# What a Dynamic Line Rating Mitigates, and What Weather Uncertainty Takes Back

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

November 10, 2026

**Abstract**

An overhead line may carry more current when the wind is stronger, so a rating computed from the weather is worth more than a fixed one — but only where a limit is actually reached, and only to the extent the weather is known. This report measures both. On a 73-bus test system driven by its own 8,784-hour load and renewable records, a rating binds in 13.0 per cent of the hours when everything is in service, on 2 lines of 104; and a single outage would leave something above its emergency rating in 1.4 per cent of the hours when a branch is out. Everywhere else a rating is a number nothing presses against. Over those states, the answer is bounded by two physical cases and no convention. Knowing the weather exactly, securing the system against any single outage costs 2.39 million dollars a year; no instrument does better than knowing the answer. Knowing nothing about it, the only defensible rating is the one calm air permits, and the same system costs 7.03 million. Every real instrument lies between those two, a factor of 2.94. The price is not uniform in the spread: it is 4.24 million a year per m/s between 0.50 and 0.75 m/s, and 0.20 million between 1.50 and 2.00 m/s, where the rating has already fallen to within 2 per cent of the calm-wind case and a better wind estimate buys almost nothing. Precision is worth paying for over a specific range and not outside it.

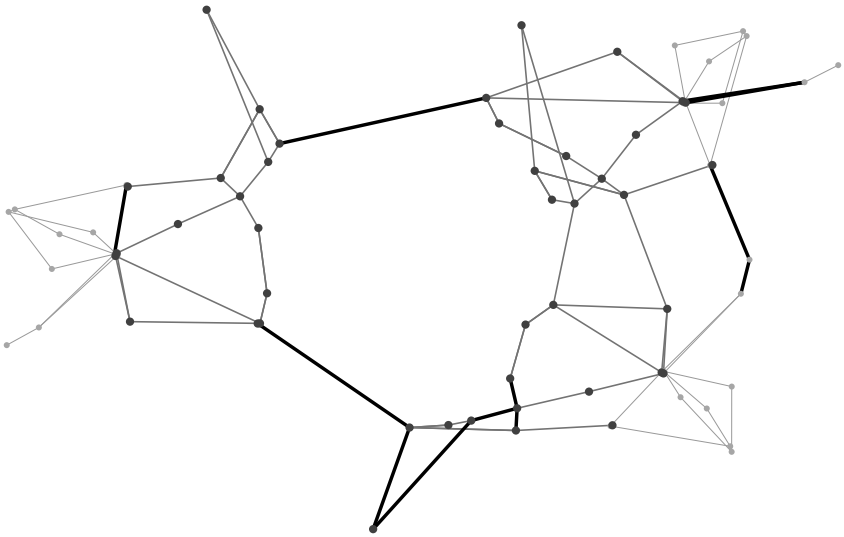
## 1 Introduction

The current an overhead line may carry without exceeding its maximum conductor temperature depends on the weather around it. Wind cools the conductor, air temperature sets the gradient that cooling works against, and sunlight adds heat. A rating computed from measured or forecast weather is therefore worth more than one computed from fixed conservative assumptions, which is why dynamic line rating is attractive. Two things complicate it.

The first is that a rating only matters where a limit is reached. A transmission system spends most of its life well inside its limits, and raising a limit that nothing is pressing against buys nothing at all. Any honest account of what a dynamic rating is worth therefore has to say first where, and how often, anything is at stake.

The second is that the weather is not known exactly. A rating that is to be published has to be conservative enough that the true capability is rarely below it, and the less certain the weather, the more conservative that rating must be. So the value of a measured rating is bounded by the quality of the measurement, and the right question is not whether it beats the fixed rating but how much of the attainable benefit a given quality of measurement keeps. The fixed rating cannot settle that, because it encodes an assumed wind and an assumed ambient chosen by convention: it is a number to be located on the axis, not the origin of it. The origin is perfect knowledge.

This report measures both on a test system, using the system's own load record, its own renewable output and its own outage statistics. The calculations use standard relations throughout: the IEEE-738 heat balance for conductor temperature, a linear direct-current dispatch, and outage distribution factors for the contingency analysis. No quantity is estimated from data and no statistical model is fitted.



dark nodes: 230 kV  light: 138 kV  dashed: transformers  heavy: the lines where a limit is reached

Figure 1: The system under study, drawn from the buses' own published coordinates: 120 branches across roughly 400 km at two voltage levels, of which 104 are overhead lines and 16 are transformers. Transformers carry no conductor and take no part in the thermal work. The lines drawn heavily are the only ones where any limit is reached in any state the year visits; everywhere else a rating is a number nothing presses against.

## 2 System and data

Figure 1 shows the system: 120 branches, of which 104 are overhead lines and 16 are transformers. Transformers carry no conductor and take no part in the thermal work.

Three quantities come from the system's own records rather than from assumption. The load is a 8,784-hour series whose annual peak is 0.931 times the static figure quoted in the bus table and whose median hour is 0.456 times it, so a study run at the static figure would examine a loading the system never reaches. The renewable output is the published hourly series for all its wind and solar units; running them at their nameplate instead, as is tempting, supplies several gigawatts around the clock and puts power on the wrong lines. Figure 2 shows both as duration curves, and the shape is the reason the bands are cut on net rather than gross load: renewables carry 35.8 per cent of the year's energy, enough that net load falls below zero in the least-loaded hours and that two hours at the same load can be a sunny afternoon and a still night with nothing in common for the thermal fleet to do. And the outage statistics give 41.2 events a year and 710 branch-hours out of service, so the probability that some branch is out at a given moment is 0.081.

The system publishes three ratings for every branch: a continuous rating and two emergency ratings, the longer of which is 10 to 27 per cent above the continuous one. The emergency rating

is the limit that applies after an outage, and using the continuous rating there — as is easy to do, because it is the one a dispatch model usually carries — overstates every violation by that margin.

The system does not say which conductor each line is strung with, so each is matched to a standard aluminum conductor and integer bundle (Table 1). The weather is not taken from data at all: it is a stated scenario, uniform across the system — mean wind 2.0 m/s at 45 degrees, air temperature 35.0 °C. Uniformity is deliberate: with the same weather everywhere, every difference between one line and another is attributable to the network. The cost is that the report can say nothing about where an instrument should be placed.

Table 1: Conductors assigned to the test system, which publishes line ratings but no conductor types. Each line is matched to the standard aluminum conductor and integer bundle whose static rating comes closest to its own; the fit column is the residual gap.

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

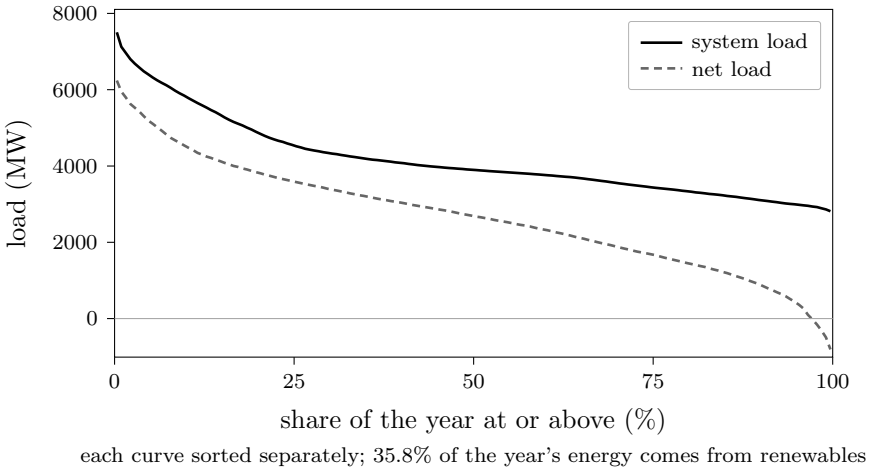



Figure 2: The year the study is built on, at the record's own hourly resolution (8,784 hours). Load is the system's own record; net load is that less the renewable output available in the same hour, which is what the thermal fleet and the network actually see, and is what the 24 dispatch bands are cut on. Net load falls below zero in the least-loaded hours, where available renewable output exceeds demand. Each curve is sorted independently, so the vertical gap between them is the difference between the two curves at the same rank and not the renewable output of any single hour; the energy share in the legend is the figure to read for that.

## 3 Methods

Figure 3 shows the sequence.

**States.** A state $s$ is a band $b$ of net load paired with a topology $\tau$: the intact network, or the network with one branch $k$ out. Net load is system load less the renewable output available in that hour, which is what the thermal fleet and the network actually see. A state's weight is the

band's hours times the probability of the topology, and the intact case takes whatever probability the outages leave.

$$h_s = h_b \, p(\tau) \tag{1}$$

$$p(\text{intact}) = 1 - \textstyle\sum_k p_k \tag{2}$$

with $p_k$ the probability that branch $k$ is out at a random hour, formed from its own outage rate and repair duration. Each state is characterised by its most loaded line, measured against the limit that applies in that state.

$$\ell_s = \max_i \frac{|f_i(s)|}{L_i(\tau)} \tag{3}$$

where $f_i$ is the flow on branch $i$ and $L_i$ is its continuous rating when the network is intact and its emergency rating when something is out.

**Why the two classes are not comparable.** The dispatch enforces the continuous rating on the intact network, so an intact state cannot exceed it and a reading of 1.0 means a rating is binding — an ordinary operating condition, and the one a dynamic rating relieves. A post-outage state is held to its emergency rating only if the schedule was chosen to hold it there, so a reading above 1.0 there is a limit exceeded. The two are reported separately for that reason.

**A secure schedule.** A schedule that merely minimizes cost says nothing about what happens when a branch trips, and its post-contingency flows are larger than any operator would run. The schedule here is instead chosen so that no single outage leaves anything above its emergency rating, by solving, screening every outage with the outage distribution factors, adding the limits that were violated, and solving again until none is. The condition added for a monitored branch $i$ and an outaged branch $k$ is

$$\left| f_i + \mathrm{LODF}_{ik} \, f_k \right| \le E_i \tag{4}$$

with $\mathrm{LODF}_{ik}$ the fraction of $k$'s flow that lands on $i$ when $k$ opens and $E_i$ the emergency rating of $i$. Both flows are already linear in the bus angles the dispatch solves for, so (4) is a pair of inequality rows rather than a second optimization, and only the limits that actually bind cost anything.

**The risk cost.** The difference between the secure schedule and the cheapest one is what security costs.

$$\Pi(\sigma) = C_{\mathrm{sec}}(\sigma) - C_{\mathrm{econ}}(\sigma) \tag{5}$$

Both terms are generation cost over the year at the same limits, so (5) isolates security from the rating. The dispatch also prices unserved energy at the value of lost load; that term is removed from both sides, and shed energy is reported separately in MWh. It is the quantity a dynamic rating reduces — a higher limit leaves the secure schedule less constrained — and therefore the quantity weather uncertainty stops it from reducing. Nothing in this report is normalized to the fixed rating: it appears as one row among the others, printed so a reader can locate the number they already hold. The two ends of the answer are physical — perfect knowledge and calm air — and the fixed rating is neither of them.

**The published rating.** A conductor's ampacity at a given wind is what the IEEE-738 steady-state heat balance allows at its maximum operating temperature: convection and radiation carry away what the sun and the current put in.

$$q_c(T_{\max}, v) + q_r(T_{\max}) = q_s + I^2 R(T_{\max}) \tag{6}$$

$$I(v) = \sqrt{\frac{q_c(T_{\max}, v) + q_r(T_{\max}) - q_s}{R(T_{\max})}} \tag{7}$$

Convection $q_c$ is the only term the wind enters, and it is the larger of a forced and a natural correlation, which cross near 1.16 m/s. A rating that is to be published cannot use (7) at a wind that is merely expected, because the wind is not known: it takes a low quantile of the ampacity the wind distribution induces, so that the true capability falls below the published figure only rarely.

$$v \sim \left|\mathcal{N}(\mu, \sigma^2)\right| \tag{8}$$

$$I_{\mathrm{pub}}(\sigma) = Q_{\varepsilon}\left[I(v)\right] \tag{9}$$

This report uses $\varepsilon = 5$ per cent, which is the same statement as publishing at the fifth percentile of the ampacity distribution and is written p5 in the figures and tables that follow. That is a stated choice and not a requirement: the facility-ratings rules ask for a documented methodology rather than a particular level of conservatism. The two physical ends of the study are both read off (9): perfect knowledge is $\sigma = 0$, where the quantile collapses onto $I(\mu)$, and calm air is $I(0)$, which takes no distribution and no quantile at all.

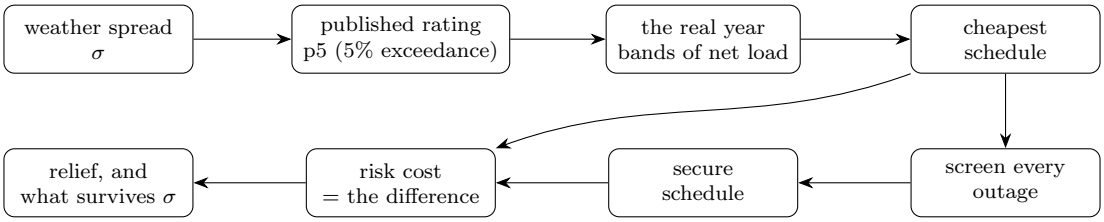



Figure 3: How the study computes its answer. The rating a given quality of weather permits is applied to every limit; the year is then dispatched twice, once as cheaply as possible and once so that no single outage overloads anything. The difference between the two is what security costs, and that is the quantity a better rating reduces. Screening the outages is linear algebra rather than a second optimization, which is what makes securing every band of the year affordable.

## 4 Results

Table 2 says where anything is at stake, and the answer is: almost nowhere. With everything in service a rating binds in 13.0 per cent of the hours, and on only 2 lines of 104. With a branch out, something is left above its emergency rating in 1.4 per cent of those hours, on 12 lines, reaching 1.88 times the emergency rating at worst. Better than half of all state-hours sit below half loading, where a rating is a number nothing presses against.

That concentration is the reason for conditioning on states at all. An annual average over this system is an average over a year that is, for the most part, uneventful. Figure 4 shows where that share is read from, and sharpens it. The two branches do not share the burden: CB-1 is at its rating in 13.0 per cent of intact hours; A27 is at its rating in 3.1 per cent of intact hours. A system-wide figure reports the pair as though they were alike, and an instrument bought on that figure would be sited on the evidence of one line while being justified by two. Table 3 and Figure 5 give what the quality of the weather does about it. The reference point is perfect knowledge, not the fixed rating: knowing the wind exactly permits a rating 1.302 times the conventional fixed one and makes securing the system cost 2.39 million dollars a year. Nothing can beat that, because no instrument does better than knowing the answer, so every other row is read as a distance from it.

The other end is equally physical, and it is not the noisiest row of the sweep. An operator who knows nothing about the wind cannot publish a rating that assumes any, so the defensible choice

Table 2: Where anything is at stake. A state is an hour and a topology; hours are grouped into bands of net load and topologies are the intact network and each credible single outage, weighted by the system's own outage statistics. Loading is the most loaded line measured against the limit that applies in that state — the continuous rating with everything in service, the emergency rating with something out. The two classes are listed separately because the same number means different things in them: the dispatch enforces the continuous rating, so an intact state at 1.0 has a rating binding, while a post-outage state above 1.0 has a limit exceeded.

| Class | Loading (× limit) | Share (%) | Hours (h/yr) | Lines (count) |
|---|---|---|---|---|
| intact | 0.000–0.500 | 44.4 | 3,583 | 3 |
| intact | 0.500–0.800 | 34.6 | 2,793 | 4 |
| intact | 0.800–0.900 | 8.0 | 645 | 2 |
| intact | 0.999–1.001 | 13.0 | 1,050 | 2 |
| post-outage | 0.000–0.500 | 56.6 | 400 | 21 |
| post-outage | 0.500–0.800 | 29.4 | 208 | 20 |
| post-outage | 0.800–0.900 | 11.9 | 84 | 14 |
| post-outage | 0.900–0.950 | 0.3 | 2 | 11 |
| post-outage | 0.950–0.999 | 0.4 | 3 | 9 |
| post-outage | 1.001–over | 1.4 | 10 | 12 |

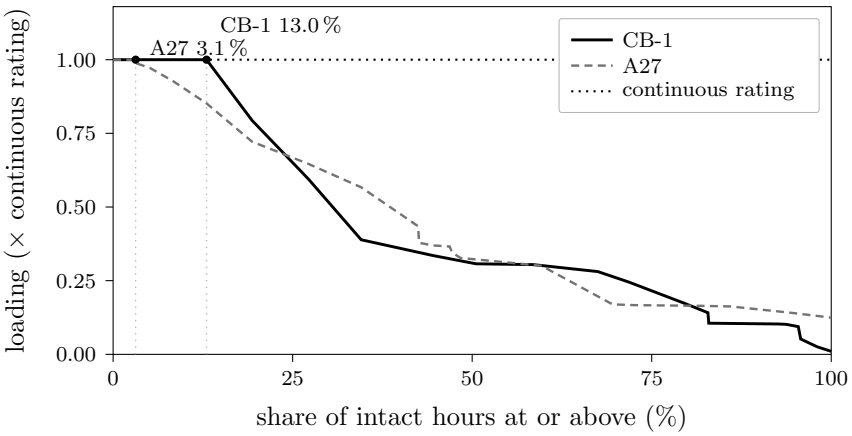



Figure 4: The headline share, in the data it is read from. Each curve gives the share of intact hours one branch spends at or above a given loading, measured against its continuous rating. The dispatch enforces that rating, so no curve rises above 1.0; the flat top of each is the hours where the rating is binding rather than exceeded, and the marked point is where it ends. Those two points are the figure: CB-1 is at its rating in 13.0 per cent of intact hours; A27 is at its rating in 3.1 per cent of intact hours, so the 13.0 per cent quoted for the system is one branch's number and the other is an order of magnitude rarer. Below about 0.9 the curves carry no claim — they cross, and a rating is a number nothing presses against there. The curves are stepped because the dispatch solves one schedule per net-load band; the steps are the bands, and the study's resolution is the band rather than the hour.

is the rating calm air permits, with the conductor cooling by natural convection alone. It takes no distribution, no spread and no exceedance — there is nothing to take a percentile of. That rating is 0.792 of the conventional fixed one and securing the system at it costs 7.03 million a year. The whole answer therefore lies between 2.39 and 7.03 million, a factor of 2.94, and every instrument

that could be installed lies somewhere inside it.

The sweep confirms that the calm case is the limit it approaches rather than a separate assumption: at 2.00 m/s the rating is already 6.96 million against 7.03, covering 98 per cent of the bracket. A conservative quantile of a wide enough wind distribution is calm air; the two extremes meet.

Two quantities carry the rest of the section. The price of the uncertainty is the distance from perfect knowledge, and the price of a unit of spread is the slope of that distance between two rows of the sweep.

$$\Delta(\sigma) = \Pi(\sigma) - \Pi(0) \tag{10}$$

$$m_{[\sigma_1,\sigma_2]} = \frac{\Pi(\sigma_2) - \Pi(\sigma_1)}{\sigma_2 - \sigma_1} \tag{11}$$

with $\Pi$ from (5). The last column of Table 3 is (10); the paragraph below is (11).

The price of spread is not uniform, and that is the part an operator can act on. Between 0.50 and 0.75 m/s each additional m/s of spread costs 4.24 million a year; between 1.50 and 2.00 m/s it costs 0.20 million. The flat end is physical rather than numerical: the forced- and natural-convection correlations cross near 1.16 m/s, so once the wind a 5 per cent exceedance permits has dropped below that, the conductor is cooling by natural convection and a better wind estimate changes the rating very little. Sharpening a poor estimate is close to worthless; sharpening a fair one is where the money is.

For a reader holding the conventional fixed rating, it is 4.64 million a year and falls between the 0.75 and 1.00 m/s rows. It is printed so the number can be located, and for no other purpose. A static rating is a convention about assumed weather — a stated wind, a stated ambient — and not a property of the conductor, so where a measured rating happens to pass it says which of two conventions is more conservative and nothing about whether the measurement is worth making.

Figure 5 draws the sweep at two exceedance levels, and the gap between them is worth stating on its own. Publishing at p10 rather than p5 — 10 per cent exceedance rather than 5 per cent — saves 1.50 million a year at 2.00 m/s — 33 per cent of everything the uncertainty costs at that spread. Put in the units of the rest of this report: at 1.00 m/s the change of level is worth 0.93 million, about what sharpening the wind estimate from 1.00 to 0.75 m/s would buy (1.05 million). The exceedance level is a policy choice with no standard behind it, and on this system it is worth about as much as a substantial improvement in the measurement. A reader comparing ratings across systems should check it before comparing anything else.

A rating is a rule, and the system has been read through one rule so far: publish a low quantile of the ampacity the wind distribution induces, so the conductor rarely exceeds a stated temperature. There is a second rule, and it protects something else. Running hot costs aluminium its tensile strength permanently, and the loss is a function of time and temperature together: above an onset near 95 °C the conductor is spending strength it never recovers, below it nothing happens however long it is held there, and a design life is an allowance of that loss. A rating can be solved to spend the allowance exactly rather than to hold a temperature.

Table 4 runs both on the same weather, the same dispatch and the same limits, so a pair of rows at one spread differs only in the rule. They neither track each other nor part company in one direction: the budget rating is 0.876 of the quantile rating at 0.75 m/s, and 1.071 of it at the spread where it is most generous.

The two rules charge uncertainty in different currencies, which is what makes them cross. The quantile rule gives up the wind: it publishes at a low percentile of the wind while holding the temperature at the maximum, so as the spread widens the wind it dares credit collapses toward

still air. The budget rule gives up the temperature: it credits the mean wind throughout and pays for the uncertainty by lowering the target until the scatter stops overspending the allowance. Which sacrifice is cheaper depends on the spread, and here it changes hands twice.

That difference is also the budget rule's exposure. Crediting the mean at every spread makes it a price on a *noisy* estimate rather than on a *biased* one, while the quantile rule hedges the wind itself and is the more defensive of the two against a mean that is simply wrong. Neither dominates, and this report does not adjudicate between them: it measures how far apart they land, and says why.

The overload column in Table 3 is the same story told physically rather than in money: the number of post-contingency limit exceedances the cheapest schedule would leave, which is what securing the system exists to prevent, and which rises with the spread for the same reason the cost does.

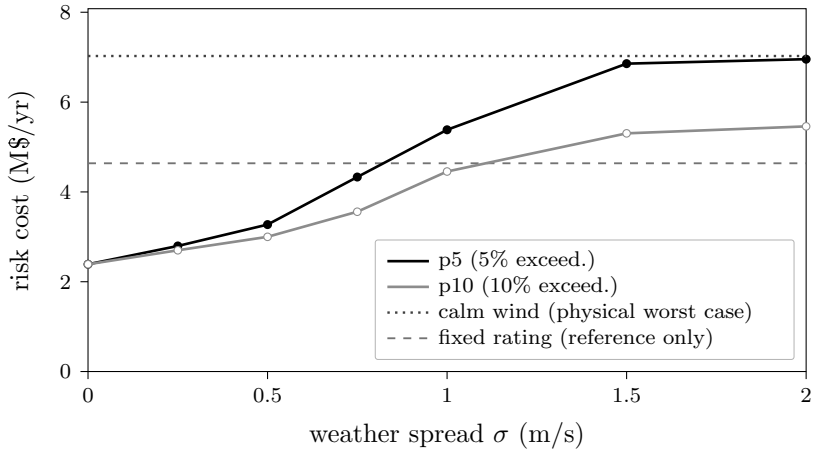



Figure 5: What securing the system costs, against how well the weather is known. The curve runs between two physical bounds. Its left-hand end is perfect knowledge, which no instrument beats; the dotted line is calm wind, the rating still air permits when nothing is known and the conductor cools by natural convection alone. The quantity to read is the height of the curve above its left-hand end, which is the price of the uncertainty, and the slope, which is the price of a unit of spread. The slope is steepest in the middle and nearly flat at the right, where the curve has already reached the dotted line and a sharper estimate buys almost nothing. The two solid curves are the same calculation at two exceedance levels. No standard fixes the level, so both are drawn: the vertical gap between them is the part of the answer that rests on that policy choice rather than on the weather, and it is not small. The dashed line marks the conventional fixed rating: a reference point a reader will recognize, not a bound and not a baseline, and the curve passing it carries no verdict.

Table 3: The risk cost, and what the quality of the weather does to it. Units are on the second header row. Risk cost is what securing the system against any single outage costs over the cheapest schedule. Rating is the multiple of each line's conventional fixed rating that the stated quality of weather permits at a 5 per cent exceedance (p5); the fixed rating is carried as a reference point a reader will recognize, not as a baseline. Curtailment is renewable energy that was available and could not be delivered. Overloads counts the post-contingency limit exceedances the cheapest schedule would leave, which securing the system is what prevents. The last column is each row's distance from perfect knowledge of the weather: the price of the uncertainty itself. The table is bounded by two physical cases and no convention. The first dynamic row is perfect knowledge, which no instrument can beat. The calm-wind row is perfect ignorance: with nothing known about the wind the only defensible rating is the one still air permits, taken with no distribution and no exceedance. Every real instrument lies between the two. The *dynamic* rows publish a low quantile of the ampacity distribution at the maximum operating temperature; the *annealing* rows publish instead the rating whose realised temperature scatter spends the design-life allowance exactly. Those are two different rules for the same job and are compared in Table 4; here they are simply two bases the same year was dispatched on.

| Basis | $\sigma$ (m/s) | Rating (× fixed) | Risk cost (M$/yr) | Curtailed (GWh/yr) | Overloads (count) | Cost of not knowing (M$/yr) |
|---|---|---|---|---|---|---|
| fixed rating | – | 1.000 | 4.64 | 531 | 118 | – |
| dynamic p5, sigma=0.00 | 0.00 | 1.302 | 2.39 | 332 | 41 | 0.00 |
| dynamic p5, sigma=0.25 | 0.25 | 1.224 | 2.80 | 371 | 58 | 0.41 |
| dynamic p5, sigma=0.50 | 0.50 | 1.134 | 3.27 | 422 | 67 | 0.88 |
| dynamic p5, sigma=0.75 | 0.75 | 1.028 | 4.33 | 507 | 105 | 1.94 |
| dynamic p5, sigma=1.00 | 1.00 | 0.905 | 5.38 | 623 | 141 | 2.99 |
| dynamic p5, sigma=1.50 | 1.50 | 0.801 | 6.86 | 803 | 179 | 4.47 |
| dynamic p5, sigma=2.00 | 2.00 | 0.797 | 6.96 | 808 | 185 | 4.57 |
| dynamic p10, sigma=0.00 | 0.00 | 1.302 | 2.39 | 332 | 41 | 0.00 |
| dynamic p10, sigma=0.25 | 0.25 | 1.242 | 2.70 | 362 | 54 | 0.31 |
| dynamic p10, sigma=0.50 | 0.50 | 1.176 | 3.00 | 394 | 65 | 0.61 |
| dynamic p10, sigma=0.75 | 0.75 | 1.100 | 3.56 | 449 | 84 | 1.17 |
| dynamic p10, sigma=1.00 | 1.00 | 1.017 | 4.45 | 516 | 109 | 2.07 |
| dynamic p10, sigma=1.50 | 1.50 | 0.914 | 5.30 | 613 | 141 | 2.92 |
| dynamic p10, sigma=2.00 | 2.00 | 0.896 | 5.46 | 635 | 148 | 3.07 |
| annealing, sigma=0.00 | 0.00 | 1.302 | 2.39 | 332 | 41 | 0.00 |
| annealing, sigma=0.25 | 0.25 | 1.268 | 2.57 | 349 | 44 | 0.18 |
| annealing, sigma=0.50 | 0.50 | 1.130 | 3.31 | 426 | 67 | 0.92 |
| annealing, sigma=0.75 | 0.75 | 0.901 | 5.44 | 628 | 143 | 3.05 |
| annealing, sigma=1.00 | 1.00 | 0.864 | 5.91 | 675 | 159 | 3.52 |
| annealing, sigma=1.50 | 1.50 | 0.855 | 6.04 | 690 | 160 | 3.65 |
| annealing, sigma=2.00 | 2.00 | 0.854 | 6.04 | 691 | 160 | 3.65 |
| calm wind (0 m/s) | – | 0.792 | 7.03 | 812 | 185 | 4.64 |

Table 4: Two rating criteria, on the same weather, the same dispatch and the same limits, so the only difference between a pair of rows is the rule that turns a wind distribution into a published rating. Cap is the p5 rating — the 5 per cent exceedance of the ampacity distribution at the maximum operating temperature. Budget solves instead for the target whose realised temperature scatter spends the design-life allowance exactly: 5 per cent of virgin tensile strength over forty years. Binds names which arm set the budget rating. Cost gap is the budget basis minus the cap basis, so a positive figure means the budget rating costs the dispatch more. Above cap is the share of realised conditions that put the conductor over the maximum operating temperature at the published current — the quantity an exceedance policy would be written against, reported because this report does not apply one. The two criteria protect different things and the table does not say which is right.

| $\sigma$ (m/s) | Cap (× static) | Budget (× static) | Ratio | Binds | Cost gap (M$/yr) | Above cap (%) |
|---|---|---|---|---|---|---|
| 0.00 | 1.302 | 1.302 | 1.000 | clearance | 0.00 | 0.0 |
| 0.25 | 1.224 | 1.268 | 1.036 | budget | -0.23 | 27.3 |
| 0.50 | 1.134 | 1.130 | 0.996 | budget | 0.04 | 4.5 |
| 0.75 | 1.028 | 0.901 | 0.876 | budget | 1.11 | 2.3 |
| 1.00 | 0.905 | 0.864 | 0.954 | budget | 0.53 | 4.5 |
| 1.50 | 0.801 | 0.855 | 1.067 | budget | -0.82 | 6.8 |
| 2.00 | 0.797 | 0.854 | 1.071 | budget | -0.92 | 9.1 |

## 5 Discussion

The result is narrower than the usual case for dynamic rating, and the narrowness is the point. On this system the benefit is real but concentrated: in a few lines, in a small part of the year, and — the part that is easiest to miss — in a particular range of measurement quality. Most of what a better instrument buys is bought between 0.50 and 0.75 m/s. Below that the curve is flattening toward its floor and there is less left to win; above it the conductor is in natural convection and the wind estimate barely moves the rating. An instrument procurement is therefore a question about which side of that range the instrument lands on, and that question can be asked without reference to any fixed rating at all.

Two limitations bear on that directly. The weather is a single uniform scenario, so the report cannot say where an instrument should go, and the range in which precision pays would move with a different assumed climate. And the realized weather in each hour is drawn independently of every other hour, so the model cannot produce a still, hot spell lasting several days — which is the situation in which a conductor is actually damaged, and in which a rating set from a forecast is most likely to be wrong for long enough to matter.

Two further exclusions are larger than they look. 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 is worth roughly a factor of two in the loss of conductor life it implies. The conductor properties are catalog values; surface emissivity alone spans a factor of four in aging between a new conductor and a weathered one, and it drifts over the life of the line. Both are systematic where weather error is random: an incorrect emissivity is incorrect in the same direction in every hour, 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.

Finally, the contingency analysis covers single outages only, the dispatch is a linear direct-current approximation without losses or reactive power, and the emergency ratings are used without their duration limits — the shorter of the two is a fifteen-minute rating and this report uses only the longer.

## 6 Conclusion

A line rating is worth something only where a limit is reached. On the system examined that is 13.0 per cent of the hours with everything in service, on 2 lines, and 1.4 per cent of the hours with a branch out. Conditioning on those states is what makes the rest of the answer legible.

In those states, the answer is bracketed by two physical cases: 2.39 million dollars a year with the weather known exactly, and 7.03 million with nothing known about it at all, which forces the rating calm air permits. Everything an instrument can buy is the difference between them, a factor of 2.94, and it is not evenly distributed across that range: most of it is bought by sharpening a fair wind estimate and almost none by sharpening a poor one, because a poor estimate already puts the conductor in natural convection, where the wind hardly matters.

The practical reading is that a dynamic rating on this system is worth having, worth targeting at a named handful of lines, and worth only as much as the weather measurement behind it. It is not a system-wide improvement, and it is not robust to a poor measurement. Every number in this report is read from a results file the study writes; none is entered by hand, and the calculation is deterministic and carries no seed. Commit `95a98178418f`, Python 3.13.1, NumPy 2.3.5.

## 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.