RIGHTTACK CONSULTING · MANUFACTURING FOOTPRINT

The cost nobody quoted is the one that picks the plant

Twelve factors decide where a manufacturing plant goes; a company sites it on the two or three that are easiest to quote and discovers the rest at full price. How to rank locations by what actually differs between them — and why the weighted scoring matrix everyone starts with belongs last, if at all.

A company decides to build a plant, draws up a shortlist, and picks the winner on the numbers that were easiest to get: the incentive package the state waved, the cheap land, the low regional construction index. It looks like a rigorous decision, and it is frequently the wrong one — because eighteen months later the plant is up and the cost that actually decides whether it works was never on the comparison sheet. The labour pool is thinner than the headline wage suggested. The permit took a year longer than the schedule assumed. The utility interconnection is queued past the ramp date. None of those were quoted, so none of them counted — and now all of them cost.

Choosing where to put a plant is one of the highest-stakes, least-reversible decisions a manufacturer makes. You will operate the answer for twenty years. And it is routinely made with a tool that is actively built to get it wrong.

The tool that gets it wrong is the one everyone starts with

That tool is the weighted scoring matrix. Twelve factors down the side, a score of one to five for each site, a weight against each factor, multiply and sum, highest total wins. It feels rigorous. It produces a confident number. And it is almost always wrong, for two reasons that compound. The first is that it starts at the end — the weighted matrix is the last step of the analysis, not the first, because every weight in it should be an output of work you have not done yet. The second is that it forces incomparable things into one column. A delivered electricity rate is dollars per kilowatt-hour. “Quality of life” is a judgement. Score the second one a 3.5 and you have not measured it; you have disguised a guess as a number and made it harder to argue with, because arguing now looks like arguing with arithmetic.

> **A soft factor scored as a 3.5 is not data. It is a judgement wearing a number.**

Rank by what changes, not by what is biggest

The governing principle is almost embarrassingly simple, and skipping it is the most common failure in site selection. A cost that does not change when you move the pin on the map cannot decide where the pin goes. So the factor that matters is not the largest one — it is the one that varies most between the specific candidates you are actually comparing. Two of your three shortlisted sites may have near-identical construction costs and wildly different labour pools; then construction, however large a number, is irrelevant to the choice, and labour decides everything. And labour usually does decide: a difference of a few dollars an hour in the fully-loaded rate, multiplied across a couple of hundred heads and twenty years, quietly dwarfs the incentive package that everyone in the room spent their energy fighting over. Find the factors that differ. Ignore the ones that don't, however big.

Hard costs go in the model. Soft costs are a gate.

The twelve factors split cleanly, and the split tells you how to use each one.

> TWELVE FACTORS, TWO JOBS
>
> **Eight hard — put them in the cost model, in money.**
> Logistics & market access · tax · incentives · tariffs & duty · regulatory & permitting · real estate · utilities & energy · labour, unions & benefits.
>
> **Four soft — use them as pass/fail screens on the shortlist.**
> Industry cluster & talent base · cost of living · quality of life · transparency & country risk.

The soft four are decisive precisely because they rarely appear as a line item — they determine whether a plant can be staffed, run and held. The mistake is to score them alongside the hard eight and blend everything into one total. Don't. Ask a yes/no question of each: does this site clear a talent threshold, a safety threshold, a rule-of-law threshold? A soft factor a site fails should eliminate it, not cost it a fraction of a point it can win back on cheaper land.

None of it ranks a plant you have not specified

All of this assumes you know what you are siting, and that is its own artifact, produced first. Not a headline volume but a band — a low case and a high case across the planning

horizon, because if both cases site the plant in the same place the decision is robust, and if they don't you have found the real risk early. The unit envelope and mass, which decide whether an outbound unit is a legal load or permitted heavy haul — a three- to fourfold difference in freight. The headcount and skill mix, the input that most often decides the answer and is least often quoted before it. The utility load, the hazard class, and a lane schedule with a row for every inbound and outbound movement. Rank sites for a plant you have not specified and you are ranking a guess dressed as a decision.

Coarse to fine, and in that order

The sequence is the discipline. You rank a region before you rank a site, because the factors that choose a region — logistics, the trade regime, the broad labour and cost gradient — are not the ones that choose a site inside it. So spend precision coarse to fine: use the gradient factors to pick a region, use the pass/fail screens to cut that region to a shortlist, and only then do the site-specific landed-cost-and-risk work that separates the finalists. Each step narrows the search; none re-opens the one before it. Start at the fine end — a detailed twelve-factor matrix across three sites you have not yet earned the right to be comparing — and you have spent all your rigour in the wrong place.

The two traps that survive a good model

Two failure modes get past even a disciplined analysis, and each deserves a named check. The first is the incentive that flips the ranking. Incentives are real money, but they are discountable, clawback-exposed, and offered hardest by the sites that are otherwise weakest. So model them net of clawback and discounted — and then run the comparison once more with every incentive set to zero. If the ranking holds, the incentive is a bonus. If it flips, the sites were essentially equal on fundamentals and you have just let a grant that may not survive a change of governor pick a twenty-year asset.

The second trap is time. Permitting and utility interconnection are the line items most often costed and least often scheduled. A site that is cheaper on paper but ready eighteen months later than its rival is, in landed cost across the life of the asset, usually the more expensive site — and that slippage never appears in the matrix that chose it.

> **Re-run the comparison with every incentive at zero. If the winner changes, an incentive just picked your plant — and incentives do not have to keep their promises.**

The assumptions age faster than the building

One last discipline, because a plant outlives the analysis that sited it. Every location decision rests on assumptions that will move — automation strips out the labour hours that

decided it, an energy transition reprices the power that didn't, trade policy rewrites the duty that was treated as a constant, the customer geography that set the region shifts. Write those down now, while the logic is fresh and nobody is yet defending a decision, as explicit re-open triggers, each with a named owner whose job is to watch it. A trigger nobody owns is not a control; it is a paragraph that reappears, unread, in the post-mortem.

None of this is exotic. It is the discipline of looking at all twelve factors so none of them ambushes you later; ranking them by how much they actually differ between your real candidates rather than by how large they are; keeping the quotable costs in a model and the unquotable ones as gates; and moving coarse to fine so the expensive analysis lands on the decision that is still open. Do that and the answer is defensible line by line — which matters, because a plant location is exactly the kind of decision a board revisits and a downturn stress-tests.

And notice it is not a spreadsheet you fill in once. It is an analysis run in the right order, on numbers most of which are estimates at the point of decision, while a shortlist of sites and the people offering incentives on them are all pulling for a particular answer. Knowing the twelve factors is the easy half. Working out which three actually decide it, and defending that against a matrix that says otherwise, is the work.

That is the work I do.

WHERE THIS GOES NEXT

Where you build it — and how you can defend it →

The manufacturing strategy and footprint method — the eight-stage sequence this location decision sits inside. Or start a conversation: rtackconsulting.com/contact