Curtailment Risk on the Cape Corridors

How much wind energy will the grid refuse, and what does that do to a project's finances?

A scenario-based screening assessment built from public data — not a project-specific forecast, and not a lender-grade grid or energy assessment

Worked example throughout: a hypothetical 100 MW curtailable Karoo wind project (illustrative location only)

Prepared by

Michael Pietersen (PhD)

StatForge · michael@statforge.co.za

Based on 1 000 resampled simulations drawn from 32 distinct historical weather-years, at hourly detail · fully reproducible

Drafted 2026-07-31 · weather data: 32 years of measured ERA5 weather (1985-2016)

Read this first: what this report is for, the three futures, and the words we use

HOW TO USE THIS REPORT — AND HOW NOT TO

Use it as: (1) a strategic screening tool showing how curtailment, compensation and grid delivery interact; (2) a scenario framework that an IPP, lender or adviser can populate with confidential project data; (3) a structured map of exactly which data project due diligence must obtain. Do NOT use it as: a project-specific curtailment forecast, or a lender-grade independent assessment. The model represents two large, electrically distinct Cape networks as a single aggregate constraint, its 2026 anchor is a reasoned assumption, and its financial layer is deliberately simplified — every project-level percentage in this document is illustrative of mechanism, not predictive of a site. Corrections, challenges and data contributions are actively invited: michael@statforge.co.za.

What this report is about. “Curtailment” means a wind or solar farm is instructed to reduce output it could otherwise have produced, because the energy cannot be transported through the network or absorbed by the wider system at that time. The energy is lost; whether the money is lost too depends on a compensation framework whose pot is fixed while the losses are not — which turns out to be the heart of the financial question (Section 6). This report puts numbers — and, more importantly, ranges — on that risk for projects connecting in the Western and Eastern Cape, where nearly all of South Africa's wind fleet sits (NTCSA, 2025).

The one thing nobody knows, made explicit. Whether curtailment gets better or worse depends almost entirely on one question: will absorptive capacity in the Cape grow fast enough by 2028, as new projects connect? Instead of guessing, every result in this report is shown under clearly labelled futures — two bookends and a middle path:

THE THREE FUTURES USED THROUGHOUT THIS REPORT

“Design-track” (best case) — the grid is upgraded as officially planned. The grid company's framework for new connections implies that, by 2028, the network can be expected to handle the new build with only ~4% of new entrants' energy curtailed (NTCSA, 2025). In this future, the upgrades and operational fixes needed to make that true actually arrive on time.

“Slow-track” (middle path) — the grid grows, but at the pace the grid company has actually demonstrated rather than the pace the plan requires. In its most recent financial year the NTCSA built 271 km of transmission lines against its own 423 km target — a 64% delivery ratio (Eskom & NTCSA, 2026) — so this future delivers 64% of the total effective relief the design case requires by 2028. To be precise: it is not a claim that 64% of the necessary kilometres of wire get built. The parameter being scaled is effective absorption capacity, which bundles wire with transformers, local load, stability limits and operating practice — the line-delivery ratio is used as a transparent, published proxy for delivery performance overall, not as a physical conversion. The fraction is an adjustable input, and this is arguably the most realistic of the three paths.

“Constrained” (worst case) — the grid stays as it is today, but the build-out proceeds anyway. “The build-out” is not our guess: it is the grid company's own published connection list — every project with its megawatts and expected connection year, growing the corridor from 3,866 MW operational to 7,205 MW by 2028, including the 1,580 MW admitted specifically under curtailable terms (NTCSA, 2025; the full list is data/capacity_pipeline.csv in the model). Connecting into a stalled grid is possible precisely because those new entrants signed curtailable connection agreements. Grid capacity remains at its (already strained) 2026 level: nothing extra breaks, nothing gets fixed. A stress pairing rather than a forecast — in reality, sustained severe curtailment would delay some of that pipeline, pushing outcomes toward the slow-track.

The gap between the bookends is not a modelling nuisance — it IS the risk. Within this model's aggregate representation, for the design-track future to hold, the corridor's implied carrying-capacity parameter must grow by roughly 70–90% in two years (from 3 121 MW to 5 592 MW — explained in Section 3; the range reflects uncertainty in how bad 2026 really is). Two honesty notes on that headline: it is a statement about a model parameter, not a physical build target — operational fixes, storage, demand shifts and reconfiguration all count toward it, so it does not mean the NTCSA must build ~80% more wire; and the slow-track exists because the delivery record argues against assuming either bookend — lines are being built, but at two-thirds of the targeted rate, while the compensation backlog (~R2 billion, delays near a year: Engineering News, 2026) shows what falling behind costs.

A note on how results are stated. Averages hide the years that hurt. Results are therefore given as a typical year (P50 — half of years are better, half worse) and a bad year (P90 — only 1 year in 10 is worse). A convention warning for project-finance readers: in energy-yield studies "P90" usually denotes an exceedance level of generation; here P90 always means the 90th percentile of the ADVERSE quantity — a 90th-percentile loss, a 10th-percentile revenue. A lender's decision rests on the bad years, which is exactly what a single assumed percentage cannot describe. Where a chart shows a spread of outcomes, bar heights are the share of simulated years landing in that bar — the bars in a panel add up to 100% — so "this bar is at 15%" reads simply as "about 1 year in 7 looks like this". One honest caveat of using real weather: the 1 000 simulations resample only 32 distinct historical years, so the effective sample is 32 corridor-years. P50 and P90 are well supported by that sample; anything rarer than roughly 1-in-30 is not, which is why this report does not headline P99-type extremes — the worst simulated outcome simply means "the worst weather actually observed since 1985".

Words used in this report

Term	Plain meaning
Curtailment	Being ordered to reduce output the plant could have produced, because the network is constrained or the wider system cannot absorb the energy during low-demand periods. The energy is always lost; how much revenue is lost depends on the compensation rules (Section 6).
Grid capacity / "headroom" (X)	How much power the system can absorb from the region in any hour, in MW. "Effective" for two reasons: it is worked out from curtailment outcomes rather than engineering drawings, and it bundles the wire together with demand and the system's ability to soak up energy in low-demand hours — at low national demand, renewables can be curtailed even with wire to spare (Section 3). Because it is a bundle, a change in X is not the same as an equal change in demand (Section 4).
The Cape corridors	The transmission routes out of the Western and Eastern Cape — home to almost the entire SA wind fleet, and where the grid is most congested.
NTCSA / NERSA / Eskom	The state grid company / the energy regulator / the utility. The GCCA is the grid company's published study of how much new generation can connect, where (NTCSA, 2025).
Curtailment tier	1,580 MW of new connections (2026–28) allowed ONLY on condition they accept being curtailed. Older projects connected with "firm access" and no such condition (NTCSA, 2025).
The "4%"	Not a cap on actual curtailment. It is the assumed curtailment level used to size a fixed compensation pot in Eskom's five-year revenue plan, and excess costs may not be recovered from consumers (NERSA, 2025). If the corridor loses more than 4%, every claim is paid only in proportion — at 8% corridor curtailment the pot covers roughly half of each claim.
Deemed energy compensation	Being paid for curtailed energy as if it had been delivered. Under the 2025 framework ALL curtailment-affected generators — public programme and private alike — are eligible, but always from the capped pot above; hence today's ~R2bn backlog (NERSA, 2025; Engineering News, 2026).
Wheeled contract	A private deal: the plant sells to a distant buyer, renting the grid to deliver. Curtailment compensation works by crediting the buyer's account as if the energy had been supplied, funded from the same capped pot (Cliffe Dekker Hofmeyr, 2025) — so wheeled projects are covered too, up to the cap. Anything beyond it is a private negotiation.
Merchant	Selling into the market with no long-term buyer. Same capped compensation eligibility; market-price risk comes on top (not modelled until the new wholesale market produces price data).

P50 / P90 / P99	Typical year / bad year (1-in-10) / very bad year (1-in-100).
DSCR (debt service cover ratio)	The year's cash available for the loan, divided by the year's loan payments due. “Available cash” means revenue minus operating costs (maintenance, insurance, staff, admin — modelled at 2% of build cost per year), before loan payments; tax is not modelled, and that is stated. Below 1.0 the project cannot pay its lenders from its own cash flow.
Monte Carlo simulation	Running the model over very many simulated weather-years (here 1 000) to see the full range of outcomes, not just the average.

1. Executive summary

Curtailment was already financially material in 2025, and escalated sharply in 2026. The hard numbers are verified: regulatory monitoring records R402–404 million in "deemed energy" compensation to 67 renewable plants in 2025 (NERSA, 2026, pp. 14 & 18 — the report itself quotes both figures), up from R46 million in the first half of 2024, with the curtailment occurring "due to low demand around midnight hours" and for grid stability (p. 14). Converting that cost into a curtailment percentage is an inference, not a measurement, and this report says so — but every input is now a regulator-published number. The numerator: dividing by the published average wind price of R1.42/kWh (Table 5, p. 14; midnight curtailment falls almost entirely on wind) gives ≈284 GWh curtailed — less if higher-tariff plants sit among the 67 claimants. The denominator: the wind fleet delivered 11,040 GWh in 2025 (Table 5, p. 14; cross-check: 3,388 MW of corridor wind × the 37% fleet load factor on p. 10 × 8,760 hours ≈ 11.0 TWh); adding back the curtailed energy gives potential output. The ratio lands at 2.2–2.5%; we calibrate at the upper end of that range and label it a proxy, because the R404m is not disaggregated by technology, province or cause — it spans the whole portfolio and includes low-demand (night-time) curtailment, not only Cape grid congestion, though night-time curtailment falls almost entirely on wind (NERSA, 2026; all caveats in Appendix A). For 2026 the evidence establishes escalation, not a rate: about R2 billion of claims sat in verification by mid-year — several times what a pot sized on 4% implies for a full year — and curtailed volumes reportedly ran an order of magnitude above 2025 (Engineering News, 2026; Yelland, 2026). No realised 2026 percentage has been published anywhere. This report therefore works with a stated assumption band of 5%–9%: bounded below by the pot-oversubscription evidence (2026 clearly ran well past 4%), bounded above by the superseded 10% planning ceiling, and carried through every result as a range rather than a point.

The headline, stated within the model's terms: the official plan implies the corridor's carrying capacity must grow roughly 70–90% in two years. Reproducing the 2026 assumption band requires an implied carrying-capacity parameter of 3 121 MW (band: 2 938 MW–3 325 MW; Section 3 explains how this is inferred, and why it is a model parameter rather than a measured grid property). For the official ~4% design intent to hold once the announced 2028 build-out is connected, that parameter must reach 5 592 MW. Growth in the parameter is not a construction target — operational fixes, storage and demand all count toward it — but whether the equivalent relief arrives is the single question that decides every number below:

Outcome in 2028 (typical year)	Design-track	Slow-track (64% delivery)	Constrained
Corridor-wide wind energy curtailed	0.9%	5.0%	20.3%
Our example project's energy curtailed*	4.1%	21.5%	43%
Its annual revenue (central contract case†)	R296 million	R282 million	R194 million
Cash available vs loan payments (DSCR)	1.32 — safe	1.24 — safe	0.79 — cannot pay
Share of years, simplified annual DSCR < 1.0	0.0%	0.0%	100%

*Table 1 — The three futures, quantified. Project figures are illustrative of mechanism, not site forecasts. *Under the "new-entry-first" sharing hypothesis — plants connected under curtailable-access arrangements absorb reductions before firm-access generators (Section 5). †Compensated from the capped pot, paid late — an illustrative fixed-budget convention consistent with the NERSA conditions, not a verified settlement rule (Section 6).*

What this means in plain terms — for a hypothetical curtailable Karoo project. If the grid upgrades arrive, the reference project is financeable on normal terms: about 4.1% of its energy curtailed in a typical year, 4.6% in a bad one, with simplified annual DSCR below 1.0 in 0.0% of simulated years. If the grid stays where it is while everyone connects anyway, the same project loses 43% of its energy — and although the framework entitles it to compensation, the pot is sized for a 4% world: at constrained-future loss rates it covers only about 21% of each claim (Section 6). The middle path is the one to study hardest: on the grid company's demonstrated delivery pace, the corridor hovers near the 4% funding threshold (typical year 5.0%), and the project still loses 21% of its energy under the contractual sharing rule — but under the pot-pro-rata reading the pot funds about 86% of each claim, holding the simplified DSCR near 1.24 — while under the harsher reading, in which each project is compensated only up to 4% of its own energy, the DSCR falls to 1.06 (Section 6 shows both; the published rules do not settle which applies). Specifically, on the middle path the project's viability rests not on the weather but on the compensation machinery — how the capped pot is divided, and whether it pays — which is exactly where diligence and contract negotiation should concentrate. The honest statement of the risk is: which future are you on, what tells you early, and what does each cost?

The finding that survives every caveat in this report. The specific percentages above are illustrative — they depend on an aggregate network representation, an assumed 2026 anchor, a hypothesised sharing rule and a simplified financial layer, all stated where used. What does not depend on any of them is the structure they reveal: this project class's economics are dominated by three variables that ordinary financial models conceal inside a single curtailment percentage — the pace and location of grid relief, the rule used to allocate congestion reductions among generators, and the funding and timing of deemed-energy compensation. Any serious diligence process should demand data on those three before trusting anyone's point estimate, including this report's.

Two practical findings fall out of the distributions. First — within the constrained future — the compensation terms move revenue by more than the weather does. To be precise about a claim that sounds extreme: the wind still earns every rand (no wind, no revenue), but its year-to-year swings around the mean are modest — inside any one contract, the gap between a calm year and a windy one (P10 to P90) is about R19 million of 2028 revenue. Across contracts, holding the weather identical, expected revenue runs from R171 million (no effective cover) through R195 million (the actual capped regime) to R293 million (fully funded) — a span of R122 million that is pure contract-and-budget mechanics. Which piece of paper you hold matters several times more than which weather you get — with one stated caveat: the weather spread is estimated from the 32 historical years in our sample (1985-2016), which cannot contain extremes rarer than the record itself; the bundled fetch script can extend the sample back to 1979 (the satellite era) for a longer view, so the comparison is fair for ordinary-to-bad years, conservative for the unprecedented ones. Second, batteries are not a universal fix: a 25 MW / 100 MWh battery recovers about 99% of curtailed energy in the design-track future but only about 29% in the constrained one, because curtailment then arrives in long spells (averaging 9.9 hours, worst over 100 hours) that fill any affordable battery early (Section 7).

2. Why a single assumed percentage fails

Standard practice prices curtailment as one fixed percentage in the financial model, with a “downside case” a little higher. That worked when curtailment was a footnote. It fails now, for three reasons. First, curtailment happens hour by hour — it is the coincidence of renewable output with insufficient network or system absorption capacity — so the yearly total swings with the weather, and the years that break a loan are precisely the swings. Second, the corridor is filling up: the same weather that causes modest losses at today's connected capacity causes severe losses at the 2028 build-out, so recent history understates what is coming. Third, when the corridor is curtailed, somebody specific loses — and the rule deciding who is not public. Operating wind farms report they cannot reconstruct why they were curtailed on a given day (Agbaje, Oni & Ibeh, 2026; Cliffe Dekker Hofmeyr, 2025). Any single project-level number therefore hides an unstated guess about that rule. This report keeps the three problems separate and explicit: the physics is a calibrated hourly simulation (Section 3–4), the sharing rule is a set of labelled scenarios (Section 5), and the contract is an overlay on top (Section 6). If a reader disagrees with a result, they can point at the exact assumption they disagree with — and change it.

3. How the model works, and how it is tied to measured facts

The engine. For every hour of a simulated year, the model asks: how much wind and solar power are the corridor's plants producing, and how much can the grid carry away? Whatever exceeds the grid's carrying capacity in that hour is curtailed. Production comes from hourly weather for three plant clusters applied to the actual list of connected and coming projects — taken project by project, with megawatts and start dates, from the grid company's published connection study, growing from 3,866 MW operational to 7,205 MW by 2028 (NTCSA, 2025). That list is predominantly wind but not exclusively: the corridor also carries about 900 MW of solar PV by 2028 (350 MW operational, plus Droërivier's 250 MW and its 220 MW curtailment-tier allocation), and the model dispatches wind and solar together against the same constraint, because that is how the network experiences them. Solar is roughly 9% of corridor energy but has an outsized effect on when the corridor binds (Section 4). This is repeated over 1 000 different weather-years, so we see not one answer but the whole spread — calm years, windy years, and the windy-year-in-a-full-corridor combinations that produce the losses lenders care about.

Why three clusters, and how plants are assigned to them. The connection study lists every project by substation, and the substations group naturally into three geographic wind climates: the Overberg / south coast (Agulhas, Bacchus — driven by the summer south-easter), the Karoo interior (Komsberg, Roggeveld, Nuweveld, Droërivier — a different regime, strongest at night), and the Eastern Cape (Grassridge, Neptune/Pembroke). Plants inside one cluster see nearly the same weather at the same time; clusters partially diverge from each other — and that partial divergence matters, because corridor congestion is worst precisely when all three peak together. Every project's cluster assignment is recorded in the model's input table (data/capacity_pipeline.csv), traceable row by row; for the oldest Western Cape plants the split between coast and Karoo is an estimate and is graded as such in the assumptions register. Finer geographic detail is possible but would outrun the resolution of the rest of the model — the constraint itself is one number, so three weather regions is the matched level of detail.

The number “X”, in plain terms — and what it is not. The Western and Eastern Cape are not one export pipe: they contain many corridors, substations, thermal and voltage limits, and the real system binds nodally. This model deliberately collapses all of that into one aggregate constraint, X — the right level of detail for corridor-scale screening from public data, and the model's principal simplification (stated again in Section 8). The model does not directly assume a numerical value for X; it infers an equivalent aggregate constraint from the selected curtailment anchors: X is set so that the model, run at a given year's connected capacity, reproduces that year's anchor. But the network representation itself remains a simplifying assumption, X inherits the full uncertainty of the anchors it is solved from, and it is best read as an implied-equivalent constraint parameter — not a measurement of physical transfer capacity. That bundling is less crude than it may sound, because the official curtailment methodology also treats connection capacity as a combined function of network limits and local load — a point worth quoting directly, since it underwrites the whole approach.

Two curtailment mechanisms, one bucket — a stated simplification. Renewables get switched off for two physically distinct reasons. Congestion: too much energy for the wire — corridor-specific, the story this report centres on. System oversupply: at low demand, inflexible coal cannot back down below its minimum stable level, so renewables are curtailed REGARDLESS of wire capacity. The record says the second mechanism is not hypothetical — it is currently prominent: the regulator describes 2025's curtailment as occurring “due to low demand around midnight hours” (NERSA, 2026, pp. 14, 18), the grid company's own modelling projects over 5 TWh a year of dumped energy by 2028 from oversupply (Eskom & NTCSA, 2026), its winter outlook shows forecast peak demand FALLING (~27 GW against ~30 GW a year earlier), and the connection study flags slow Western Cape load growth as a risk to its own 4% design (NTCSA, 2025). This model does not separate the two mechanisms: X absorbs both, so it is best read as effective ABSORPTION capacity — wire, local demand and the system's flexibility floor bundled together — and a stagnating or falling demand trend tightens X just as surely as missing wire does. Two consequences are worth stating plainly: the 2025 anchor is not a pure congestion observation — it includes a material low-demand oversupply component — and none of the three futures models demand explicitly — a demand-stagnation path would push outcomes toward the constrained end even if lines are built on schedule. Separating the mechanisms requires a residual-demand layer (national hourly demand minus wind and solar against the fleet's flexibility floor), which is the model's next structural upgrade (Section 8). The regulator draws the same line, which is a confirmation rather than a complication: the approved

4% service covers CONGESTION curtailment, while over-frequency and other low-demand interventions continue under existing power purchase and connection agreements (NERSA, 2025). Two distinct roles for demand follow, and they should not be confused. Local demand is an INPUT to the capacity calculation — it determines how much generation a substation can host, and is therefore already inside the 4% design case (quoted below). National low-demand oversupply is a separate curtailment MECHANISM, sitting outside the 4% service. Only the second is missing from this model — and it is the reason the 2025 anchor, which is all-cause, and the design intent, which is congestion-specific, are not quite the same population.

The grid company's own connection-capacity methodology states, in its assumptions:

> *"The reduction in the load forecast due to the downturn of the economy will lower the available capacity at substations as less of the generation will be absorbed by the local load and more of it will have to be exported to the transmission network."*
>
> — Eskom Transmission, GCCA-2023 Phase 1 (June 2021), §2 Assumptions, p. 2

And in its methodology for the first level of the capacity hierarchy:

> *"When a generator is connected to the secondary busbar in a substation, the power generated is first absorbed by the local load and the excess is fed upstream through the transformers. … The capacity of a substation's transformation is fixed, but the substation load varies throughout the day. This means that the lower the load connected at a substation, the lower the generation capacity that can be connected at the substation secondary busbar."*
>
> — Eskom Transmission, GCCA-2023 Phase 1 (June 2021), §3.1 Level 1: Transformer Capacity, p. 3

Why this matters for reading every number in this report. South African connection capacity is not a wire rating. It is a hosting-capacity calculation in which local demand absorbs generation first, only the surplus is exported, and the binding limit is the lowest of the transformer, substation-transfer, local-area and supply-area constraints. Demand is therefore already inside the official 4% design intent — which is precisely what this model's X approximates, and why treating X as effective ABSORPTION rather than as a wire is the faithful reading rather than a convenient simplification.

The solve, step by step. (1) Take the anchor year's connected fleet — for 2025, the operational megawatts in the connection list — and build its hourly generation across hundreds of simulated weather-years using the weather engine. (2) Pick a trial value of X. (3) For every hour, whatever generation exceeds X is curtailed; summing over the year and dividing by potential energy gives that trial's modelled curtailment rate. (4) If the modelled rate is above the anchor, X is too small — raise it; if below, lower it; halve the interval and repeat (bisection) until the modelled rate matches the anchor to within 0.05 percentage points. (5) X(2025) matches the 2.5% proxy at the 2025 fleet; X(2026) matches the assumption band's midpoint at the 2026 fleet, and the same solve at the band's 5% and 9% edges produces the range carried through the headline. Because the rate is averaged over hundreds of weather-years, X is matched to the expected rate, not to one lucky or unlucky year.

Three solves pin the parameter down:

Solve	Anchored to	Result
X(2025)	The 2025 proxy anchor of 2.5% (inferred from the R404m deemed-energy cost: NERSA, 2026; caveats in Appendix A)	3 400 MW
X(2026)	The 2026 working-assumption band of 5%–9% (reasoned bracket — see below; solved at both edges and the midpoint)	3 121 MW
X(design)	The official ~4% design intent for new connections, at the full 2028 build-out (NTCSA, 2025)	5 592 MW

Table 2 — The model's implied-equivalent constraint parameter, solved three ways. Each solve inherits the evidential weight of its anchor — proxy, assumption, and official design intent respectively.

Two observations come from the solves themselves, before any simulation of the future. First, within the model's aggregate representation, reproducing the assumed 2026 band requires a lower equivalent parameter than reproducing the 2025 proxy (3 121 MW vs 3 400 MW): connected capacity grew only ~8% between the anchors while assumed curtailment roughly tripled, and a model calibrated to 2025 alone yields just 4.3% for 2026. What that reflects physically is not identifiable from public data — it could be network outages, tighter operating margins, stricter constraint application, low demand, system-wide oversupply curtailment mixed into the claims, or differences in how the two anchors were constructed. It does not necessarily mean physical transfer capacity fell. Second, the distance from X(2026) to X(design) — roughly 80% growth in the parameter in two years — is the size of the equivalent relief the official framework quietly assumes.

What the 2026 evidence establishes, and what it does not. The public record firmly establishes escalation: the grid company acknowledged ~R2 billion of curtailment claims in verification by mid-2026 (Engineering News, 2026), curtailed volumes in H1-2026 reportedly ran an order of magnitude above the whole of 2025, and some projects reported revenues ~9% below budget (Yelland, 2026; Agbaje, Oni & Ibeh, 2026). None of that is a curtailment rate. The 9% is a revenue shortfall blending lost energy with reimbursement delays and project-specific factors; the "ten times" compares claim volumes across different periods and portfolios, not like-for-like annual rates; and the sector's standard statistics cannot help: CSIR publishes generation net of curtailment (CSIR, 2025), and the programme's own quarterly reports through December 2025 contain no curtailment or deemed-energy data at all — checked document by document (IPP Office, 2025–26). Our 5–9% band is therefore a reasoned working assumption, and here is the reasoning. The floor: a compensation pot sized on 4% of corridor energy implies very roughly R0.5–0.7 billion of claims per year at fleet tariffs — the ~R2 billion in verification by mid-year means 2026 ran well past the 4% level. The ceiling: the 10% figure sometimes quoted is the modelling ceiling of the superseded 2024 planning study, not an observation, and the system kept operating. We solve X against both edges, so the headline is a range — required growth of 68.2% (if 2026 ends at 5%) to 90.3% (if at 9%) — and the conclusion survives the whole band. The single document that would replace this assumption with a measurement is the NTCSA's six-monthly implementation report to NERSA, which the approval conditions require and which reportedly served at NERSA's March 2026 committee (Energize, 2025): it must contain implemented curtailment levels, incident records, connected capacity and costs. Obtaining it is the highest-value data acquisition available to this model, and recalibrating when it lands is a one-line input change.

Will the grid actually grow? The official record cuts both ways — and it is now quantified. Against timely delivery: in the financial year to March 2026 the NTCSA built 271 km of transmission lines against its own 423 km target, citing contractor financial constraints and underperformance; its Transmission Development Plan requires an average of roughly 1,450 km per year (peaking above 2,100) against current local construction capacity it assesses at ~800 km per year; and transformer lead times run 36–48 months (Eskom & NTCSA, 2026). The corridor's effective capacity also moved the wrong way during 2026 (Section 3), the compensation pot is already oversubscribed with ~R2 billion in claims under verification (Engineering News, 2026), and the framework window itself ends in March 2028, so the regime governing the out-years is not yet written. In favour: transformer commissioning is ahead of target (4,000 MVA delivered against 3,750 in FY26), more than half the land the TDP needs has been secured (7,700 km of servitudes), contractor-incubation and independent-transmission-project mechanisms are being stood up (Eskom & NTCSA, 2026), NERSA designed its approval conditions expressly to incentivise acceleration (Energize, 2025), and the required growth is in "effective" capacity — operational fixes and re-dispatch count toward it, not only new wire. This model does not forecast which future wins — but it prices both, and it names the early indicators: each NTCSA six-monthly curtailment report, the annual line-kilometre delivery figures against TDP targets, and whether the MYPD-period compensation allocation is topped up.

The weather. Wind does not blow randomly hour to hour: windy spells last, windy days cluster into windy years, and the three Cape clusters often peak together — which is exactly what fills a corridor. Each cluster's long-run output is anchored to a regional target — 36% of maximum output for the Overberg, 38% for the Karoo, 35% for the Eastern Cape, 26% for solar. Two honesty notes about those targets. They are calibration assumptions derived from published fleet-level evidence (CSIR, 2025; Lloyd's Register & SAWEA, n.d.), regionalised by resource-quality judgement — no public dataset reports per-region fleet output, so they are graded as assumptions in Appendix A, not measured statistics; plant-level generation data mapped to the three regions would convert them into measurements, and is the second-most valuable data upgrade this model has (Section 8). And published fleet output is net of historical curtailment, so anchoring to it can embed a sliver of past grid constraint into the weather engine — at the ~2.5% curtailment levels of the calibration period the distortion is under one percentage point, and it is largely absorbed when the grid capacity X is solved against the same period's measured curtailment. In this run the weather is 32 years of measured ERA5 reanalysis weather (1985-2016; Hersbach et al., 2020), fetched via the bundled script. Two adjustments are applied and disclosed. First, a single reanalysis grid point is not a wind farm — reanalysis smooths terrain and real turbines stand on the best hills — so each region's wind speeds are scaled by one solved factor so that long-run fleet output matches the published figures above; the measured data keeps what it is uniquely good at (the hourly, seasonal and year-to-year structure of real weather) while the output level is anchored to measured fleet performance. Second, the archive is restricted to the 1985-2016 window: diagnostics revealed a step change in interior wind speeds aligned exactly to calendar 2017 — strongest at the terrain-sensitive Karoo point, absent at the coast — which is the signature of a data-source change, not climate; mixing the two regimes would distort the bad-year tail, so the long internally-consistent segment is used. A note on how the resampling works: each simulated year draws one whole historical year, taken for all three regions simultaneously. This is deliberate — windy years are windy corridor-wide, and that coincidence is exactly what fills the wire — so the effective pool is 32 corridor-years, not 32 per region. Mixing regions or months across different years would enlarge the pool arithmetically while quietly destroying the correlations that create bad years, thinning precisely the tail a lender needs to see.

4. The corridor outlook: what the whole region loses

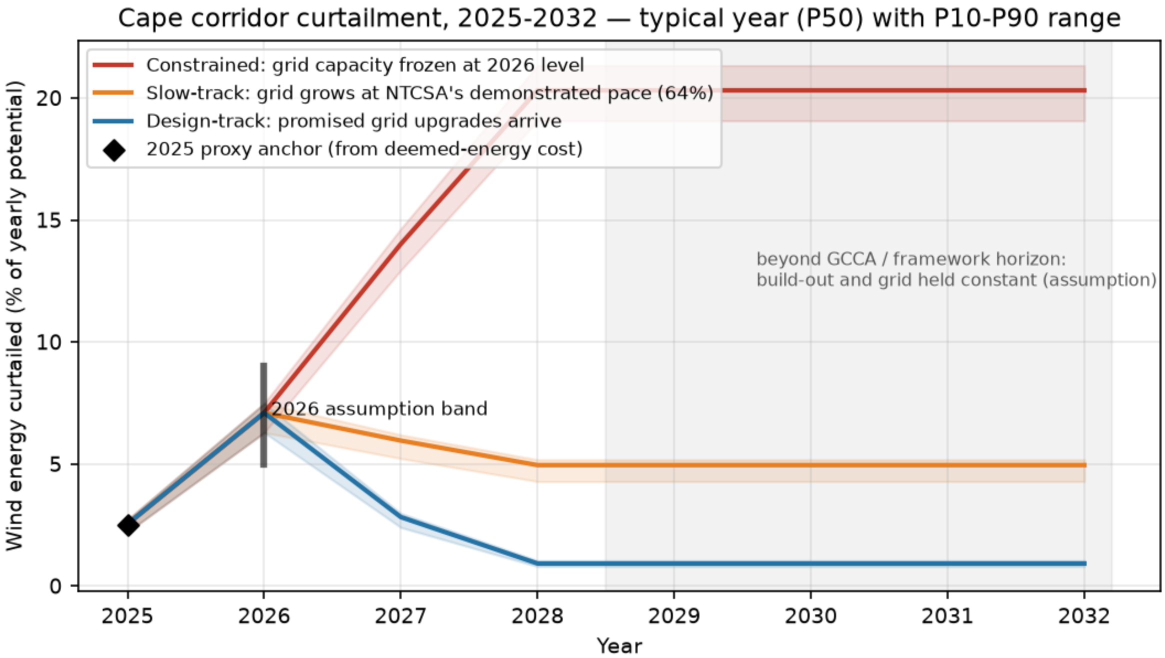

Figure 1 — Corridor-wide wind curtailment under the three futures. Solid line: typical year (P50). Shaded band: the P10–P90 range of years. Black diamond: the 2025 proxy anchor. Vertical bar: the 2026 assumption band. Grey region: beyond the GCCA/framework horizon — build-out and grid held constant by assumption.

Wind energy curtailed (corridor-wide)	2025	2026	2027	2028	2030	2032
Design-track, typical year	2.5%	7.1%	2.8%	0.9%	0.9%	0.9%
Design-track, bad year (P90)	2.7%	7.5%	3.0%	1.0%	1.0%	1.0%
Slow-track, typical year	2.5%	7.1%	6.0%	5.0%	5.0%	5.0%
Slow-track, bad year (P90)	2.7%	7.5%	6.2%	5.2%	5.2%	5.2%
Constrained, typical year	2.5%	7.1%	14.0%	20.3%	20.3%	20.3%
Constrained, bad year (P90)	2.7%	7.5%	14.6%	21.3%	21.3%	21.3%
Relief-2031 sensitivity, typical year	2.5%	7.1%	2.8%	0.9%	0.9%	0.0%

Table 3 — 2025 and 2026 are calibrated to the anchors, so the futures only diverge from 2027. "Relief-2031" adds further grid growth from 2031 on top of the design-track.

Why the lines flatten after 2028 — read this before trusting them. The flat right-hand side of Figure 1 is a held assumption, not a forecast, and the charts shade it grey to say so. Both the connection study and the curtailment framework itself run out in 2028 (the framework window closes in March 2028: Energize, 2025), and beyond that nothing is committed: no approved connection list, no approved grid works, no compensation regime. Rather than invent them, the model freezes each future's generation-versus-grid balance at its 2028 state. That is more coherent than it may look — under the current mechanism new connections are only unlocked where grid studies create room, so a future with no new grid is also a future with few new connections, and a future where the grid grows to plan is one where connections grow with it. But it does mean the post-2028 numbers answer a conditional question — "what does each year look like if the 2028 balance persists?" — and the year-to-year wiggle you would expect from a live system (new projects, new lines, rule changes) is deliberately absent. The relief-2031 sensitivity shows one alternative path; a new GCCA or the post-2028 framework, when published, would replace this assumption the same way the six-monthly report replaces the 2026 band.

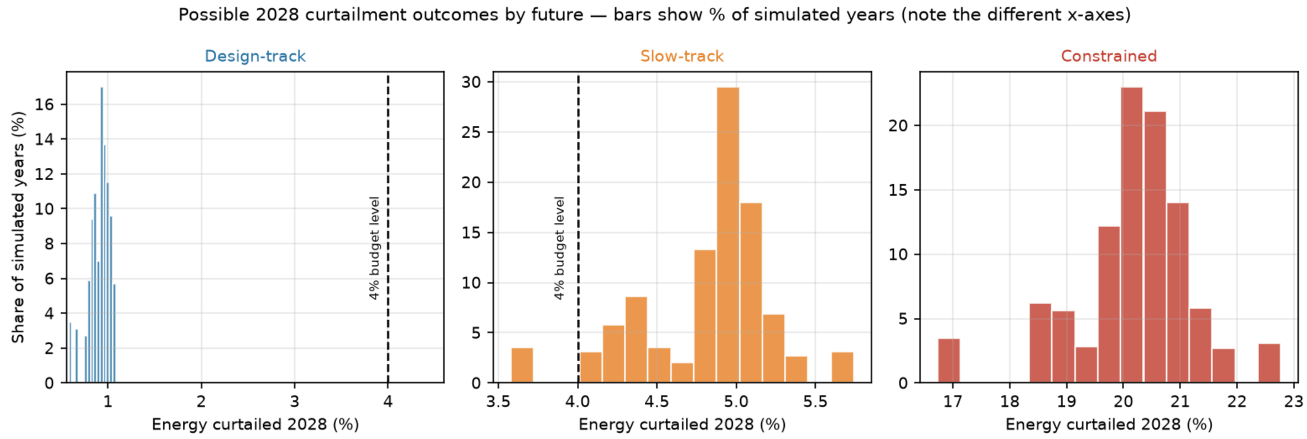

Figure 2 — The spread of possible 2028 outcomes under each future. Bar heights are the share of simulated years (each panel sums to 100%). Note the different x-axes: the futures live on different scales, which is itself the finding. The dashed line is the 4% budget-sizing level, shown where it falls in range.

Read Figure 2 against the dashed 4% line, remembering what that line is: the assumption used to size a fixed compensation pot, not a limit on curtailment (NERSA, 2025; NTCSA, 2025). In the constrained future the entire distribution lies far beyond it — meaning the pot would be oversubscribed every single year, the existing payment backlog (~R2 billion in verification, with delays approaching a year: Engineering News, 2026) would compound, and the compensation every project in the corridor relies on would be diluted in exactly the years it is needed most (Section 6 prices this dilution). In the design-track future the corridor sits comfortably below the line. The 4% figure is thus best understood as a funding-shortfall trigger, and this model prices how often it trips. One external cross-check on the scale: the grid company's own system modelling projects that nationally "dumped" (excess) energy can exceed 5 TWh a year by 2028 if the planned solar build materialises (Eskom & NTCSA, 2026) — the same order of magnitude as this model's constrained-future corridor estimate of ~4.3 TWh. The mechanisms differ (system-wide oversupply versus corridor congestion), so the figures are not additive — but the operator's modelling shows that multi-TWh annual non-absorption of renewable energy already appears in published planning scenarios, and is therefore not inherently an extreme scale.

How much does curtailment move if the system's ability to absorb energy changes? Section 3 established — in the grid company's own words — that connection capacity is a hosting calculation in which local demand absorbs generation first and only the surplus is exported. Demand therefore sits inside X, and a persistent change in it moves absorption by a near-constant block of megawatts in every hour rather than during one window. That makes one question cleanly answerable, and it is worth being exact about which one. The model can answer: IF effective corridor absorption changes persistently by N megawatts, how does curtailment respond? It cannot answer: what happens if a particular industrial plant closes. The second question needs the pass-through — how much of a national demand change actually reaches this corridor, through merit order and network flows — and that ratio is not estimated anywhere in this model. So what follows is a response curve, not a scenario: the model supplies the response, and the reader supplies the mapping from their own view of demand, storage or operations to megawatts of corridor absorption.

The response, and the shape of it. Holding the slow-track grid and the weather draws fixed, a typical 2028 year moves from 2.0% corridor wind curtailment with 600 MW of additional absorption, to 5.0% at the baseline, to 9.3% with 600 MW less. Two features matter more than the levels. First, the response is steep: around the baseline each 100 MW of lost absorption adds roughly 0.7 percentage points of curtailment, so a change equivalent to a few percent of the corridor's absorption is enough to decide which side of the 4% funding threshold it lands on — grid construction is not the only variable that empties the compensation pot. Second, the response is convex and asymmetric: losing absorption costs about 0.7 points per 100 MW while gaining it saves only about 0.5, and the adverse side steepens as losses deepen. Adverse movements therefore compound with grid under-delivery rather than merely adding to it — the same asymmetry that runs through this whole report.

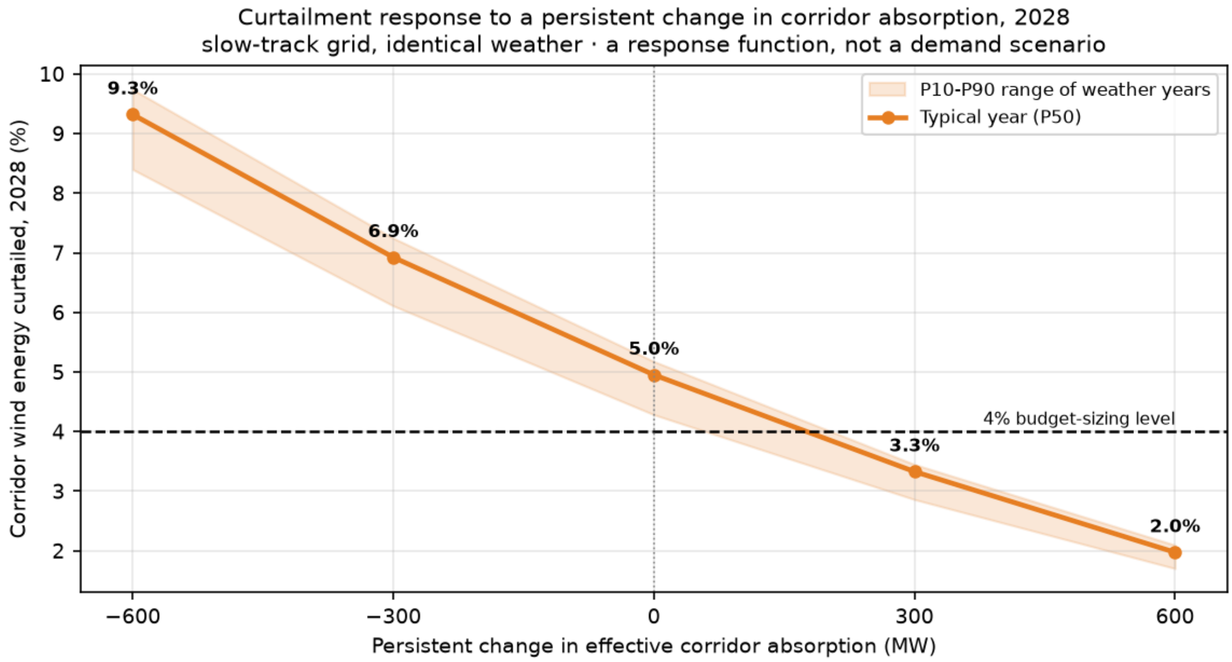

Figure 3 — Curtailment response to a persistent change in effective corridor absorption. Slow-track grid and weather draws held identical, so the curve is the pure absorption effect; shaded band is the P10–P90 range across weather years. The horizontal axis is the model's own parameter, NOT a demand quantity: the drivers that could move it — regional demand, national demand transmitted through the export balance, local storage or flexible load, operating practice — are not distinguished, and no pass-through from any of them is estimated. Read it as a response function: supply your own view of how many megawatts of corridor absorption a scenario removes, and read the curtailment off the curve (Appendix A 7.6, Grade C).

How to use it, and what would sharpen it. For scale, the range tested spans roughly ±13% of the corridor's modelled 2028 absorption. A transmission planner or a client's own load study can convert a demand, storage or operating scenario into a megawatt figure on this axis far more credibly than this model can from public data — which is exactly the division of labour intended. Two upgrades would replace the reader's judgement with the model's: a residual-demand layer, which would separate system-oversupply curtailment from corridor congestion and let demand vary explicitly rather than through X; and a nodal representation, which would replace the single corridor constraint with the substation-level structure the connection study actually uses. Until then, the curve states a relationship, not a forecast.

A counterfactual worth noting: solar drives the binding hours, wind pays for them. The model's 2028 curtailment is concentrated around the middle of the day — 30% of curtailed energy falls between 10:00 and 14:00 — yet 93% of the curtailed energy is wind. To test whether corridor solar actually causes that midday binding rather than merely coinciding with it, the year was rerun with corridor PV removed: total curtailment falls from 1012 GWh to 695 GWh — a 31% reduction — and the midday share of what remains drops from 30% to 22%, with the balance shifting to the windy overnight hours (15% → 22%). So midday solar output materially increases the number of binding events, while most curtailed megawatt-hours are allocated to wind under the pro-rata sharing rule. For a wind project the practical implication is uncomfortable and rarely priced: a material part of its curtailment risk is created by other developers' solar build-out, over which it has no control and may have limited contractual recourse.

5. From the corridor to one project: who eats the loss?

Sections 3–4 say how much energy the corridor as a whole loses. But a lender finances one project, and the rule that decides how the corridor's loss is shared among projects is not published. Affected wind farms say they cannot tell why they were curtailed on a given day (Agbaje, Oni & Ibeh, 2026; Cliffe Dekker Hofmeyr, 2025). No honest model can pretend to know this rule. Instead, we compute the project's loss under four clearly labelled possible rules, and report the range. One point of honesty about the "new-entry-first" rule before reading the table: the observed record already contradicts any idea that incumbents go untouched — the 2024–25 deemed-energy compensation flowed to EXISTING projects (NERSA, 2026), because oversupply curtailment and congestion predate the new framework. New-entry-first describes the contractual priority the new framework's curtailable connection agreements create at the margin, not current observed practice — it is the concentrated end of the range, pro-rata the dispersed end, and reality will sit between them. The width of the range is itself information: it measures how much the market's opacity costs, and it is the single number that a project's own connection agreement and substation data would narrow most.

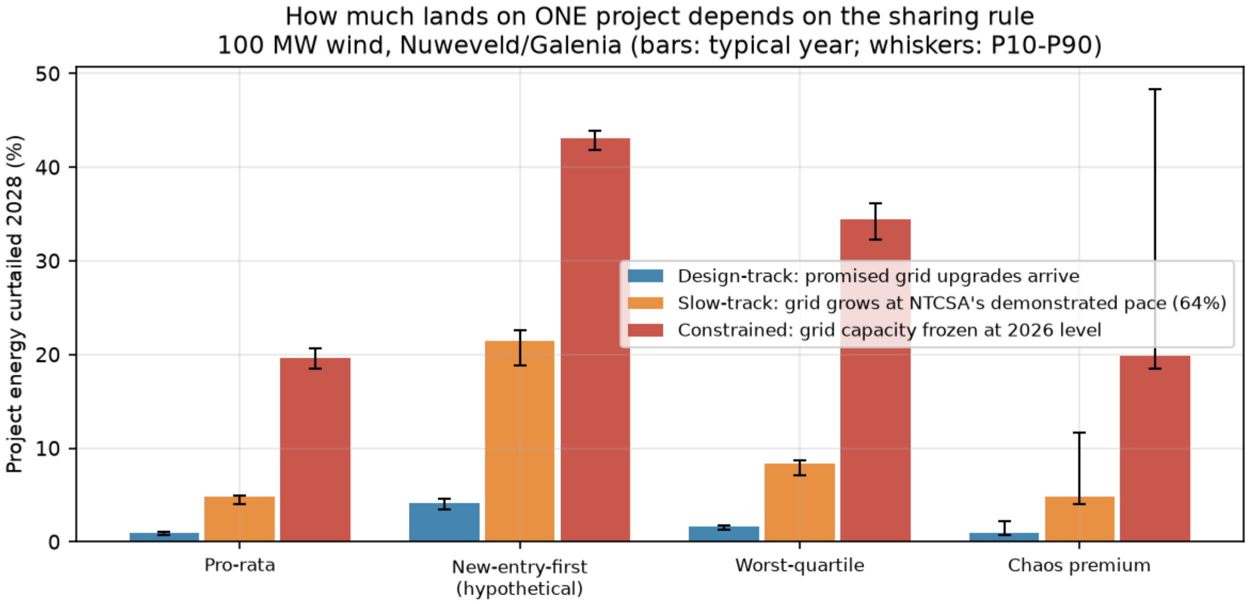

Figure 4 — The same corridor loss, shared four ways: what lands on the 100 MW Nuweveld/Galenia example in 2028, under each future.

Sharing rule (2028: typical / bad year)	Design-track	Slow-track	Constrained
Pro-rata — everyone loses the same fraction: if the corridor loses 5%, every plant loses 5% of its own output	0.9% / 1.0%	4.8% / 5.0%	20% / 21%
New-entry-first (hypothetical) — plants connected under curtailable-access arrangements absorb reductions before firm-access generators. Plausible from the connection framework's structure, but NOT yet confirmed against actual connection or curtailment agreements	4.1% / 4.6%	21.5% / 22.6%	43% / 44%
Worst-quartile — badly placed plants lose ~1.75× the corridor average	1.6% / 1.8%	8.4% / 8.8%	34% / 36%
Chaos premium — identical to pro-rata in a typical year BY CONSTRUCTION; what it adds is a tail: a 15% yearly chance of being singled out for 2.5× losses (opacity/governance risk). Read this row against the bad-year number, not the typical one	0.9% / 2.2%	4.9% / 11.7%	20% / 48%

Table 4 — "New-entry-first" is used as the reference reading in Sections 6–7 because the example project connects under the curtailment tier, whose connection agreements accept curtailment while firm-access incumbents' do not — but it remains a hypothesis until the applicable agreements and operating procedures are sighted. Curtailment instructions may in practice also depend on electrical effectiveness and system-security conditions.

6. Money: revenue and the loan

The wire does not care who you sold your power to — but the compensation rules decide who pays for what the network or wider system could not absorb, and those rules have a precise shape. Under the approved framework, every generator financially hurt by curtailment is eligible for compensation — public programme and private wheeled projects alike (for wheeled deals, the buyer's account is credited as if the energy had been supplied). But the pot is capped: it sits inside Eskom's five-year revenue determination, sized on the assumption that curtailment runs at 4%, and costs above the cap may not be recovered from consumers (NERSA, 2025, reported in Energize, 2025; mechanism per Cliffe Dekker Hofmeyr, 2025). How the capped pot is divided among claimants is NOT publicly specified, and the results below therefore show TWO readings side by side. Reading one (pot pro-rata, our convention): the annual pot equals the cost of 4% corridor curtailment and eligible claims are paid pro-rata from it — at 8% corridor curtailment each claim receives roughly half (4 ÷ 8); at 20%, roughly a fifth. Reading two (per-project cap, the harsh reading): each project is compensated only up to 4% of its OWN energy — “your allocation is 4%, that is all you get” — and losses beyond it are simply unfunded. Nothing in the published rules excludes the harsh reading, and one arithmetic fact argues it may be the intended one: if every project is capped at 4% of its own energy, aggregate claims can never oversubscribe the 4%-sized pot — it is exactly the rule a budget-constrained administrator would write. Actual settlement may instead depend on verification order, tariffs, allocation periods, contractual terms and regulatory true-ups, and whether unfunded amounts survive as receivables is unknown — reconstructing the real cash-flow waterfall from the NERSA decision, the practice note and actual agreements is a precondition for any investment-grade use of these figures. Payments also arrive late — about R2 billion is currently stuck in “verification and settlement”, with delays approaching a year (Engineering News, 2026) — and money that arrives a year late is worth less to a project that must pay its loan now. The model prices all three effects: the pro-rata convention, the delay, and the loss itself.

The results below therefore show a central case bracketed by two bounds. Central: compensated from the capped pot, paid late — the framework as it actually operates. Upper bound: compensation always paid in full and on a normal lag — the framework as intended, always topped up. Lower bound: no effective cover at all. The money assumptions are deliberately ordinary and fully adjustable (Appendix A): a selling price of R0.90 per kWh, building cost of R20 million per MW, 70% of the project financed with a 15-year loan at 11%, operating costs of 2% of build cost per year. Revenue scales in a straight line with the selling price, so replacing our price with a real term sheet changes the level of every number but none of the risk conclusions.

THE POT, IN NUMBERS — how a project can lose 21% of its energy yet have ~86% of its claim funded (slow-track, 2028)

The pot is sized to pay for 4% of the CORRIDOR's energy. In a typical slow-track year the corridor loses 5.0% — total claims exceed the pot, and every claim is paid the same fraction: roughly 4 ÷ 5.0 ≈ 81% in that year. Averaged over all simulated years (including sub-4% years, when claims are paid in full), the funded share is about 86%.

Why is our project's loss (21%) so much bigger than the corridor's (5.0%)? Because the losses are concentrated, not shared: under the new-entry-first hypothesis the curtailment tier — about 22% of corridor energy — absorbs nearly all of it, while firm-access incumbents lose almost nothing. Check the average: 22% of the corridor losing ~21% plus the rest losing almost nothing ≈ the corridor's 5.0%. This is exactly why the pot can pay most of a heavily-hit project's claim: total claims corridor-wide equal the corridor's 5.0% — not anyone's 20% — so the 4%-sized pot funds ~86% of every rand claimed, including every rand of our project's large claim, months late.

The two percentages answer different questions: 21% is how much energy this project loses; 86% is what share of each rand claimed the pot can pay UNDER THE PRO-RATA READING. A project can then be simultaneously heavily curtailed and largely compensated — until the corridor as a whole deteriorates and the pot dilutes for everyone, which is exactly the constrained-future mechanism (funded share ~21%).

THE HARSH READING CHANGES EVERYTHING for a heavily-hit project: capped at 4% of its own energy, our project is paid on 4 of the 21 points it loses — roughly 19% of its loss covered instead of ~86%. On the slow-track, expected

2028 revenue drops from R285 million (pro-rata reading) to R248 million (harsh reading), and the simplified DSCR from 1.24 to 1.06. Which reading applies is a question for a lawyer and the actual agreements, not for a weather model — and it is worth roughly R37 million a year to this project. Also note the interaction that softens the harsh reading: it hurts precisely when losses are concentrated (new-entry-first); under pro-rata sharing the project's own loss sits near the corridor rate and the two readings nearly coincide.

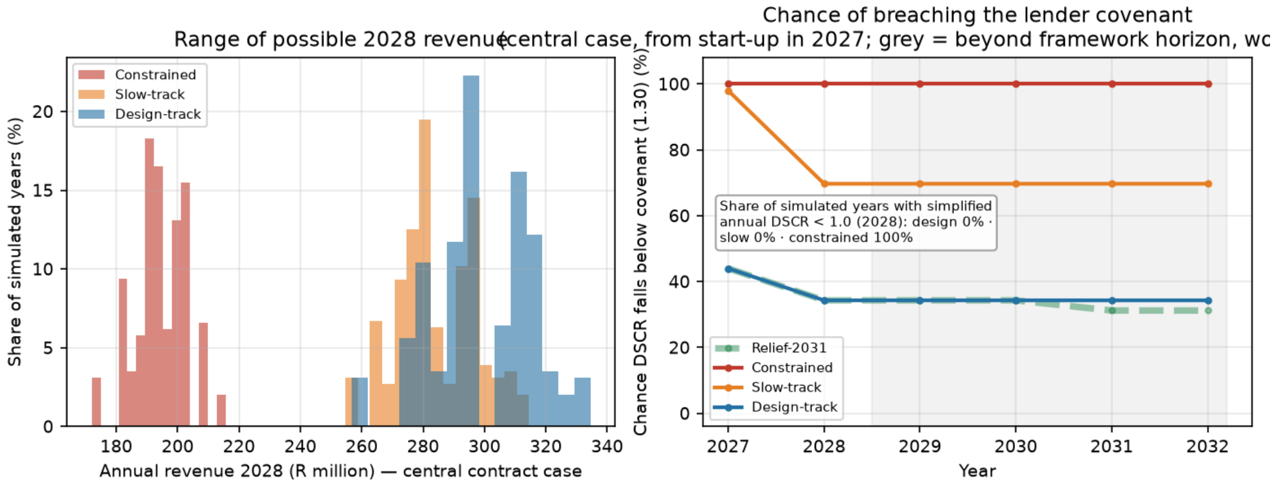

Figure 5 — Left: the spread of possible 2028 revenues under each future (central case: compensated from the capped pot; bars are the share of simulated years). Right: the chance of breaching the 1.30 lender covenant — the level at which loan agreements typically lock up distributions and trigger consent rights. Outright default (DSCR < 1.0) is annotated: compensation keeps it near zero on the design- and slow-tracks and near certainty in the constrained future.

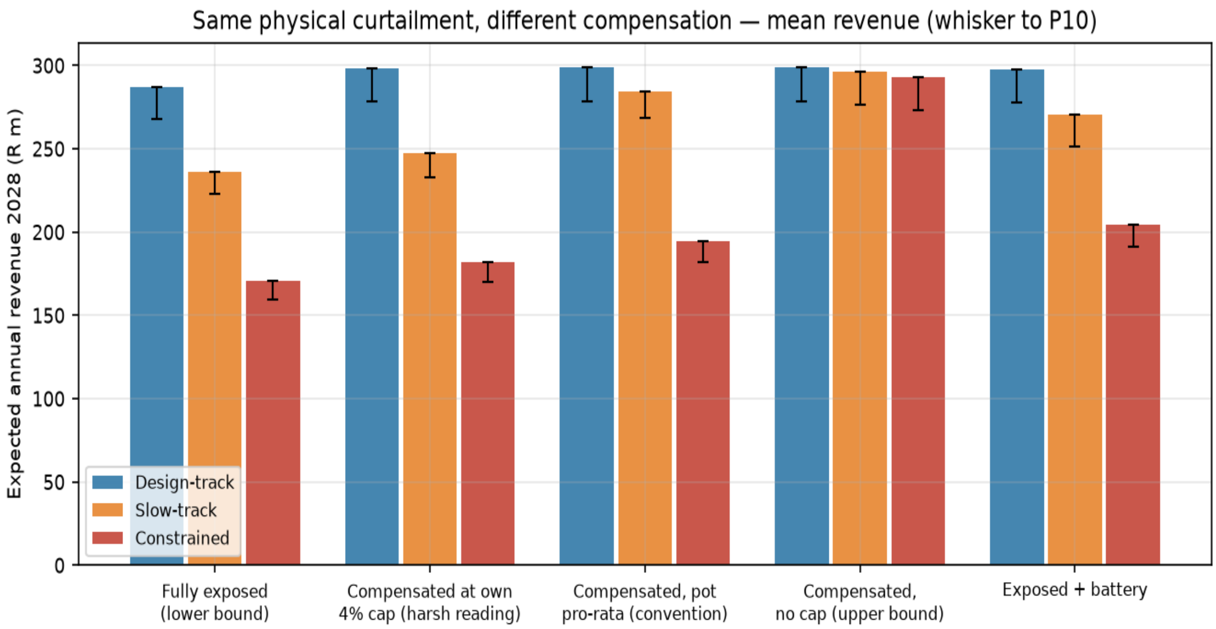

Figure 6 — Same weather, same wire, different compensation: expected 2028 revenue, with the whisker showing a bad (1-in-10) year.

2028, example project	Expected revenue	Bad-year revenue (P10)	DSCR, typical year	Share of years, DSCR < 1.0
Compensated in full (upper bound) — constrained	R293 million	R273 million	1.29	0.0%
Compensated, pot pro-rata (convention) — constrained	R195 million	R182 million	0.79	100.0%

Compensated at own 4% cap (harsh reading) — constrained	R182 million	R170 million	0.73	100.0%
No effective cover (lower bound) — constrained	R171 million	R160 million	0.67	100.0%
No cover + battery — constrained	R205 million	R191 million	0.84	100.0%
Compensated, pot pro-rata (convention) — slow-track	R285 million	R268 million	1.24	0.0%
Compensated at own 4% cap (harsh reading) — slow-track	R248 million	R233 million	1.06	12.5%
Compensated, pot pro-rata (convention) — design-track	R299 million	R279 million	1.32	0.0%
Compensated at own 4% cap (harsh reading) — design-track	R298 million	R278 million	1.31	0.0%

Table 5 — Compensated rows price the settlement delay (6 months; 12 when the pot is oversubscribed) plus the applicable cap. Under the pro-rata reading the pot covers on average ~21% of each 2028 claim in the constrained future and ~100% on the design-track; under the harsh reading each project is paid at most 4% of its own energy regardless. Which reading applies is unresolved in the public record — the spread between the paired rows measures what resolving it is worth.

Two readings of Table 5 matter. In the design-track future, the compensation question is almost moot — losses are small enough that even the uncovered case remains financeable. In the constrained future, the framework's promise and its funding part company: a project is entitled to compensation, but the pot covers only about 21% of each claim, so the central case sits much closer to the uncovered bound than to the fully funded one — and the project cannot reliably service its debt despite being "compensated" on paper. Financeability in that future therefore hinges on something that does not currently exist: a mechanism that tops up the pot, or a creditworthy private counterparty contractually absorbing what the pot does not. That is the precise, quantified form of the bankability alarm now circulating in the industry (Agbaje, Oni & Ibeh, 2026): projects are not uncompensated — they are compensated from a pot that arithmetic dilutes precisely when they need it most.

7. Can a battery fix it?

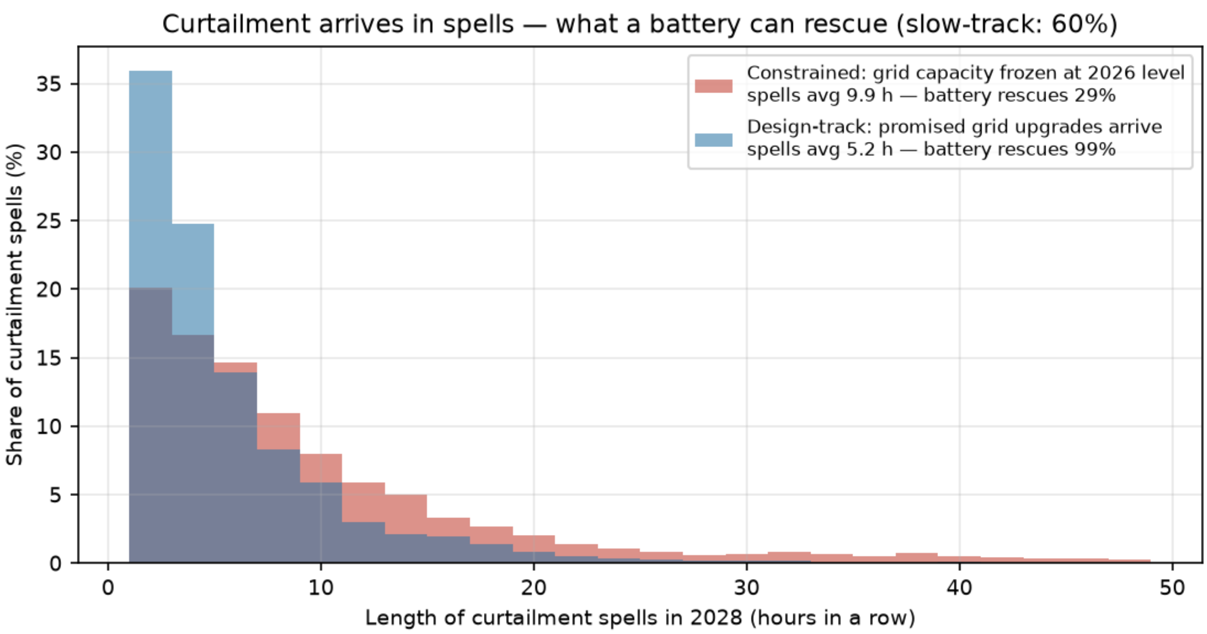

Figure 7 — How long curtailment spells last in 2028, and the computed share of curtailed energy a 25 MW / 100 MWh battery can rescue.

Partially — and whether "partially" means 90% or 23% depends on the future you are in, not on the battery. One scope note first: this section measures physically recoverable energy, which is a necessary input to a storage business case but not the business case itself — capex, degradation, cycling limits, grid-access rules and market settlement can make technically recoverable energy uneconomic to recover. The battery result here is computed hour by hour from the simulation, not assumed: the battery charges whenever the project is being curtailed and sells the energy back once effective absorption headroom returns. In the design-track future, curtailment comes in short spells (about 5.2 hours on average) with plenty of clear hours between them, and the battery rescues about 99% of the lost energy — storage effectively neutralises the residual risk. In the constrained future the corridor is curtailed for about 2 971 hours a year in spells averaging 9.9 hours (the worst exceed 100 hours). The battery fills up early in every long spell and sits full while the losses continue: it rescues only about 29%. The practical point for anyone sizing storage: the decision depends on the length of curtailment spells — which only an hourly simulation reveals — not on the average annual percentage.

8. What this model cannot tell you (and what would sharpen it)

This assessment is built entirely from public information, and its honesty about that is deliberate. (1) The weather is currently measured (ERA5). (2) The grid is modelled as a single carrying capacity for the corridor, not a wire-by-wire network study; differences between substations enter through the sharing rules of Section 5. A transmission engineer's review is a stated precondition before investment-grade use. (3) The sharing rule is genuinely unknown — the scenario range measures that, and one project's actual curtailment logs (every affected plant must lodge claims within 48 hours, so the logs exist) would replace scenarios with measurement. The model is built to ingest exactly that data. (4) The rules themselves can change — caps, budgets and compensation mechanics sit inside regulatory determinations that get revised, and the current curtailment framework's approval window itself ends in March 2028, so the out-years run on a regime not yet written (NERSA, 2025; NTCSA, 2025). (5) The 2026 anchor is a provisional, part-year estimate band, because per-connection figures are not public — which is why the headline is stated as a range across that band. (6) The regional output targets the weather is anchored to are Grade-B assumptions regionalised from national fleet evidence — plant-level generation and curtailed-energy data mapped to the three regions would turn them into measured pre-curtailment capacity factors (and would also validate the composite power curve's upper tail, which drives congestion hours). (7) The financial layer is deliberately thin: annual revenue less flat operating costs against a level annuity — no tax, reserves, escalation, degradation, sculpting or intra-year timing — which is why this report says "share of years with simplified annual DSCR below 1.0", never "default probability". (8) The scenarios hold the published 2028 build-out fixed; in reality connection is conditional on financing and construction, and severe expected curtailment would itself delay or deter entrants — the constrained future's combination of full build-out with zero relief is a stress pairing, and real outcomes likely sit between it and the slow-track. (9) The model conflates two curtailment mechanisms — corridor congestion and system oversupply at low demand — inside the single parameter X, and holds demand implicit and constant (Section 3 states the evidence that oversupply is currently prominent). A stagnating or falling demand trend would tighten effective absorption capacity even with lines built on schedule — the reduced-form sensitivity in Section 4 (Figure 3) bounds that effect by shifting absorption in persistent MW blocks. The structural fix, and the model's next planned upgrade, is a residual-demand layer separating oversupply curtailment (demand-driven, system-wide) from congestion curtailment (wire-driven, corridor-specific). The core structural finding survives all of these caveats: within the model's representation, the official plan implies roughly 70–90% growth in equivalent carrying capacity in two years, and the cost of the equivalent relief not arriving is measured in Sections 4–6. Until the underlying data gaps close, this document should be quoted for its structure, not its point estimates.

References

Agbaje, T., Oni, D. & Ibeh, J. (2026). South Africa's curtailment surge is rewriting which renewables a bank will fund. Electron Intelligence Research Platform, 9 July 2026.

Cliffe Dekker Hofmeyr (2025). NTCSA releases practice note and updated GCCA on congestion curtailment. Projects & Energy Alert, 3 December 2025. www.cliffedekkerhofmeyr.com.

CSIR (2025). Statistics of utility-scale power generation in South Africa 2024. Council for Scientific and Industrial Research, January 2025. www.csir.co.za.

Energize (2025). Regulator backs curtailment to relieve grid congestion. By N. Schmidt, 13 May 2025. www.energize.co.za/article/regulator-backs-curtailment-to-relieve-grid-congestion. [NERSA decision conditions: compensation eligibility for all affected generators, MYPD6 cap, EC/WC scope, April 2025 – March 2028 window, six-monthly reporting.]

Engineering News (2026). NTCSA acknowledges 'temporary bottleneck' of R2bn in curtailment payments to IPPs. 18 June 2026. www.engineeringnews.co.za.

Eskom & NTCSA (2026). State of the System — 2026 Winter Outlook Briefing. 22 April 2026. www.eskom.co.za. [TDP delivery: 271 km built vs 423 km target FY26; required ~1,450 km/yr vs ~800 km/yr industry capacity; transformer lead times 36–48 months; 7,700 km servitudes secured; dumped-energy projection >5 TWh/yr by 2028; 26 GW forward-booked renewables.]

Hersbach, H., Bell, B., Berrisford, P., et al. (2020). The ERA5 global reanalysis. Quarterly Journal of the Royal Meteorological Society, 146(730), 1999–2049. (Accessed via the Open-Meteo historical weather archive, open-meteo.com.)

IPP Office (2025–26). IPPPP quarterly reports (Q3 Dec 2024 through Q3 Dec 2025). www.ipp-projects.co.za. [Checked directly: these quarterlies carry programme capacity and generation but contain NO curtailment or deemed-energy data — confirming NERSA (2026) as the sole official source for the 2025 curtailment cost, and that no official 2026 curtailment figure exists in programme reporting through December 2025.]

Lloyd's Register & SAWEA (n.d.). How well are South African wind farms performing? Whitepaper, South African Wind Energy Association. sawea.org.za.

Eskom Transmission (2021). Generation Connection Capacity Assessment of the 2023 Transmission Network (GCCA-2023), Phase 1, June 2021. www.eskom.co.za. [Methodology, verified directly: generation is absorbed by local load first with only the excess exported (§3.1); “the lower the load connected at a substation, the lower the generation capacity that can be connected”; a lower load forecast reduces available capacity (§2); connection capacity is the LOWEST of transformer, substation-transfer, local-area and supply-area limits (Fig. 1). Confirms that South African connection capacity is a combined network-and-load quantity — the basis for this model treating X as effective absorption.]

NTCSA (2024). The South African Wholesale Electricity Market (SAWEM) Market Code, Initial Draft 1.0, 18 April 2024. www.eskom.co.za. [Primary source for market mechanics, read directly: definition (181) — the System Marginal Price is set by “the incremental price of the marginal unit scheduled to run in the Unconstrained Schedule in a scheduling period”; definitions (196) and (28) treat the Unconstrained and Constrained Schedules as distinct constructs (Market Code §9.6); definition (116) sets a Market Price Cap on the SMP.]

Natural Justice, groundWork & LETA (2026). FAQs: South Africa's Wholesale Electricity Market — Market Code and Rules, June 2026. naturaljustice.org. [Trading expected to commence end-2026; the Central Purchasing Agency purchases IPP output under existing PPAs and trades it into SAWEM on their behalf, recovering the market-versus-contract difference through a legacy charge; Eskom generation and distribution are compulsory market participants.]

NERSA (2025). Reasons for Decision: NTCSA application for Congestion Curtailment to be a Constrained Generation Ancillary Service. National Energy Regulator of South Africa, June 2025. www.nersa.org.za.

NERSA (2026). Report on Monitoring Performance of Renewable Energy Power Plants, Issue 27, March 2026. National Energy Regulator of South Africa. www.nersa.org.za/files/files/2026/04/UpdatedReportonMonitoringPerformanceofREPowerPlantsIssue27__March2026_Final.pdf. [Page-verified against the PDF: total REI4P energy 17,808 GWh and load factors — wind 37%, PV 26%, CSP 32% — p. 10 (Fig. 4, p. 11); Table 5 energy/cost/price by technology incl. wind 11,040 GWh at R1.42/kWh, p. 14; deemed-energy cost R402m to 67 plants, p. 14, stated as ~R404m in the conclusions, p. 18; midnight/low-demand curtailment description, pp. 14 & 18. The series' H1-2024 issue recorded R46.2m across 40 IPPs.]

NTCSA (2025). Generation Connection Capacity Assessment (GCCA) 2025 Update. National Transmission Company South Africa. [Capacity tables, node-level curtailment-tier allocations, 4% budget-sizing language, 1,580 MW unlocked capacity.]

Yelland, C. (2026). Renewable energy on a financial cliff — and Eskom holds the cards. News24, 15 June 2026. [H1-2026 curtailed volume roughly an order of magnitude above the whole of 2025; projects ~9% below budgeted revenue; payment backlog approaching R1bn per industry sources.]

Note on citation practice and evidence hierarchy: primary regulatory and system-operator documents (NERSA decisions and monitoring reports, NTCSA/GCCA publications, Eskom briefings) govern rules and quantities; journalism and analyst platforms are cited only for stakeholder claims and recent developments not yet in primary documents, and carry correspondingly less weight. Figures described as "measured" or "published" are cited at the number; figures described as "derived" or "inferred" are computed by us from cited inputs, with the derivation stated. Scenario parameters carry no citation because they are assumptions — graded and listed in Appendix A.

Appendix A — Key assumptions and their confidence grades

Full register with sources and notes: ASSUMPTIONS.md in the model folder. Grades: A = measured or published primary source; B = published secondary source or industry convention; C = explicit estimate, adjustable in config.yaml.

Assumption	Value	Grade
2025 anchor — inferred proxy from deemed-energy cost, not measured	2.5% of wind potential	C
2026 corridor band — reasoned working assumption, not observed	5%–9%	C
Compensation mechanism (all affected generators)	capped pot; funded share = min(1, 4% ÷ corridor rate); paid 6–12 months late	A/B
Capacity build-out & node tiers (GCCA 2025)	3,866 → 7,205 MW by 2028	A
"4%" = budget-sizing assumption, not a cap	Eskom 5-yr revenue determination	A
Regional wind fleet output levels	35–38% of maximum, by region	B
Year-to-year wind variability	σ ≈ 5.5% (global convention; no SA series exists)	B
Cross-region correlation; sharing-rule parameters	scenario / estimate	C
Price, build cost, loan terms, running cost	R0.90/kWh; R20m/MW; 70% @ 11%, 15y; 2%	C
Effective grid capacity X (2025 / 2026 / design)	3 400 MW / 3 121 MW / 5 592 MW	solved

Appendix B — The capacity build-out, row by row

Every scenario in this report runs the same build-out: the corridor grows from 3,866 MW operational to 7,205 MW by 2028. This appendix shows exactly where that comes from. The named projects, provincial landing totals (WC 3,527 MW + EC 2,098 MW standard access), the 1,580 MW curtailment tier and its node allocations are taken from the grid company's published connection study (NTCSA, 2025). Where the GCCA gives only totals, the table carries explicit reconciliation rows — including a −226 MW adjustment where named WC projects overshoot the published landing point — and where phasing is not published (the curtailment tier's connection years), the split is an estimate and is graded as such. The full table with source notes is data/capacity_pipeline.csv in the model; the grade column follows Appendix A's A/B/C convention.

Project / row	Zone	Tech	MW	Year	Tier	Grade

WC operational wind overberg	overberg	wind	1100	2024	standard	C
WC operational wind karoo	overberg	wind	768	2024	standard	C
WC operational pv	karoo	pv	350	2024	standard	A
WC operational hybrid	karoo	wind	128	2024	standard	B
EC operational wind	eastern_cape	wind	1520	2024	standard	A
Grassridge cluster	eastern_cape	wind	492	2025	standard	A
Bacchus 1	overberg	wind	103	2025	standard	A
Bacchus 2	overberg	pv	80	2026	standard	B
Bacchus 3	overberg	wind	124	2026	standard	B
Droerivier pv	karoo	pv	250	2027	standard	A
Komsberg 1	karoo	wind	150	2027	standard	A
Komsberg 2	karoo	wind	320	2028	standard	A
Agulhas wind	overberg	wind	380	2028	standard	A
WC aggregation adjustment	overberg	wind	-226	2028	standard	B
EC pipeline residual	eastern_cape	wind	86	2027	standard	B
Bacchus tier	overberg	wind	140	2026	curtailment	A
Agulhas tier	overberg	wind	100	2027	curtailment	A
Droerivier tier	karoo	pv	220	2027	curtailment	A
Nuweveld Galenia tier a	karoo	wind	360	2027	curtailment	A
Nuweveld Galenia tier b	karoo	wind	360	2028	curtailment	A
Neptune Pembroke tier a	eastern_cape	wind	200	2027	curtailment	A
Neptune Pembroke tier b	eastern_cape	wind	200	2028	curtailment	A

Table C1 — The build-out as the model uses it (cumulative by connection year). Grades: A = GCCA-published figure; B = published total with estimated split or technology; C = estimated allocation (e.g. the coast/Karoo split of the pre-2024 WC fleet, tier phasing).

About the author

Michael Pietersen (PhD) is an independent quantitative modelling and risk-analysis consultant (StatForge, Pretoria). His work applies statistical and time-series methods — distributions rather than point estimates, volatility, clustering and tail risk — to commercial decision problems. This assessment was produced end-to-end by the author and is fully reproducible: the data assembly, hourly simulation engine, validation suite and this document itself regenerate from a fixed seed, and every input traces to the graded register in Appendix A.

Contact: michael@statforge.co.za — for discussion of this assessment, corrections or data contributions, or bespoke analysis built on the same machinery: portfolio-specific curtailment exposure, renewable-surplus valuation and capture-price modelling under SAWEM, and contract-structure risk quantification.