Asia Climate Risk Screener: A Zero-Backend Browser Tool for Asset-Level Physical and Transition Climate-Risk Screening

Ka Wah Philip Ng
PhD, FCPA (Aust. & HK), FCMA, CESGA
Independent Researcher, Hong Kong
ngkwphilip@gmail.com · ORCID: 0009-0009-3955-010X

Abstract—Asset-level climate-risk analytics remain priced for large institutions, leaving many financiers, insurers and plant owners without an accessible first-pass screen. This paper presents the Asia Climate Risk Screener, a single self-contained web page that screens physical and transition climate risk for more than 15,000 geolocated power stations across Asia — all operating Asian coal plants in the GEM census, the wider satellite-estimated fossil fleet of gas and oil assets, plus the solar, wind, hydropower, nuclear, geothermal and bioenergy fleet — using only open satellite and climate data. Selecting any asset builds a plain-language risk profile in seconds from four live public feeds — six decades of measured weather (ERA5), river-discharge records since 1984 (GloFAS), climate-model projections to 2050 (CMIP6), and satellite-derived solar and wind resource (NASA POWER) — joined to a bundled satellite-emissions asset layer (Climate TRACE). Feature engineering condenses multi-decadal daily weather and river records into decision-ready signals, and a transparent heuristic scoring model flags heat stress, flood trends and transition exposure, with explicit "feed unavailable" states rather than silent gaps. I evaluate the hazard-data pipeline against documented Asian climate disasters using a two-arm recall-and-specificity harness, detrended and graded at the 95th percentile. Cross-matching the coal fleet against an independent facility census confirms 96% coverage. The system is zero-backend — no server, account or API key — is deployed as a static site, and is released openly.

Index Terms—*climate risk, physical risk, transition risk, decarbonisation, power sector, Asia, open data, climate finance, decision support, zero-backend web application, evaluation.*

I. Introduction

A lender weighing a loan to a power station in coastal Bangladesh, an insurer pricing cover for a wind farm in Vietnam, or an investor screening a portfolio of Indonesian coal plants all face the same question: what will climate change do to this specific asset at this specific location? Answering it today typically means commissioning consultants or subscribing to commercial climate-risk platforms priced for large institutions. Yet the raw ingredients for an answer are already public: satellite-estimated emissions inventories, six decades of reanalysed weather, global river-flow records, downscaled climate-model projections, and satellite-derived renewable-resource climatologies. What is missing is the join — a tool that brings these feeds together at the level of an individual asset and states the result in plain language.

The Asia Climate Risk Screener is that join. It is a single self-contained web page that maps more than 15,000 geolocated Asian power stations — Climate TRACE's satellite-estimated fossil fleet cross-matched to every operating coal plant in the Global Energy Monitor (GEM) census (its pipeline and mothballed records are also mapped), plus the solar, wind, hydropower, nuclear, geothermal and bioenergy fleet from GEM's 2026 tracker releases — and builds a live risk profile for any of them in seconds. Each profile combines measured physical hazards (heat stress, extreme rain, river flooding), one climate-model's projected shift to mid-century, the site's solar and wind replacement resource, and a transition-exposure line drawn from the plant's fuel, age and any planned retirement. A composite screening score summarises the profile; a plain-English summary states what changed, what is projected, and what a decision-maker should look at next.

The entire system is zero-backend: one HTML file plus small static data bundles, fetching four public CORS-enabled APIs directly in the viewer's browser, joined to a bundled satellite-emissions asset layer. There is no server, database, account, or API key; deployment is a drag-and-drop onto a static host. This paper describes the data estate (Section II), the architecture (Section III), the indicators and their alignment with industry frameworks (Section IV), a live two-arm evaluation harness (Section V), a measured coverage analysis (Section VI), a worked case study (Section VII), practitioner lessons (Section VIII), and limitations (Section IX).

The contributions are threefold: (1) an integrated, asset-level physical-plus-transition screening workflow over a current (2026-vintage) open census of more than 15,000 Asian power stations; (2) an honestly graded two-arm evaluation harness — the tool tests both the sensitivity and the specificity of its own hazard-data pipeline against documented disasters and control location-years, and reports failures openly rather than hiding them; and (3) a zero-backend architecture with explicit degraded states, reproducible by anyone with a browser. I claim no algorithmic novelty: the indicators are deliberately standard, and the contribution lies in the integration, the measured evaluation, and the accessibility.

II. Data Sources

The asset estate combines one emissions inventory with six GEM facility trackers, all current 2026 releases, all permitting reuse with attribution (Table I). Climate TRACE provides 5,341 Asian fossil-fuel power assets with satellite-estimated annual emissions [1]. GEM's Global Coal Plant Tracker provides the complete Asian coal fleet represented in the census — 2,265 plants aggregated from unit-level records with capacities, unit ages, planned retirements and owner chains [2] — and GEM's solar, wind, hydropower, geothermal, bioenergy and nuclear trackers provide 10,116 operating clean-energy plants (50 MW and above, except nuclear, for which all sizes are kept) [3], [4], [5]. All live hazard feeds are keyless and carry open licences: Open-Meteo's ERA5 archive API (weather since 1961), the GloFAS river-discharge API (records since 1984), the CMIP6 downscaled climate API, and NASA POWER's satellite-derived solar and wind climatology [6]–[11].

Table I. Data sources at a glance.

Source	Content	Vintage	Licence
Climate TRACE v7 API [1]	5,341 fossil power assets, satellite-estimated emissions	2025 release	CC BY 4.0
GEM Coal Plant Tracker [2]	2,265 Asian coal plants, unit-level, owners, retirements	Jul 2026	CC BY 4.0
GEM Solar & Wind Trackers [3]	8,221 operating projects ≥50 MW	Feb 2026	CC BY 4.0 (selected solar records CC BY-NC 4.0 via TransitionZero)
GEM Hydro & Geothermal Trackers [4]	1,589 operating plants ≥50 MW	Mar 2026	CC BY 4.0
GEM Bioenergy (GBPT V3) & Nuclear [5]	256 bioenergy plants; all 50 operating nuclear plants	2026	CC BY 4.0

Source	Content	Vintage	Licence
Open-Meteo Archive / Flood / Climate [6]	ERA5 weather; GloFAS discharge; CMIP6 projections	live	CC BY 4.0 (© ECMWF/Copernicus/JRC)
NASA POWER [11]	Monthly solar irradiance and 50 m wind, 2020–24 means	live	NASA open data

Two sources are bundled statically rather than fetched live, for different reasons. Climate TRACE assets are harvested once through its paging API (5,341 records) and shipped as a compact JavaScript bundle, both for load speed and because a browser client should not page a public API two dozen times per visit. The GEM trackers are distributed through a download form rather than an open API; the one-off downloads are therefore bundled at build time, with release dates stated wherever the data appear. All other feeds are fetched live on every use, so the tool requests the latest data exposed by the source APIs on each use.

III. System Architecture

A. Zero-backend static design

The dashboard is one HTML file (roughly 1,000 lines of HTML/CSS/vanilla JavaScript with Leaflet and Chart.js from CDNs [15]) plus four static data-bundle files (fossil assets with city benchmarks, the coal census, clean plants, evaluation events). All hazard computation — six decades of daily weather, river records, climate projections — happens in the viewer's browser against the live feeds. There is no backend to secure, scale, or pay for; the whole site deploys by dragging a zip onto a static host and also runs from the local filesystem. Fig. 1 shows the asset map the user lands on.

Fig. 1. The asset map (Tab 1) of the live site: more than 15,000 Asian power stations. Coloured solid dots are fossil-fuel stations sized by satellite-estimated annual emissions; clean-energy plants are sized by capacity (solar yellow, wind teal, hydro blue, nuclear pink, geothermal red, bioenergy green); hollow black rings mark the coal census. The CARTO basemap extends beyond the plotted assets, so continent labels (Europe, Africa, Oceania) are visible at the frame edges; all plotted stations are within the Asian study region. Captured 2026-08-21.

B. Honest degradation

Public endpoints are routinely unreachable from corporate or filtered networks. Every panel therefore has an explicit degraded state: a grey "feed unavailable" label naming the failed source, never a silent blank or a fabricated flat line. A status panel on the method tab re-tests all live feeds on demand and reports reachability from the user's own network (Fig. 2). A feasibility spike on a filtered corporate network in August 2026 measured five of six candidate feeds tested during development as reachable both with and without a VPN (typical latency 0.3–1.4 s); the sixth (the OpenStreetMap Overpass API) was unreachable on all four public mirrors and was replaced by statically bundled context rather than allowed to fail silently.

Live feed status

Feed	Status (this network, now)
Weather history (ERA5)	reachable
River discharge (GloFAS)	reachable
Climate projections (CMIP6)	reachable
Solar/wind resource (NASA POWER)	reachable
Asset data (Climate TRACE, bundled)	reachable

Re-test feeds now

Fig. 2. Live feed status panel: all live feeds reachable from the author's filtered corporate network, re-testable on demand. A failed feed would show its honest unavailable state here and in every dependent panel.

C. Polite-client behaviour

Free APIs rate-limit burst traffic. After evaluation runs intermittently hit Open-Meteo throttling, the client was re-engineered: requests fetch only the variables each indicator needs, every feed call gets one automatic retry after an eight-second backoff, and evaluation events are paced three seconds apart. Post-fix, a full evaluation run completes in about a minute with no failed feeds.

IV. Indicators and Methods

A. Measured physical hazards

Heat stress is characterised by two complementary indicators so that the metric remains informative across both tropical and temperate assets. The primary index is location-relative: following the ETCCDI TX90p convention [12], I compute the frequency of days on which daily maximum temperature exceeds the site's own calendar-day 90th percentile, calibrated on the 1961–1990 reference period; the implementation reports this exceedance frequency as days per year (equivalently the TX90p fraction expressed against a 365-day year) so that the figure is directly interpretable to a practitioner, and compares the most recent twenty-year normal against that baseline. A secondary, absolute index — the count of days per year with maximum temperature at or above 35 °C — is reported alongside it for warm-climate assets, where an absolute physiological threshold is directly interpretable; for assets whose full-record 35 °C count is near zero this absolute index is labelled "not informative at this location" rather than reported as a spurious zero-trend. Trend significance on both series is assessed with the non-parametric Mann–Kendall test, and the magnitude of change is estimated with the Theil–Sen slope; a change is flagged only where the Mann–Kendall two-sided p-value is below 0.05, replacing the earlier fixed-magnitude threshold. Where a discrete step is of interest, a Pettitt change-point test locates the most probable single break and reports its year and significance.

Extreme-rain trend is computed on the annual maximum one-day precipitation (Rx1day) series over the full ERA5 record. Rather than ranking a single recent value — which is dominated by interannual noise — I fit a Mann–Kendall/Theil–Sen trend to the Rx1day series and report the per-decade slope and its significance, so the indicator reflects a genuine tendency rather than one wet year.

River-flood trend compares recent annual peak discharge from the matched GloFAS cell against the 1984–2014 reference norm, but only after a hydrological-relevance gate. The matched cell must satisfy a minimum long-term mean discharge and lie within a maximum channel distance of the asset; the matched cell's distance and mean flow are printed on the risk card. Where no cell clears the gate, the panel shows an explicit "no hydrologically relevant river cell within tolerance" state rather than reporting a percentile against a minor tributary.

B. Modelled future shift by delta method, with benchmark

Because raw absolute output from a single climate model carries substantial local bias — particularly for threshold-exceedance counts at coastal grid points — the projected shift is applied by the delta (change-factor) method rather than read from raw model absolutes. For each indicator I take the model's change between the 1995–2014 baseline window and the 2031–2050 projection window and apply that change to the observed ERA5 baseline:

$$X_{projected} = X_{observed\ baseline} + (X_{model,\ 2031–2050} - X_{model,\ 1995–2014}).$$

For mean hottest-day temperature the change is additive as above; for annual rainfall a multiplicative change factor is used so that the model's fractional change, not its biased absolute, is transferred to the observed record. Both the observed baseline and the model baseline are displayed, so any model bias is visible to the reader rather than silently propagated. A single high-resolution model (EC-Earth3P-HR, SSP5-8.5-class) is used for the live single-asset view; because one model is one plausible future, every projected shift is shown against a regional benchmark — the median delta across 24 representative Asian cities (measured: +0.99 °C hottest-day warming, +1.3 hot days per year, +8.4% annual rain) — and, where the run completes within the rate budget, against the inter-model range of a small HighResMIP ensemble, reported as a spread rather than a point estimate. The benchmark is internal to the same model: it demonstrates spatial consistency, not model skill.

C. Replacement resource and transition exposure

For fossil assets, satellite-estimated emissions are converted to implied annual generation with clearly labelled emission factors (0.95/0.45/0.70 t CO_2 per MWh for coal/gas/oil), noting that the inventory is reported as CO_2-equivalent and therefore yields a modest over-estimate of generation. The approximate land area required to replace that generation with utility solar at the local resource is

$$A = E_{annual} / (CF \times 8760 \times \rho),$$

where E_{annual} is implied annual generation (MWh), CF is the NASA POWER solar capacity-factor proxy, 8760 is hours per year, and ρ is an empirically calibrated area power density. I adopt ρ ≈ 40 MW/km², consistent with the footprint of operating utility-scale parks in comparable Asian resource, in place of a purely optical yield; the resulting figure is presented explicitly as an order-of-magnitude land-feasibility check with a stated uncertainty of roughly a factor of two, not a design figure. Under this calibration a coal asset with ~9,700 GWh/yr implied generation at a 17% capacity factor requires on the order of 160 km² rather than the ~19 km² produced by the earlier expression, which assumed an unrealistically high land yield.

The 50 m NASA POWER wind speed is extrapolated to a nominal 100 m hub height by the logarithmic wind profile before the viability screen, since modern utility turbines operate well above 50 m; the coarse (~0.5°) resolution of the wind field and its insensitivity to local terrain are stated as caveats alongside the ERA5 heat-grid caveat. Transition exposure combines fuel class with plant age from the coal census (oldest unit versus a ~40-year typical life) and any planned retirement already on the books.

D. Composite score, banding, and standards alignment

Each qualitative rating pill is mapped to a numeric band value on a 0–100 scale before aggregation, so the composite is a transparent function of its inputs rather than an opaque grade. The mapping is deliberately coarse — four levels — because the underlying indicators are screening signals, not calibrated risk quantities. Table III states the mapping in full.

Table III. Ordinal-to-numeric banding for the composite score.

Rating pill	Band value (0–100)	Interpretation
lower / low	20	at or below the location's historical norm
moderate	50	a detectable but non-extreme shift
elevated	70	a significant shift flagged by the indicator
high / severe	90	extreme shift or unambiguous transition exposure
feed unavailable	excluded	dropped from the mean; never scored as zero

The physical-hazard line is the mean of the available hazard band values (heat stress, extreme-rain trend, river-flood trend, and the modelled mid-century shift); the transition line is the band value of transition exposure. Indicators in a "feed unavailable" state are dropped from the mean rather than imputed. The composite screening score is the weighted combination

$$S = wphys \cdot Hphys + wtran \cdot Htran, \quad wphys + wtran = 1,$$

with equal weights (wphys = wtran = 0.5) as the disclosed default. The score and every rating pill are labelled heuristic screening grades, not investment, lending or insurance advice. The methodology aligns with established practice where possible — WMO climate-normal conventions, ETCCDI-style extreme indices [12], and the TCFD / IFRS S2 taxonomy separating physical from transition risk [13], [14] — while the banding thresholds themselves are explicitly non-standard and stated as such in the tool.

E. Weight sensitivity

Because the equal-weight default is a disclosed convention rather than a derived optimum, the tool reports how sensitive the composite is to the physical/transition split. For an asset with a moderate measured-hazard line and a high transition line — the common profile of an ageing coal plant in a benign climate — the two weightings bracket a wide range: a transition-led financier applying wtran = 0.7 obtains a materially higher score than an operations-led user applying wphys = 0.7, even though the underlying data are identical. The interface exposes the weight as a user control and prints the score under the equal-weight default alongside the 70/30 and 30/70 endpoints, so the central modelling choice is visible and adjustable rather than buried in a constant.

F. Reproducibility: all disclosed constants in one place

Every constant used by the scoring pipeline is fixed and disclosed; none is fitted to the validation events. Table IV collects them so that any rating pill, land-feasibility figure, or composite score in this paper can be reproduced by hand from the raw feeds. Where a constant is a deliberate simplification rather than a physically derived value, it is marked as heuristic.

Table IV. Disclosed parameters for full reproducibility.

Parameter	Value	Role	Type
Absolute hot-day threshold	35 °C	warm-climate heat count threshold	heuristic
Relative hot-day percentile	calendar-day 90th (TX90p)	primary heat index	ETCCDI
Heat baseline period	1961–1990	reference normal	WMO
Recent normal window	most recent 20 yr	comparison window	WMO
Trend test	Mann–Kendall, $p < 0.05$	trend significance	non-parametric
Slope estimator	Theil–Sen	trend magnitude	non-parametric
Change-point test	Pettitt	single-break detection	non-parametric
Flood reference norm	1984–2014	GloFAS baseline	reference
River-cell gate	≥100 m^3/s long-term mean, within ~30 km	hydrological relevance gate	heuristic
CMIP6 model	EC-Earth3P-HR (SSP5-8.5-class)	projection	HighResMIP
Projection windows	1995–2014 vs 2031–2050	delta-method windows	—
Emission factors, coal/gas/oil	0.95 / 0.45 / 0.70 t CO_2/MWh	implied generation	heuristic
Solar area power density (ρ)	40 MW/km^2	replacement land	heuristic
Solar CF 'strong' mark	≥16–18%	resource screen	heuristic
Wind viability (100 m hub)	≥6.5 m/s	resource screen	heuristic
Wind height extrapolation	50 m → 100 m, log profile z_0 = 0.03 m (×1.092)	hub-height screen	standard profile
Band values, lower/mod/elevated/high	20 / 50 / 70 / 90	ordinal → numeric	heuristic
Composite weights (default)	0.5 / 0.5	aggregation	disclosed default
Recall percentile bar	≥95th	evaluation hit threshold	protocol
Specificity expected rate	~5%	control hit rate expected	protocol

V. Evaluation Against Documented Disasters

A screening tool must satisfy two requirements that a single-arm harness cannot separately test: documented disasters should rank in the extreme tail of the affected location's record (sensitivity/recall), and ordinary location-years should not (specificity). The evaluation harness therefore runs both arms. In the recall arm, the full measured record at each documented disaster is downloaded and the event year's rank in the location's own history is computed; a hit is recorded where the event clears the 95th percentile, a stricter and more defensible bar than a one-in-five-year value. In the specificity arm, every remaining location-year in the same records — those not among the documented events studied here, rather than years independently verified to be free of any climate impact — is graded by the identical rule as a control; because the extreme-tail definition is the top 5% by construction, a well-behaved indicator should flag close to 5% of these control years, and a materially higher control hit-rate would indicate that the indicator is flagging the secular trend rather than the event.

To prevent the underlying warming or wetting trend from inflating both arms, each series is detrended (Theil–Sen slope removed) before percentile ranking, so the reported rank reflects the anomaly relative to the local trend rather than the trend itself. Results are reported as a recall rate and a control hit-rate with their sample sizes, and, given the small number of documented events, individual event grades are treated as diagnostic rather than as a headline accuracy percentage. Misses are reported openly with their diagnostic explanation; the Bangkok 2024 case is retained as an example of grid-resolution and indicator-selection limits (night-time and humidity extremity under-captured by a daytime maximum-temperature index over a diluted coastal cell) rather than removed.

A clarification is essential here, because the harness deliberately tests a different and harder question than the one the operational heat indicator answers. The shipped heat indicator (Section IV-A) is distributional: it compares the recent twenty-year normal against the 1961–90 baseline and flags a site where that shift is significant — which is precisely why Castle Peak is graded elevated at $p < 0.001$. The recall arm, by contrast, detrends each series and then asks whether a single named year stands out against its own local trend. These are not the same test, and they need not agree: an event whose extremity is largely the secular warming trend will be correctly flagged elevated by the operational distributional indicator yet score "not confirmed" as a single-year anomaly once that trend is removed. The four heat "misses" in Table II are exactly this case — at Jacobabad, Chongqing and Bangkok the shipped indicator does flag the site as elevated, and the harness is reporting only that the specific event year is not an outlier relative to the local trend. I report the stricter, detrended single-year test because it is the more conservative and more falsifiable of the two; a reader should therefore read the recall figure as a lower bound on the tool's sensitivity to distributional heat shift, not as evidence that the operational heat indicator fails on these sites.

Table II shows a full run (graded 2026-08-21, recomputed live from the raw records): recall 4 of 8 (50%), with 3.8% of 477 control location-years flagged — consistent with the expected false-positive rate under the percentile-based protocol, and indicating that the tail test is not simply flagging the secular warming trend. The pattern of misses is instructive. All four confirmed events are rainfall or flood events; all four misses are heat events. At Jacobabad, Chongqing and Bangkok the event year still ranks between the 85th and 94th percentile of the detrended record — top-decile-to-top-15% anomalies — but a large share of those heatwaves' raw extremity is the secular warming trend itself, which the detrending step deliberately removes; at Delhi the 2024 event ranks only 57th, its extremity concentrated in a single season and in night-time minima, which an annual daytime-maximum count dilutes. I read the 50% recall not as a failure of the shipped indicator — which flags all four heat sites as elevated — but as a quantified statement of what the stricter single-year anomaly test measures: it captures events that are anomalous relative to the local trend, and it shows its workings openly when an event's extremity is mostly the trend.

Table II. Evaluation scorecard under the two-arm protocol (detrended series, ≥95th-percentile bar), full live run.

Event	Where / when	Signal	Event-year value	Rank (detrended)	Grade
Zhengzhou extreme rainfall	Henan, China, 2021	rain	169.6 mm	95th pct	CONFIRMED
Pakistan super-flood	Sindh, 2022	flood	59 m³/s	97th pct	CONFIRMED
India–Pakistan heatwave	Jacobabad, 2022	heat	34 hot days	85th pct	NOT CONFIRMED
Yangtze heatwave & drought	Chongqing, 2022	heat	102 hot days	94th pct	NOT CONFIRMED
Southeast Asia heatwave	Bangkok, 2024	heat	138 hot days	92nd pct	NOT CONFIRMED
North India heatwave	Delhi, 2024	heat	25 hot days	57th pct	NOT CONFIRMED
Cyclone Michaung rain	Chennai, 2023	rain	254.9 mm	98th pct	CONFIRMED
Kyushu floods	Kumamoto, 2020	rain	169.9 mm	97th pct	CONFIRMED

Event	Where / when	Signal	Event-year value	Rank in detrended record	Grade
Zhengzhou extreme rainfall ('21·7' Henan flood)	Zhengzhou, China, 2021	rain	169.6 mm	95th percentile	CONFIRMED
Pakistan super-flood (Indus, Sindh)	Sindh, Pakistan, 2022	flood	59 m³/s	97th percentile	CONFIRMED
India–Pakistan heatwave	Jacobabad, Pakistan, 2022	heat	34 hot days	85th percentile	NOT CONFIRMED
Yangtze basin heatwave & drought	Chongqing, China, 2022	heat	102 hot days	94th percentile	NOT CONFIRMED
Southeast Asia heatwave	Bangkok, Thailand, 2024	heat	138 hot days	92th percentile	NOT CONFIRMED
North India heatwave	Delhi, India, 2024	heat	25 hot days	57th percentile	NOT CONFIRMED
Cyclone Michaung extreme rain	Chennai, India, 2023	rain	254.9 mm	98th percentile	CONFIRMED
Kyushu floods	Kumamoto, Japan, 2020	rain	169.9 mm	97th percentile	CONFIRMED

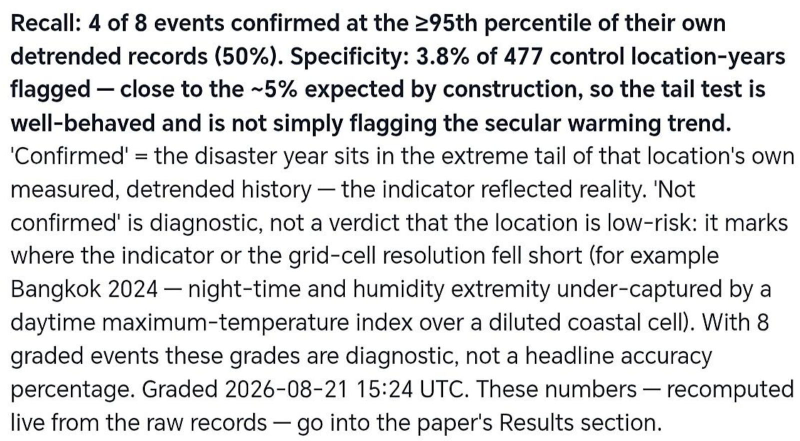

Fig. 3. The evaluation harness (labelled "validation" in the interface) after a full run under the two-arm protocol: documented disasters graded live at the 95th-percentile bar on detrended records, with the recall rate, the control (specificity) hit-rate, and the diagnostic framing printed beneath the score.

VI. Coverage Analysis

Coverage was measured, not asserted. Cross-matching the coal census against the emissions inventory at five-kilometre tolerance within the same country: 96% of operating Asian coal plants in the GEM census have a matched Climate TRACE satellite emissions estimate, and 75% of Climate TRACE coal assets match a census plant — the gap reflecting assets below the census threshold and naming differences. Because a five-kilometre nearest-neighbour match can generate false pairs in dense plant clusters, a precision audit accompanies the match rate: of the 1,966 matched pairs, 84% show exact or near-exact name agreement after normalisation (80% exact), while the discordant remainder mixes operator-versus-plant naming for the same complex with a residue of false pairs, so the 96% match rate should be read together with this name-audited precision.

The clean-energy layer is fully sourced from GEM's 2026 tracker releases, so the post-2020 renewables build-out is present. To parallel the coal cross-match, a random sample of 400 GEM solar plants commissioned by 2023 was cross-validated against the independent, satellite-derived TransitionZero Solar Asset Mapper registry (Q1 2024 vintage): 71% have a registry asset within five kilometres (median separation 2.0 km across the sample), the residual reflecting disagreement in either source — satellite under-detection, commissioning-date uncertainty, or coordinate error — rather than evidence of phantom plants. Earlier versions of the tool used the WRI Global Power Plant Database (a ~2020 snapshot) for the clean layer; reviewer testing caught the missing post-2020 build-out, and the layer was rebuilt — exactly the kind of failure an honest coverage statement exists to surface.

VII. Case Study: Castle Peak Power Station, Hong Kong

Castle Peak is one of the world's largest coal-fired stations: the emissions inventory records 9.221 Mt CO_2e for 2025, and the matched census record adds the detail a risk officer needs — a gross installed capacity of 3,058 MW across five units (operator figure: 1 × 350 MW and 4 × 677 MW; the census record states 3,060 MW), the oldest operating since 1985 (about 41 years, at the end of a ~40-year typical life), a planned retirement on the books for 2035, and the owner chain (Castle Peak Power Co Ltd). Clicking the plant produces the profile in Fig. 4 within seconds, computed live on 2026-08-21.

Measured hazards: on the primary TX90p index the site runs at 82 hot days a year over the most recent twenty years against a 1961–90 normal of 34 (≈2.5× — flagged elevated), with the Mann–Kendall trend highly significant ($p < 0.001$, Theil–Sen slope +6 hot days per decade) and a significant Pettitt change-point in 1998. The absolute 35 °C index is near zero at this site and is labelled not informative at this location — a live demonstration of why the primary heat index is location-relative. The extreme-rain trend is moderate (+2.7 mm per decade on Rx1day; Mann–Kendall $p = 0.08$, not significant). The river-flood line shows the hydrological-relevance gate working as designed: Hong Kong has no major river, and the best-matched GloFAS cell (25 km east) has a long-term mean flow of 0.7 m^3/s — far below the 100 m^3/s gate — so the panel declines to grade riverine flood against a minor channel, shows the explicit gate state, and excludes the line from the composite rather than imputing it. The delta-method future shift (+0.9 °C) is modest and in line with the +0.99 °C regional benchmark. The replacement-resource panel (Fig. 5) shows a usable solar resource (≈17% capacity factor) and marginal wind (5.4 m/s at 50 m, ≈5.9 m/s at a 100 m hub); replacing the plant's implied ~9,700 GWh/yr would take on the order of 160 km^2 of solar at that resource under $\rho \approx 40$ MW/km^2 — a land-feasibility sanity check with a stated factor-of-two uncertainty, not a design figure. Transition exposure is rated high: coal, end of typical life, with a planned retirement already scheduled.

Under the published banding (Table III) — heat elevated 70, rain moderate 50, modelled shift lower 20, river flood excluded by the gate — the physical-hazard line is 46.7; with the transition line high at 90, the equal-weight composite is 68.3, reported as 68/100. The interface prints the weight-sensitivity endpoints alongside: 60 under a physical-led 70/30 split and 77 under a transition-led 30/70 split — same underlying data, disclosed weighting. The plain-English summary draws the decision-relevant conclusion: measured heat risk is elevated, and the long-term case for the asset depends on a credible retrofit-or-retire plan. For a financier, the entire first-pass screen — four live feeds, a bundled emissions estimate, the coal census, six decades of weather — took under a minute and cost nothing.

Fig. 4. Risk profile for Castle Peak power station, computed live 2026-08-21: emissions estimate, matched census record (capacity, units, age, planned retirement, owners), measured and modelled hazard pills under the published banding, transition exposure, and the composite screening score with the weight control, sensitivity endpoints and disclaimer.

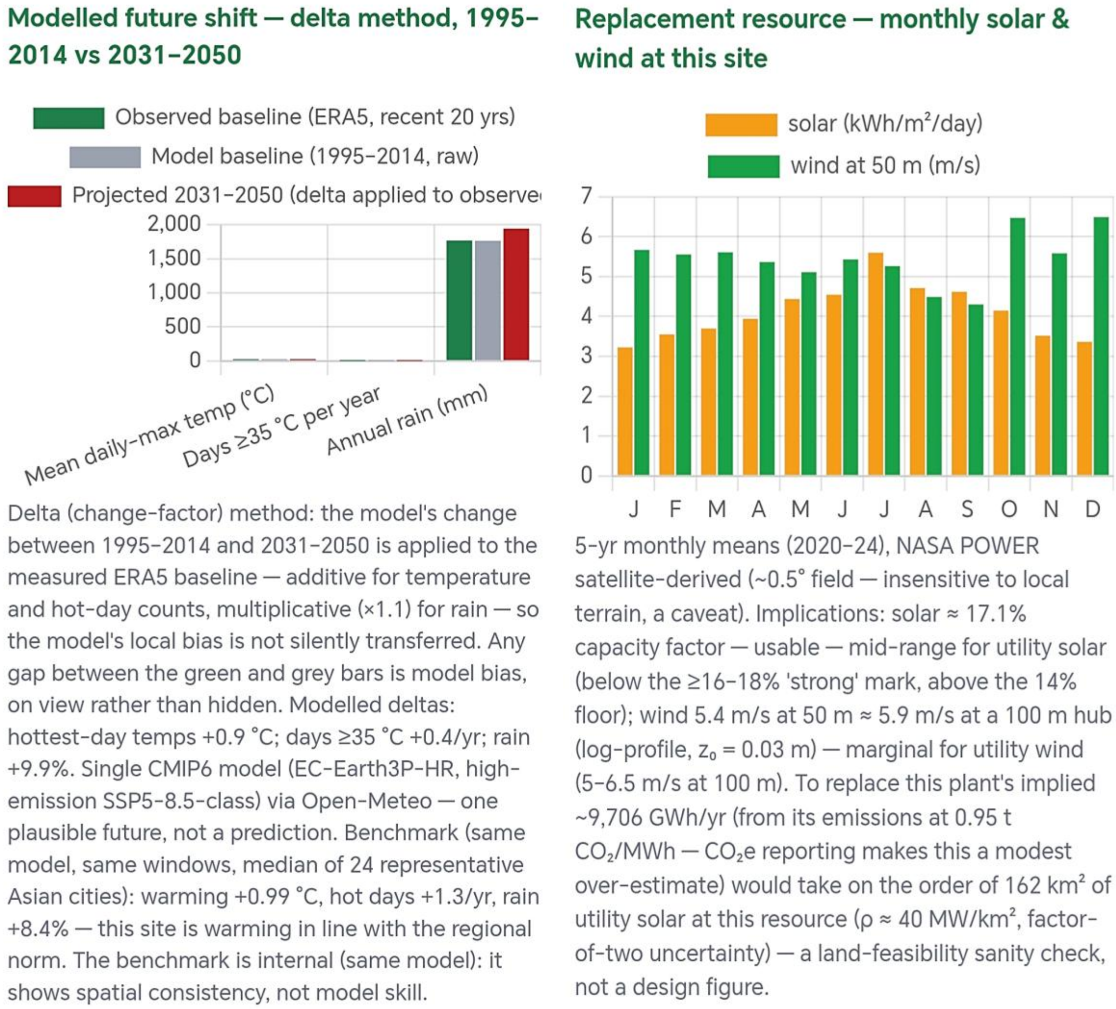

Fig. 5. Evidence panels behind the profile: the delta-method 2031–2050 projection (left: observed ERA5 baseline, raw model baseline, and the delta-applied projection shown side by side so model bias is visible), and the monthly solar/wind replacement resource with its implications (right). Captured 2026-08-21.

VIII. Lessons Learned

1) Silent parameter mismatches are the norm. The emissions API silently ignores singular parameter names (sectors= works, sector= does not) and its country filter returned unsorted global data until filtering moved client-side. Defensive parsing must verify what actually came back, not what was requested.

2) Country codes surprise. Hong Kong plants are filed under HKG, not CHN; a country filter built on intuition silently dropped them. Testing with assets you personally know catches what no schema will.

3) Climate APIs have hard temporal edges. Every CMIP6 model exposed through the weather API used here is a HighResMIP run ending in 2050; requesting a 2041–2060 window returns HTTP 400. The fix was a 2031–2050 window, stated honestly as "to 2050" everywhere in the interface.

4) Free feeds rate-limit bursts; engineer for it. Lighter per-indicator requests, one retry with backoff, and three-second pacing turned an intermittently failing evaluation run into a reliable one-minute one.

5) Some public infrastructure is simply unreachable from real corporate networks. All four public mirrors of one GIS query API failed. Bundling static context at build time is a legitimate substitute — but only because the interface discloses it.

6) Honest degraded states preserve trust. Throughout testing, failed feeds were discussed as data, not as embarrassment; a fabricated flatline would have been discovered and remembered.

IX. Limitations and Future Work

Intended use — the Asia Climate Risk Screener is a first-pass screening and education tool. Its ratings are heuristic grades computed from public data; they are not investment, lending, underwriting or insurance advice, and they are not a substitute for engineering due diligence.

The primary live projection remains based on a single high-emission CMIP6 model applied by delta method and presented as one plausible future, benchmarked rather than validated; systematic multi-model ensemble reporting, rather than the current rate-limited supplementary range, is an important next step. Gridded reanalysis smooths local extremes, as the Bangkok evaluation miss shows. The hazard set does not yet include cooling-water/drought stress (relevant to the 2022 Yangtze case and to river-cooled thermal and hydro assets), tropical-cyclone wind, or coastal sea-level-rise and storm surge; these are the hazards most likely to strand a specific power asset and are prioritised future work. Fuel labels in the emissions inventory are occasionally wrong or name a neighbouring complex; every fossil risk card carries this caveat, and the measured hazard lines are location-based and unaffected. The replacement-area figure is a heuristic with a stated factor-of-two uncertainty. The emissions inventory is a 2025 snapshot. No uncertainty is yet propagated to the composite score; coarse confidence bands are a planned addition. Non-power assets and plants below tracker thresholds are out of scope. Further planned extensions include the GEM oil-and-gas plant tracker to enrich the fossil layer and portfolio-level batch screening.

X. Conclusion

Asset-level climate-risk screening for Asia's power sector need not sit behind a paywall. With only open satellite and climate data, a browser, and no infrastructure, the Asia Climate Risk Screener covers more than 15,000 stations, states its numbers in plain language, discloses every heuristic in a single reproducibility table, and tests its hazard-data pipeline against documented disasters with a two-arm harness that reports both recall and specificity, openly showing the misses alongside the hits. The tool is deployed as a static site; source and data bundles are released openly.

Acknowledgment

The author thanks Climate TRACE, Global Energy Monitor, Open-Meteo, the ECMWF/Copernicus and JRC teams behind ERA5 and GloFAS, and NASA POWER for maintaining open data services, and the OpenStreetMap and CARTO contributors for the basemap.

Disclosure of AI Assistance

This system was designed and implemented with AI assistance (Kimi AI) acting as a pair-programming and drafting tool under continuous human direction. All engineering decisions, data-validation steps, and the final text of this manuscript were reviewed, tested, and approved by the human author, who takes full responsibility for the content.

Data and Code Availability

Source code and data bundles are released under the MIT licence, permanently archived on Zenodo (doi: 10.5281/zenodo.22058897), and maintained at github.com/drngkwphilip/asia-climate-risk-screener; the live demonstration is at https://asia-climate-risk-screener.netlify.app. Data: Climate TRACE [1]; Global Energy Monitor [2]–[5]; Open-Meteo [6]; NASA POWER [11].

References

[1] Climate TRACE, "Emissions inventory and open API," climatetrace.org, 2025 release, accessed Aug. 20, 2026.

[2] Global Energy Monitor, "Global Coal Plant Tracker," July 2026 release, globalenergymonitor.org, accessed Aug. 20, 2026.

[3] Global Energy Monitor, "Global Solar Power Tracker" and "Global Wind Power Tracker," February 2026 releases, globalenergymonitor.org, accessed Aug. 20, 2026.

[4] Global Energy Monitor, "Global Hydropower Tracker" and "Geothermal Power Tracker," March 2026 releases, globalenergymonitor.org, accessed Aug. 20, 2026.

[5] Global Energy Monitor, "Global Bioenergy Power Tracker (GBPT V3)" and "Global Nuclear Power Tracker," 2026 releases, globalenergymonitor.org, accessed Aug. 20, 2026.

[6] P. Zippenfenig, "Open-Meteo.com Weather API," Zenodo, 2023, doi: 10.5281/zenodo.7970649.

[7] H. Hersbach et al., "The ERA5 global reanalysis," Q. J. R. Meteorol. Soc., vol. 146, no. 730, pp. 1999–2049, 2020, doi: 10.1002/qj.3803.

[8] L. Alfieri et al., "GloFAS — global ensemble streamflow forecasting and flood early warning," Hydrol. Earth Syst. Sci., vol. 17, pp. 1161–1175, 2013, doi: 10.5194/hess-17-1161-2013.

[9] V. Eyring et al., "Overview of CMIP6 experimental design and organization," Geosci. Model Dev., vol. 9, pp. 1937–1958, 2016, doi: 10.5194/gmd-9-1937-2016.

[10] R. J. Haarsma et al., "High Resolution Model Intercomparison Project (HighResMIP v1.0) for CMIP6," Geosci. Model Dev., vol. 9, pp. 4185–4208, 2016, doi: 10.5194/gmd-9-4185-2016.

[11] NASA Langley Research Center, "POWER Data Access Viewer," power.larc.nasa.gov, accessed Aug. 20, 2026.

[12] X. Zhang et al., "Indices for monitoring changes in extremes based on daily temperature and precipitation data," WIREs Clim. Change, vol. 2, no. 6, pp. 851–870, 2011, doi: 10.1002/wcc.147.

[13] Task Force on Climate-related Financial Disclosures, "Recommendations of the TCFD," Financial Stability Board, 2017.

[14] IFRS Foundation, "IFRS S2 Climate-related Disclosures," ISSB, 2023.

[15] Chart.js v4.4.3, chartjs.org; Leaflet, leafletjs.com; basemap © OpenStreetMap contributors, © CARTO, accessed Aug. 20, 2026.