How Much Weather Certainty a Transmission Rating Policy Requires

Ricardo Marquez, Ph.D.
rmarquez123@gmail.com

November 10, 2026

Abstract

An overhead transmission line may carry more current when the wind is stronger, and a rating computed from a weather forecast is therefore only as good as the forecast. An operator who rates permissively relieves congestion but, on the occasions when the weather falls short of the forecast, runs the conductor above the temperature it was designed for and spends its service life faster than intended. An operator who rates conservatively protects the conductor and pays in generation cost and, past a point, in unserved energy. This report measures how well the weather must be known before that choice stops being forced. On a 73-bus test system driven by its own 8,784-hour load record, with contingency exposure taken from its own outage statistics, the two requirements can be met together up to a wind standard deviation of 0.75 m/s. The answer depends far more on the assumed climate than on anything in the network: over a plausible range of mean wind and air temperature the threshold spans 0.00 to 1.50 m/s. Uncertainty is therefore not paid for in feasibility on this system but in money: the least expensive rating that respects the design life costs \$0.000M to \$0.750M more per year than the rating that ignores it. Results are reported as that premium and as a break-even conductor replacement cost rather than as a multiple of design life, because those are the quantities an asset owner already knows.

1 Introduction

The current an overhead line may carry without exceeding its maximum conductor temperature depends on the weather around it. Wind cools the conductor by forced convection, air temperature sets the gradient that convection and radiation work against, and sunlight adds heat. A rating computed from measured or forecast weather is worth substantially more than a rating computed from conservative fixed assumptions, which is why dynamic line rating is attractive. The difficulty is that the weather is not known exactly, and the consequences of the two directions of error are not symmetric.

If the weather is better than assumed, the line is under-used and the system pays in generation cost. If it is worse, the line carries a current its true cooling cannot support, the conductor runs above its design temperature, and it loses service life at a rate that rises exponentially with temperature. The second error is invisible in the short run. Nothing trips, no alarm sounds, and the cost appears years later as a conductor that must be replaced earlier than planned.

This report asks a specific question: how small must the uncertainty in the wind be before an operator can choose a rating policy that neither sheds load nor consumes conductor life faster than the conductor was designed for? The question is answered on a network rather than on a single line, because the two requirements are coupled through the power flow. Rating one line conservatively pushes power onto others, which then run hotter. The life consumed by one conductor depends

on the uncertainty attached to a different one, and a calculation performed line by line cannot see that.

The calculations use standard relations throughout: the IEEE-738 heat balance for conductor temperature, an Arrhenius relation for loss of service life, and a linear direct-current dispatch with outage distribution factors for the contingency analysis. The work is a forward propagation of a stated weather distribution through those relations. No quantity is estimated from data, and no statistical model is fitted.

2 System and data

Figure 1 shows the system. It has 120 branches, of which 104 are overhead lines and 16 are transformers. Transformers carry no conductor and are excluded from the thermal and aging calculations while remaining in the power flow.

Three quantities are taken from the test system's own records rather than assumed, and each one changes the answer materially. The first is the load. The system ships a 8,784-hour record whose annual peak is 0.931 times the static load figure quoted in its bus table and whose median hour is only 0.456 times it. A study run at or above that static figure would be examining a loading the system never reaches, and would overstate annual aging by more than an order of magnitude, because conductor heating is concentrated in a small number of high-load hours. The record is grouped into 32 bands of net load — load less the renewable output available in that hour, shown sorted in Figure 4 — which preserves the hours of the year; Table 3 is insensitive to that choice, varying by a factor of 1.1345 across band counts from 16 to 128.

The second is contingency exposure. The system quotes a permanent outage rate and a repair duration for every branch, which together give 41.2 outage events per year and 710 branch-hours out of service. These weight the contingency part of the calculation, so the question of whether contingencies matter is settled by the system's own reliability data rather than by an assumption.

The third is the set of line ratings, which the dispatch uses directly. The system does not state which conductor each line is strung with, so each line is matched to the standard aluminum conductor and integer bundle whose static rating is closest to its own quoted rating (Table 1). The residual mismatch is small but not zero, and it is what causes the two line classes to sit at different temperatures when both are loaded to the same fraction of their rating.

The weather is not taken from data. It is a stated scenario, uniform across the whole system: mean wind 2.0 m/s at 45 degrees to the conductor, air temperature 35.0 °C, and solar irradiance 1,000 W/m^2. Uniformity is deliberate. With the same weather everywhere, every difference between one line and another in the results is attributable to the network, which is the object of study. The cost is that the report can say nothing about where an instrument should be placed.

3 Methods

Figure 2 shows the physical situation and Figure 3 the sequence of calculations.

Conductor temperature. The steady-state temperature T_c of a conductor carrying current I solves the heat balance

$$q_j(I, T_c) + q_s = q_c(v, \phi, T_a, T_c) + q_r(T_a, T_c) \tag{1}$$

$$q_j = I^2 R(T_c) \tag{2}$$

in which q_j is resistive heating, q_s solar gain, q_c convection and q_r radiation, all per unit length. Convection is the larger of a forced and a natural correlation, so the balance has a corner where

the two cross, at roughly 1.2 m/s for the conductors used here. The balance is monotone in T_c, so it is solved by bisection, which the corner does not trouble.

The rating policy. An operator who believes the wind is distributed as $|\mathcal{N}(\mu, \sigma^2)|$ publishes the p-th percentile of the resulting ampacity. A larger σ pushes that percentile down and a larger p pushes it up, so the pair (σ, p) fixes the rating. The dispatch then enforces $|f_k| \le R_k$ on every branch k.

Why the realized weather must be separated from the believed weather. The rating is defined as the current that holds the conductor at its maximum temperature, and the aging relation below is referenced to that same temperature. If the conductor were assumed to experience exactly the weather the rating assumed, a line loaded to its rating would age at exactly its design rate and no policy could ever exceed the design rate. The calculation would return the same answer whatever the inputs. The rating is therefore set on the believed distribution while the conductor is exposed to an independent draw from it, and the difference between the two is the entire subject of this report.

Loss of life. Conductor aging follows an Arrhenius relation in which continuous operation at the reference temperature T_{ref} consumes one whole service life in D_{ref} hours,

$$\frac{\mathrm{d}L}{\mathrm{d}t} = \frac{1}{D_{\mathrm{ref}}} \exp\left[\frac{E_a}{R_g}\left(\frac{1}{T_{\mathrm{ref}}} - \frac{1}{T_c}\right)\right] \tag{3}$$

with temperatures in kelvin. Taking the reference to be 100 °C over a forty-year design life makes the annual budget 0.02498 of a service life. Because the relation is convex in temperature, the loss of life expected over a spread of weather always exceeds the loss of life at the average weather; the ratio of the two, written J in Table 3, measures what treating the weather as a single number costs.

Annual accounting. Each line accumulates loss of life in three situations, weighted by how much of the year each occupies: normal operation, the 0.5 hour window after a neighboring branch trips and before the operator restores a secure state, and the remainder of that branch's repair, during which the system has been re-dispatched around it. The first is computed from the dispatch, the second from the linear outage distribution factors, and the third from a fresh dispatch on the reduced network.

What the contingency figures are, and are not. The dispatch used here is economic, not security-constrained: it enforces the rating on every branch of the intact network and says nothing about what the flows become after a trip. A real schedule is chosen so that no single outage overloads anything, which is exactly the constraint omitted here. Post-contingency flows taken from an unsecured schedule are therefore much larger than any operator would permit, and the loss of life they imply bounds the exposure rather than describing an operating outcome. It is reported as such and kept out of the threshold test below; the normal-operation figures carry no such caveat.

The threshold. A policy meets the two requirements when it sheds no load and leaves no line above its annual budget in normal operation. The threshold $\sigma^\star$ is the largest weather spread at which some policy still meets both. Above it, the operator must decide which requirement to break.

The weather distribution is represented by equally spaced quantiles rather than by pseudo-random draws. This makes the results reproducible without reference to a seed, covers the tail evenly, and keeps a threshold search from moving across the boundary for reasons unconnected to the physics.

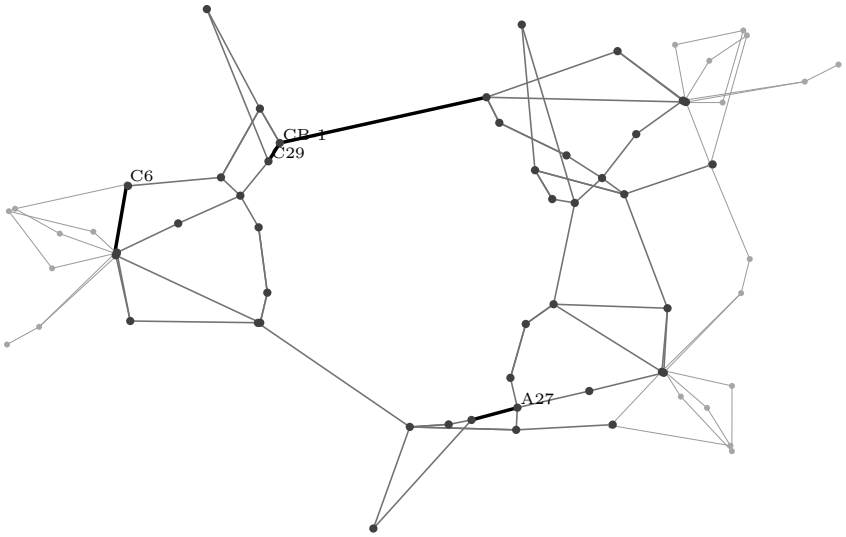


Figure 1: The system under study, drawn from the buses' own published coordinates: 73 buses across roughly 400 km, at two voltage levels, with 104 overhead lines and 16 transformers. Transformers carry no conductor and take no part in the thermal or aging work. The heavy lines are those that spend the most conductor life, and they are few: the exposure this report is about is concentrated in a handful of circuits rather than spread across the network.

Table 1: Conductors assigned to the test system. The source data gives line ratings but no conductor types, so each line is matched to the standard aluminum conductor and integer bundle whose static rating comes closest to its own. The fit error is what makes the two line classes reach different temperatures at the same fraction of rating.

Conductor	Bundle	kV	Rating (MW)	Fit error (%)	Lines
Bluejay 1113 ACSR	1	230	500	-3.5	67
Hawk 477 ACSR	1	138	175	-2.7	37

4 Results

Table 2 is the answer. For each level of weather uncertainty it gives the largest rating that can be published while every conductor still reaches forty years of service. Nothing in that table rests on a chosen level of conservatism: the rating is found by searching on the rating itself against the conductor's own design life, and the exceedance quantile the answer happens to fall at is reported as an output. That quantile does not settle anywhere, which is worth noticing on its own — a ratings rule that fixes the exceedance will waste capacity at some spreads and spend conductor life at others.

The usable rating erodes steadily as the weather becomes less certain, and what that erosion costs is worth reading carefully, because the obvious quantity is not the one that responds. Between perfect knowledge and the widest spread examined the publishable rating falls by about a quarter, and the annual generation cost rises by about half of one per cent. Nothing is shed and no conductor is worked harder; the conductor is held at the same life throughout. What changes is the renewable energy that cannot be delivered, which nearly triples over the same range.

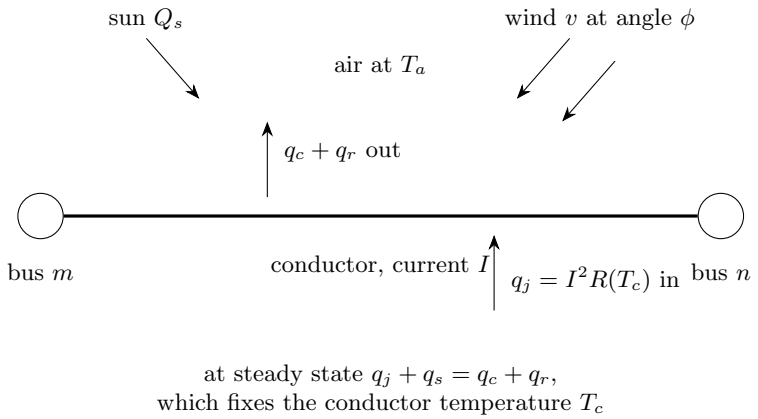


Figure 2: What the weather does to a conductor. Current heats it; wind and radiation cool it; the sun adds to the load. The balance fixes the conductor temperature, and the wind term dominates the cooling, so an error in the assumed wind is an error in the temperature.

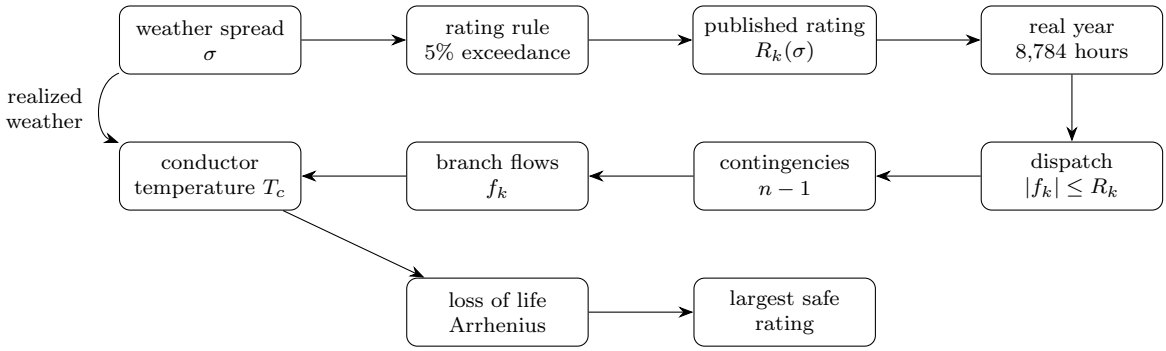


Figure 3: How the study computes its answer. The spread σ enters twice: once through the rating the operator publishes at a fixed 5 per cent exceedance, which the dispatch obeys, and once through the weather the conductor actually experiences. Separating the two is what makes the question non-trivial; treating them as one makes it vacuous.

That is the shape of the answer on a system whose cheapest generation is also its most variable. A conservative rating does not make the electricity much more expensive, because the thermal fleet that fills the gap is large and not badly priced. It spills wind and solar behind a constraint instead. A study that reported only the bill would have concluded that weather uncertainty costs almost nothing, and would have been looking at the wrong column.

The rating never falls below the value the line already carries. A static rating — the fixed figure the system ships, computed at conservative weather assumptions — sits at the nameplate quoted in the caption, and every row of Table 2 is above it. Even at the widest spread examined a rating targeted at the conductor's design life is worth more than the fixed one it would replace.

That is not true of a rule that fixes the exceedance instead. Table 3 shows a policy publishing the fifth percentile of the ampacity distribution running from thirty per cent above nameplate at perfect knowledge to twenty per cent below it at the widest spread: past about three quarters of a meter per second such a rule gives back a rating worse than the fixed one, and the measurement it rests on buys nothing. The loss belongs to the rule rather than to the weather, and it is the same point the wandering exceedance column made earlier: a fixed percentile does not track the

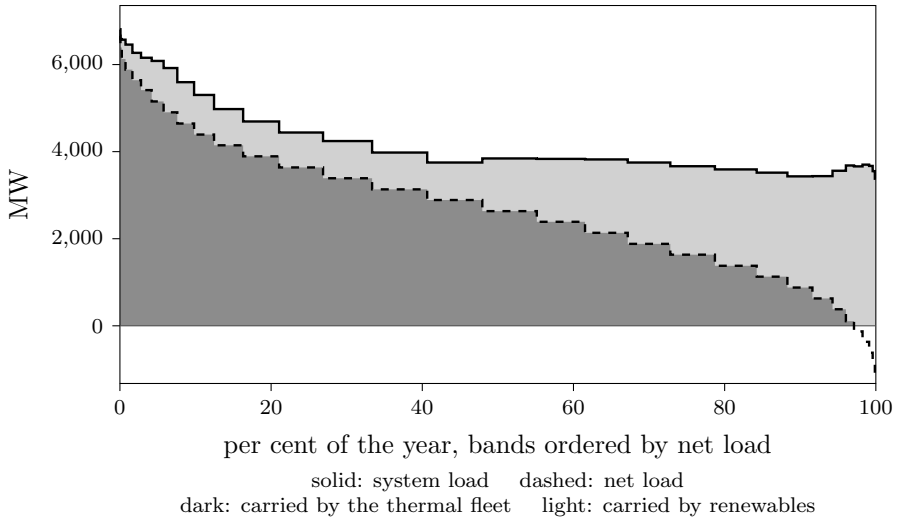


Figure 4: The year as the model sees it: 32 bands of net load, each drawn across the share of the year it occupies. The solid line is the band's mean system load and the dashed line its mean net load — what is left for the thermal fleet once the renewable output available in those hours is taken out — so the light band between them is that renewable output. The dashed line falls below zero in the last band, where renewables alone exceed demand. Only the dashed line is sorted, it being the quantity the bands are ordered on. The solid line is not monotone and is not expected to be: it rises again toward the right, because the bands with the most renewable output are not the bands with the least demand. A mild, windy, sunny afternoon carries moderate load and very large renewable output, which places it at the right-hand end of this axis while leaving it in the middle of the range on load alone. That is the whole case for regrouping the year: were load monotone here, ordering by load and ordering by net load would be the same thing.

constraint that actually binds.

Table 3 gives every policy examined. Reading down a column of fixed weather spread, a more permissive rating quantile lowers the dispatch cost and raises the loss of life; reading across a row of fixed quantile, a larger spread raises the loss of life while leaving the dispatch unchanged, because the dispatch sees only the published rating. The spread and the quantile are therefore not two independent controls. The spread is the condition the operator is in and the quantile is the response available.

Figure 6 shows that trade at one spread. The shortest life observed anywhere in the sweep is 269.6 years of service life, at $\sigma = 0.75$ m/s and the p10 rating, against a design value of forty.

The two requirements can be met together up to a wind standard deviation of 0.75 m/s. The best policy there is the p5 rating, at \$489.435M per year and 0.073 of the annual budget on the worst line. The first spread at which no policy suffices is 1.00 m/s.

Two extremes bracket every number above, and it is worth naming them. The upper one is perfect knowledge: if the weather were known exactly the first row of Table 2 is what could be published. The lower one is the wind stopping. IEEE-738 in still air is natural convection alone, and a conductor on this system then carries about three quarters of its nameplate — less than the static rating, which is computed against a light breeze rather than against no wind at all. Everything this report computes sits between those two, and every rating it recommends sits in the upper half.

Table 4 gives the consequence, and it is the least comfortable result here. The design-life

criterion is an expectation: it holds the conductor to forty years on average over the weather it is expected to see. It says nothing about a spell in which the wind simply stops. At the rating this study recommends for a spread of one meter per second, a still afternoon of eight to eighteen hours spends one per cent of the conductor; at the rating perfect knowledge would justify, the same one per cent goes in minutes. A rating chosen on expected life therefore needs a floor underneath it, and this study does not supply one — it quantifies why one is needed.

Figure 5 shows how unevenly the year contributes: bands are ordered by how fast they consume conductor life and their share of the annual total accumulated against the share of hours they occupy. The departure from the diagonal is what makes a peak-hour calculation multiplied out over a year so badly wrong.

Table 5 separates where each line spends its life, and the separation is the most striking thing here: under an unsecured schedule the contingency terms dominate the totals completely. That is a statement about the missing security constraint rather than about the weather, and those columns are upper bounds. What they do establish is that the exposure worth managing on this system sits in the minutes after a trip and not in normal running. Line C6 ages almost entirely in normal operation, whereas line C6 takes 98.3% of its loss of life from the periods when some other branch is out of service. Across the whole system contingencies account for 85.4% of annual aging, so they are not the dominant term in aggregate, yet for an individual line they can be very nearly the only term. A study restricted to normal operation would report that line as untroubled.

The cost of that erosion can be stated as a replacement cost. Loss of life here is a fraction of a whole conductor life, so the ratio of what a more permissive rating saves in a year to the extra loss of life it causes in that year is a break-even price per conductor, and the duration of the condition cancels out of it. A more permissive rating is worth taking when conductors cost less than that figure and worth refusing when they cost more, which is a comparison an asset owner can make against a number they already hold.

Figure 7 shows how the threshold moves with the assumed climate. It ranges from 0.00 to 1.50 m/s over the mean winds and air temperatures examined, and 0 of 20 climates admit no satisfactory policy at all. The threshold rises with mean wind, because a windier climate keeps the conductor further from the still-air regime where cooling collapses, and falls as the air warms, because a hotter climate leaves less thermal margin. The variation across this figure is larger than the variation produced by anything the operator controls, which means the threshold is a property of the climate at least as much as of the network and should be quoted with the climate it belongs to.

5 Discussion

The largest gap is the absence of a security-constrained dispatch. Because the schedule is chosen without regard to what a single outage would do, the post-contingency flows it produces are far beyond what a real operator would allow, and the conductor temperatures that follow are correspondingly extreme. Quantifying contingency aging properly needs a schedule secured against single outages and evaluated against emergency rather than continuous ratings; the test system publishes both, and this work uses only the continuous one. Until that is done the contingency figures here bound the exposure rather than predict it. The normal-operation results do not depend on any of this.

The sharpest limitation in the weather treatment is this. The realized wind in each hour is drawn independently of every other hour, so the model cannot produce a still, hot spell lasting several days. That is precisely the situation in which conductors are damaged: load is high, cooling

Table 2: The largest rating that can be published while every conductor still reaches its design life of forty years. The test system has two classes of line, rated 175 MW and 500 MW as it ships, so the answer is given for each. The rating is found by searching on the rating itself against the conductor's own design life, so nothing here rests on an assumed level of conservatism. Curtailment is renewable energy that was available and could not be delivered, in GWh per year; it is the quantity that moves, while the generation cost, in millions of dollars per year, barely does.

σ (m/s)	138 kV line (MW)	230 kV line (MW)	Life (yr)	Curtailed	Cost
0.00	272	778	40.5	102	486
0.25	267	764	40.4	102	487
0.50	248	709	40.2	125	487
0.75	219	627	40.6	202	487
1.00	208	595	40.5	231	488
1.50	201	573	40.6	253	488
2.00	200	570	40.6	256	488

is poor, and the condition persists long enough for the exponential in the aging relation to do real work. Independent hourly draws average that away. The figures here should be read as a lower bound for that reason, and a serial weather model is the single most valuable extension.

The weather is also uniform in space, which isolates the network effect cleanly but leaves the report unable to say where an instrument should go. Any such recommendation requires a spatial weather field, which this work does not have.

Two error sources are excluded and both are larger than they appear. The flows here come from a dispatch model; in operation they would come from a state estimator, and a five per cent error in an estimated flow changes the inferred loss of life by roughly a factor of two, since temperature rises with the square of current and aging rises exponentially with temperature. The conductor properties are catalog values; surface emissivity alone ranges over a factor of four in aging between a new conductor and a weathered one, and it drifts over the life of the line. Both of these are systematic where weather error is random: an incorrect emissivity is incorrect in the same direction in every hour of the year, so it shifts the annual total rather than averaging out of it. Recovering them from measurements is a different kind of problem from the one solved here and is left to separate work.

It is worth stating plainly what the uncertainty in this report is and is not. It is not introduced by the method. The wind does vary, the conductor properties are not the catalog values, and the flows are estimates. A rating computed from fixed assumptions does not remove any of that; it declines to report it. The quantities here are a statement of how large an exposure was already being carried.

Finally, the test system offers only two distinct line ratings, so the diversity available among conductors is between two classes rather than within them, and the dispatch is a linear direct-current approximation without losses or reactive power. Contingency analysis covers single outages only; simultaneous outages are not modeled, which the exposure figures show to be a small omission here but would not be on a system with longer repair times.

6 Conclusion

An operator choosing how conservatively to rate a line is choosing between generation cost and conductor life, and the exchange rate between them is set by how well the weather is known. On the system examined, the two requirements can be met together up to a wind standard deviation of

Table 3: Every rating policy examined, at the stated weather. Columns, in order: the weather spread in m/s; the exceedance the rating is published at; the rating itself, as a multiple of the line's own rating in the source data; curtailment in GWh per year; the extra generation cost in millions of dollars per year; and a dimensionless factor described below. Curtailment is renewable energy that was available and could not be delivered, and it is the quantity that responds: it moves by a factor of nearly three down the table. Generation cost is given as the difference from the cheapest policy here, because the total is about 487 million dollars a year in every row and almost all of that is the unavoidable cost of serving load. The last column is dimensionless: the factor by which loss of life computed at the average wind understates the loss of life expected over the spread of winds, on an identical dispatch. No policy examined shed any load, and none brought a conductor near its forty-year design life, so neither is tabulated.

σ	Exceedance	Rating	Curtailed	Δ cost	Understatement
0.00	p5	1.302	174	0.00	1.00
0.25	p5	1.224	212	0.45	1.06
0.25	p10	1.242	204	0.34	1.06
0.50	p5	1.134	259	1.16	1.24
0.50	p10	1.176	236	0.78	1.32
0.75	p5	1.028	315	2.26	1.96
0.75	p10	1.100	279	1.51	2.38
1.00	p5	0.905	380	4.42	1.96
1.00	p10	1.017	321	2.40	3.11
1.50	p5	0.801	458	7.10	1.99
1.50	p10	0.914	374	4.24	2.76
2.00	p5	0.797	462	7.14	2.04
2.00	p10	0.896	389	4.63	2.69

0.75 m/s. Where a threshold exists it is far more sensitive to the assumed climate than to anything in the network, and in some climates it does not exist at all, which is a useful answer in itself: there no instrument purchase helps and the loading must change instead.

Two results are worth separating from the threshold. Treating the weather as a single number rather than a distribution understates loss of life by a factor of up to 3.1 on an identical dispatch. And individual lines age for entirely different reasons: some in normal operation, others almost exclusively in the periods after a neighboring branch trips. The second group is invisible to any analysis of normal operation, and identifying it requires the network, the contingency set and the outage statistics together.

7 Appendix A: Reproduction

The study is run with

```
python docs/terra_graph_engine/case_studies/transmission/
       weather_certainty_aging/run.py
python docs/terra_graph_engine/case_studies/transmission/
       weather_certainty_aging/make_report.py --pdf
```

which writes `results.json` and renders this document from it. Every figure in the text is read from that file. The calculation is deterministic and carries no seed. Commit `60af0c478a2b`, Python 3.13.1, NumPy 2.3.5.

Table 4: The calm-wind extreme. Every other result in this report is an expectation over an assumed wind distribution; this one assumes no distribution at all, because IEEE-738 in still air is natural convection alone and the answer is definite. The last column is how long a still spell has to last, at that rating, to spend one per cent of the conductor. An annual multiple says nothing useful about a spell lasting an afternoon.

Line	Rating	× nameplate	Calm T_c (°C)	Hours per 1% of life
138 kV	nameplate	1.00	138	128.1
	$\sigma = 1$	1.19	177	8.3
	perfect knowledge	1.56	270	0.1
230 kV	nameplate	1.00	131	228.8
	$\sigma = 1$	1.19	165	17.6
	perfect knowledge	1.56	249	0.1

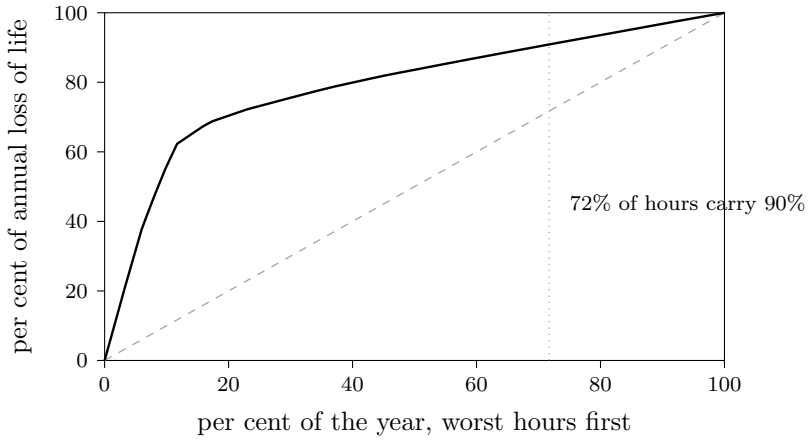


Figure 5: How unevenly the year ages the conductors. Bands of the year are ordered by the rate at which they consume conductor life, worst first, and the curve accumulates their share of the annual total. The diagonal is what an evenly aging year would look like. The departure from it is the reason a calculation performed at the peak hour and multiplied by the hours in a year overstates the answer by more than an order of magnitude, and the reason the bands must resolve the hot end finely.

References

[ieee738] IEEE Std 738-2012, Standard for Calculating the Current-Temperature Relationship of Bare Overhead Conductors, IEEE, 2013.

[ieee1283] IEEE Std 1283-2013, Guide for Determining the Effects of High Temperature Operation on Conductors, Connectors, and Accessories, IEEE, 2013.

[rtsgmlc] C. Barrows et al., The IEEE Reliability Test System: A Proposed 2019 Update, IEEE Transactions on Power Systems, vol. 35, no. 1, pp. 119–127, 2020.

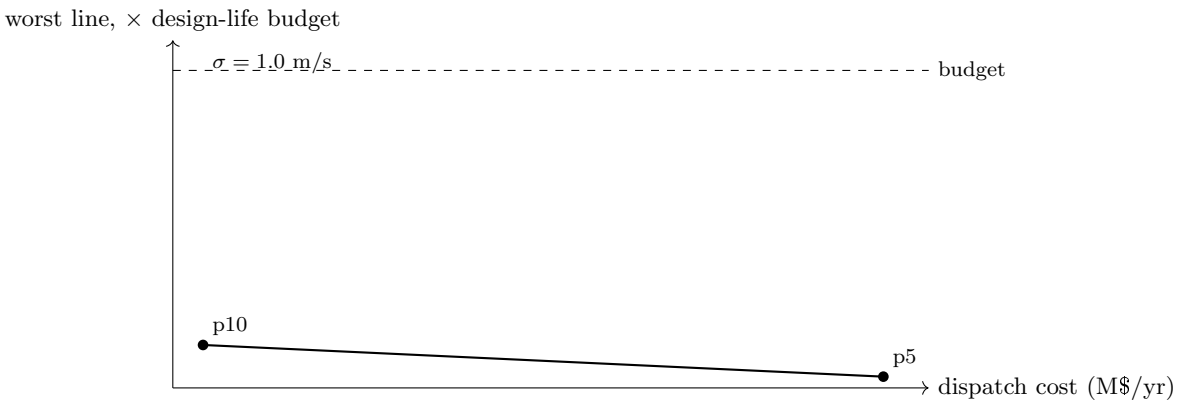


Figure 6: The trade at one level of weather knowledge. Each marker is a rating quantile: a more permissive rating moves left (cheaper) and up (more conductor life spent). The dashed line is the design-life budget. The curve is steep rather than gradual, so the choice between adjacent policies is not a small one.

Table 5: Where each line spends its life, at $\sigma = 1.0$ m/s. The second column is the thermal service life the line would reach on normal operation alone, against a design value of forty years, reported as a bound above 200 years for the reason given under Table 3. The three percentage columns split the loss of life by when it is incurred, and the two contingency columns are upper bounds for the reason given in Section 4. They are physically different situations and imply different remedies. Some lines age almost entirely in normal operation; others age almost entirely in the minutes after a neighboring line trips, and no study of normal operation would see them at all.

Line	Thermal life in normal use (yr)	Normal (%)	After a trip (%)	During repair (%)
C6	>200	2	98	0
CB-1	>200	89	4	8
A27	>200	91	1	8
C29	>200	92	0	8
A28	>200	90	2	9
A11	>200	86	7	7
C2	>200	90	1	8
B12-1	>200	89	2	9

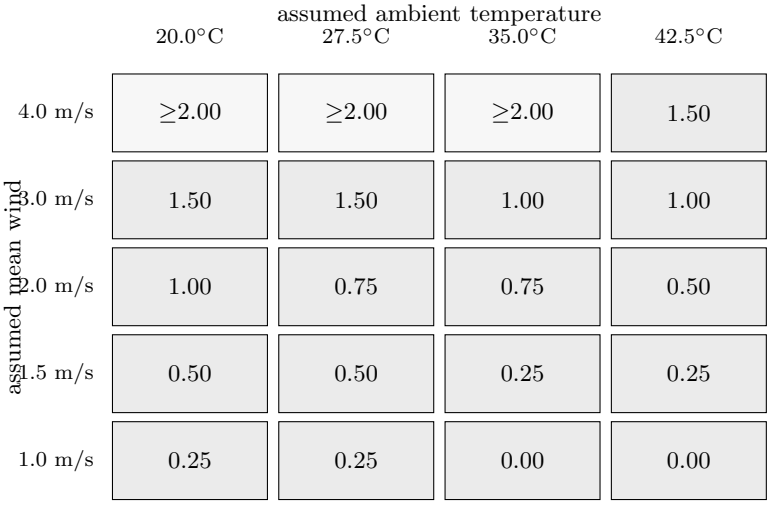


Figure 7: The largest weather spread σ (m/s) at which some rating policy still satisfies both requirements, over the assumed weather. Darker cells marked *none* have no such policy at any σ: the loading itself is the problem there, and no instrument fixes it. The threshold moves more across this table than it does across anything the network controls.