ccPlanning

ccplanning.net

WORKFORCE PLANNING MASTERCLASS

The Accuracy Trap

Why chasing forecast accuracy is the wrong goal — and what to measure instead

John Casey

Workforce Planning Expert · ccPlanning

A masterclass in the ccPlanning white-paper series · read the rest at ccplanning.net

Contents

1. Accuracy is a means, not an end

Ask a planning team what they are trying to do and “improve forecast accuracy” comes back almost reflexively. It feels unarguable. It is also the wrong target.

A more accurate forecast feels like a better forecast, so the instinct is to chase the number down and assume the operation improves with it. But the real goal is a plan that delivers the right service at the right cost — and beyond a certain point, a more accurate forecast barely moves that, while a great deal of effort goes into shaving decimal places off a metric that was never the thing that mattered.

This is not a licence for sloppy forecasting. It is an argument for knowing what accuracy is **for**: measuring the part of it that actually drives decisions, and spending the saved energy on the thing that determines whether a plan succeeds — how well it survives the error that will always remain.

> *The goal was never an accurate number. It was a good decision. Accuracy only matters up to the point where it stops changing one.*

2. Error has a floor you already know

No forecast of human behaviour is exact. Even a perfect model of demand runs into the fact that arrivals are random: the same Tuesday, lived a hundred times, produces a hundred different contact counts. That irreducible scatter — the Poisson floor — sets a hard limit on how accurate any forecast can be.

Chasing accuracy past that floor is chasing noise, and noise does not respond to better models, more history or longer planning meetings. The first discipline is recognising where the floor sits for each queue, so you stop spending effort below it. Once you accept that some error is permanent, the question changes — from “how do I eliminate the error?” to “given that error remains, is my plan robust to it?” That is a far more useful place to put your attention.

3. Not all error costs the same

The standard accuracy metric is MAPE — the average percentage you were off, period by period. It is easy to quote and quietly misleading. It treats every miss as equal: a 10% error on a dead Sunday counts the same as a 10% error on your busiest Monday, even though only one of them leaves you short when it matters. And it says nothing about **direction**.

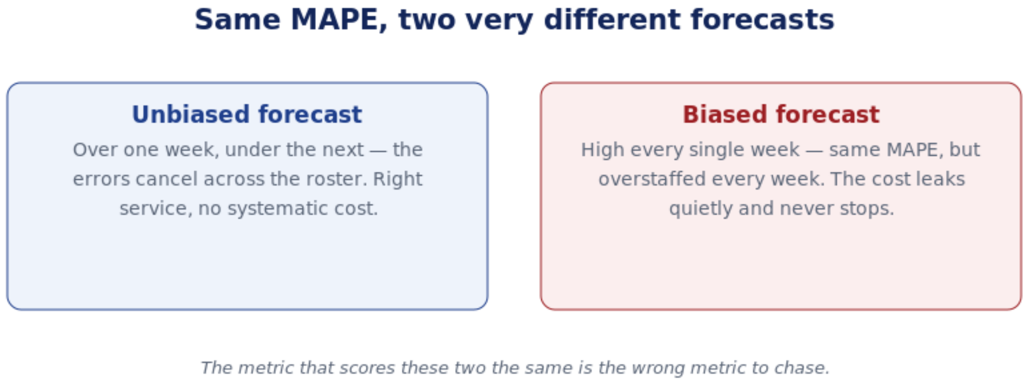

Two forecasts, identical MAPE. The biased one overstaffs every week; the unbiased one self-corrects. Accuracy alone cannot tell them apart — bias can.

A forecast that runs 6% high every week and one that is 6% high one week and 6% low the next have the same MAPE — but they are not the same forecast. The first quietly overstaffs every week; the second roughly cancels out. That directional error — **bias** — is the one that compounds, because it mis-sizes the whole operation in one direction, week after week, while noise averages out across a roster.

So if you measure one thing about your forecast, measure bias. If you measure two, add **WAPE** — the volume-weighted error — because it counts your busy periods for what they are worth instead of letting quiet days flatter the average. MAPE makes a comforting headline; bias and WAPE tell you whether the forecast is actually fit to drive staffing.

> *Bias is the error that compounds. Chase it before you chase the last point of MAPE — a lean costs you every single week.*

4. Which input is wrong matters more than how wrong the forecast is

Even a flawless contact forecast is only half the staffing equation. The number of people you need is driven just as hard by average handle time, by shrinkage, and by the service target you plan to — and an error in any of those moves the roster too. Pour all your effort into the contact forecast and you can still be badly staffed because handle time ran heavy or shrinkage beat the policy figure you used.

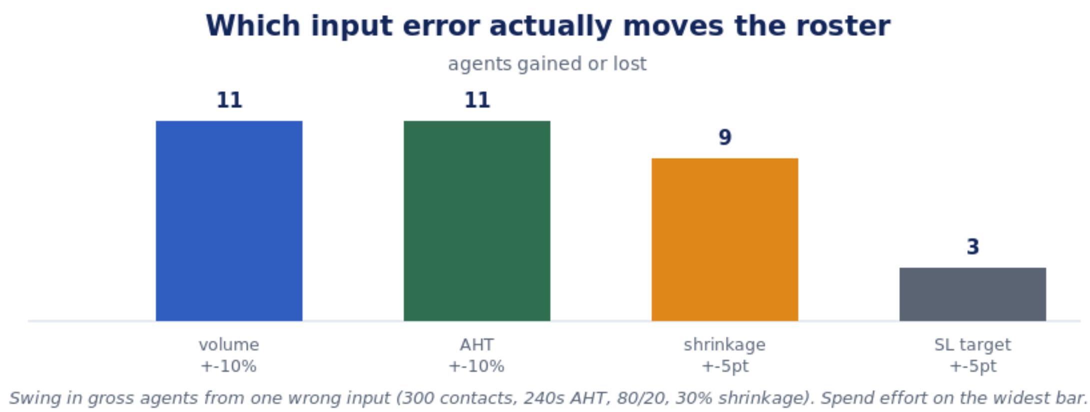

A single wrong input moves the roster by very different amounts. Volume and AHT lead; shrinkage is the quiet third; the service target trails.

Volume and AHT usually carry roughly equal leverage, because both feed straight into the workload. Shrinkage is often the quiet third lever — and the one most likely to be set from a stale assumption rather than measured. The service-level target, the thing people argue about most, typically moves staffing the least. The right question is not “how accurate is the contact forecast?” but “which input is my plan most exposed to, and have I buffered against that one?”

Knowing that order — for **your** operation, not in the abstract — is worth more than another point of forecast accuracy, because it tells you where to put your contingency and, just as usefully, where to stop worrying.

5. A useful forecast is one you can act on

The reframe is simple to state and hard to live by: the job of a forecast is to produce good decisions, not accurate numbers. A forecast that is 3% more accurate but arrives too late, in a form no one can roster to, or with no sense of its own uncertainty, is worse than a slightly looser one that is timely, decision-ready and honest about its range.

Three things make a forecast actionable, and none of them is accuracy. It has to be **timely** — available while there is still time to hire, train or move shifts. It has to be at the **right grain** — an annual total is useless for a roster, which lives on the intraday profile, while an interval forecast is the wrong tool for a budget. And it has to carry its **uncertainty** — a range, not a single number, so the people downstream plan for the spread instead of being ambushed by it.

On Monday, this means replacing “improve forecast accuracy” on the scorecard with two sharper targets: drive bias towards zero, and keep WAPE inside a band you have agreed is good enough for the decisions it feeds. Report bias first, run a sensitivity check so you know your real exposure, forecast material numbers as ranges, and when accuracy is already past its useful floor, have the confidence to stop polishing it and put the time into the plan. None of this lowers the bar — it moves it to where it belongs.

> *The best planners do not have the lowest MAPE. They have plans that hold up when the forecast is wrong — because they always assumed it would be.*

About the author

John Casey has spent more than 30 years in contact-centre workforce planning, including roles as a workforce-planning manager and operations director, and now writes and teaches at ccplanning.net. The views in this paper are his own. It represents general professional experience and is not affiliated with, nor does it represent, any current or former employer.

ccplanning.net · Free guides, calculators and the rest of the white-paper series.