For financial CfDs, we apply the same procedure but, as discussed in Appendix A.1, in some specifications we also optimize over the contract factor α_i to study how different parameter choices affect leverage and LCOE. For the fleet-optimal specification we solve the model for common contract factors from 0.35 to 1.25 in steps of 0.025 and select the factor with the lowest average LCOE across the 63 parks ($\alpha^* = 0.525$). For the park-optimal specification we solve, for each park, for the fixed payment and contract factor combination that maximizes debt capacity subject to the break-even condition, which minimizes the cost of electricity.

A.3 Data

A key strength of the study is the use of real production data from operational turbines rather than simulated or synthetic profiles. Real output is shaped by constraints that are often unobserved or hard to model, so using realized feed-in lets us observe them directly and produce more realistic estimates of revenue volatility. The data also capture heterogeneity that simulations miss: not only weather and price years, but site quality and turbine technology.

The core dataset consists of turbine-level production records for 381 turbines, aggregated into 93 wind parks, in Germany. These turbines vary in size and age, with an average rated power of 2.1 MW; the oldest was commissioned in December 1995 and the newest in May 2024. To obtain a long, balanced panel, we exclude parks entering operation after 2014 and remove parks with missing or persistently very low production (below an hourly average of 0.06 MWh per MW of rated power). The resulting sample covers a ten-year period (2014–2023) over which all parks are continuously observed, and includes major market events such as the low prices of Covid-19 and the high prices of the 2022–2023 energy crisis, alongside a wide range of weather years. The final estimation sample contains 268 turbines across 63 parks with a combined capacity of about 510 MW.

The simulations use three additional time series: hourly electricity prices from the ENTSO-E Transparency Platform; potential-generation feed-in profiles from Netztransparenz, used to build the national hourly production profile for financial CfDs; and installed capacity from Open Power System Data and ENTSO-E, used to normalize feed-in into capacity factors.

Although uniquely detailed, all assets are operated by a single company and are geographically concentrated in north-eastern Germany. Figure A1 shows the location, commissioning year, and installed capacity of the 63 parks, highlighting both regional clustering and heterogeneity in turbine vintages.




Figure A1. Location, commissioning year (color), and installed capacity (marker size) of the wind parks in the sample.

To gauge representativeness, we compare each park's hourly normalized production profile with the national onshore-wind profile. Correlations range from about 0.3 to 0.8, with a mean near 0.6 (Figure A2), which indicates substantial diversity in production patterns. Figure A3 shows how the sample sits against the market value of onshore wind over the period. Most parks earn slightly less per MWh than the national average. Four earn more: two by less than 2 €/MWh and two by about 17 and 20 €/MWh. The sample therefore includes both kinds of project, though not in equal numbers, and the results in Section 4.2 should be read with that in mind. We regard this heterogeneity as representative of a history of projects being built with production-based support contracts, yet informative of real-world investment opportunities.




Figure A2. Histogram of correlations between each park's hourly normalized output and the national onshore-wind profile; the vertical line marks the sample mean.




Figure A3. Distribution of capture prices across the 63 wind parks, measured as merchant revenue per MWh of potential output over 2014–2023. The dashed line marks the market value of onshore wind over the same period (56.9 €/MWh), computed on the national feed-in profile under the same assumption about negative-price hours. Four parks lie above it.

For each turbine we observe hourly power feed-in, on-site wind speed, and a log of outages with cause and duration (grid constraints, maintenance, permitting, marketing decisions), including both full and partial curtailment. On average, turbines produce about 91% of their potential output, with roughly 9% curtailed for operational or grid-related reasons, and grid-induced curtailment accounts for the largest single share (Figure A4).




Figure A4. Distribution of lost generation (as a share of potential output) by outage category. Bars show the average curtailed generation for each cause; whiskers indicate the central 90% (5–95%) of the distribution.

In Germany, energy curtailed for grid reasons is typically compensated as if it had been delivered, since such constraints are outside investors' control. To keep parks comparable and to match this compensation regime, all simulated revenue streams are based on potential generation, that is, the output that would have occurred without grid-related curtailment, while retaining other causes of curtailment.