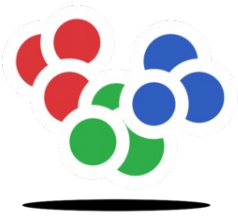

ccPlanning

ccplanning.net

WORKFORCE PLANNING MASTERCLASS

Shrinkage, Honestly

The building block most corrupted by wishful thinking — how to measure it from reality, forecast it, and stop it quietly wrecking your staffing

John Casey

Workforce Planning Expert · ccPlanning

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

Contents

1. The most corrupted number in the plan

Of the three building blocks of a staffing plan — volume, AHT and shrinkage — shrinkage is the one most likely to be wrong, the one most often wrong in the same hopeful direction, and the one whose error does the most damage. It deserves a paper of its own.

Volume gets the modelling attention. AHT gets some. Shrinkage gets an inherited percentage, typed into a cell, carried flat across the year, and quietly chosen to be the figure the operation wishes it had rather than the one it actually runs. That is the corruption: not fraud, just optimism — the target sickness rate instead of the real one, last year's number instead of this year's, an annual average instead of a seasonal profile. And because shrinkage is the last operation in the staffing calculation, every point of that optimism flows straight through to the number of people you ask for.

This masterclass is about doing it honestly: what shrinkage really is, why a small error in it scales the whole plan, how to split it into the two very different things it contains, how to measure it from reality and forecast it by period, and how to reduce the part you genuinely control without pretending away the part you don't.

> *Shrinkage is rarely wrong by accident. It is wrong in the hopeful direction — the figure you wish you had, not the one you run.*

2. What shrinkage actually is

Shrinkage is the share of paid time that is not available to handle contacts — everything that takes an agent away from the queue. Holidays, training, team meetings, coaching, one-to-ones, breaks, sickness, no-shows, system downtime, the lot. If 30% of paid time is lost this way, each agent delivers only 70% of their contracted hours to contacts, and your staffing requirement has to be grossed up to cover the difference.

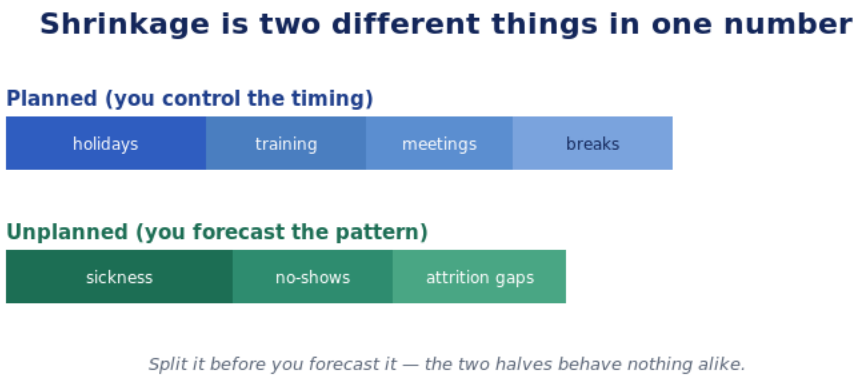

Split shrinkage into planned and unplanned before you forecast it — the two halves behave nothing alike.

Two definitions matter. First, be clear what is in your number: shrinkage definitions vary wildly between operations, and comparing your 28% to someone else's 22% is meaningless unless you both count the same things. Second, distinguish shrinkage from occupancy and utilisation — occupancy is about busy-versus-waiting while available; utilisation usually folds shrinkage in. Shrinkage specifically is the gap between contracted time and available-to-contact time, and it is the gross-up that turns a productive-hours requirement into a contracted-FTE requirement.

3. Why it is where plans die

Here is why shrinkage punishes error more than volume or AHT do: it sits at the very end of the sum. You work out the workload from volume and AHT, add a service buffer, and then — last of all — gross the result up for shrinkage. There is nothing downstream to absorb a mistake, so an error does not average out; it scales the whole requirement.

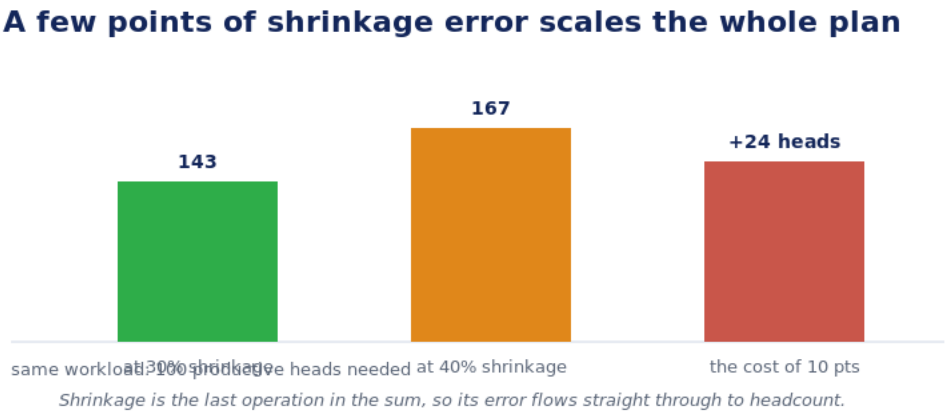

The same workload needing 100 productive heads requires 143 at 30% shrinkage and 167 at 40% — a 24-head swing from one assumption.

Work it through. A base workload that needs 100 productive heads, grossed up at 30% shrinkage, requires about 143 rostered staff. Assume shrinkage is really 40% — a leave-heavy or flu-hit period — and the same workload needs about 167. That is 24 extra heads conjured by a ten-point assumption, with a perfect volume forecast sitting alongside it looking immaculate. Apply your flat annual figure to a high-shrinkage week and you are not a little short; you are double-figures short, every day, and the staffing model gives you no warning because the error is in the one input nobody is watching.

> *Shrinkage is the last thing the requirement touches, so it is the last place you want a lazy number.*

4. Planned versus unplanned

The single most useful move in shrinkage planning is to stop treating it as one number and split it into two, because planned and unplanned shrinkage are different problems with different levers.

Planned shrinkage — holidays, training, team meetings, coaching, project time — is largely yours to control. You decide when it happens, which means you can steer it away from your busiest weeks rather than letting it fall randomly across the year. Planned shrinkage badly placed is self-inflicted: schedule a wave of training into your peak and you have manufactured a shortage. Planned shrinkage well placed is nearly invisible.

Unplanned shrinkage — sickness, no-shows, sudden attrition gaps — you cannot schedule, but you can absolutely forecast, because it has its own seasonality: winter sickness, summer leave pressure, the Monday-morning absence bump, the post-bonus attrition spike. Treating it as an unpredictable surprise is a choice, not a fact. It is noisy, but it is patterned, and patterns can be forecast.

Keep the two on separate lines in the plan. Planned shrinkage you manage by placement; unplanned shrinkage you manage by forecast and buffer. Blend them into one inherited percentage and you lose the ability to do either.

5. Measure it from reality, not the target

The corruption this paper is named for lives here. It is tempting — and common — to plan on the shrinkage you wish you had: the target sickness rate rather than the actual one, the leave allowance rather than the leave taken, the training plan rather than the training that actually consumed time. Optimistic shrinkage looks efficient on the spreadsheet and under-staffs you in the room, every single week it is wrong.

The discipline is to measure from reality. Pull the actuals: how much time really went to each category last year, by week, not the policy figure for it. Categorise honestly, including the awkward lines — the over-runs, the unbooked coaching, the “quick” team huddles that were not quick. Reconcile the total against the simplest possible check: contracted hours minus hours actually on contacts. If your built-up shrinkage and that top-down gap disagree, something is missing from your categories, and what is missing is almost always real lost time you would rather not count.

> *Plan on the shrinkage you have, not the shrinkage you wish you had. The wish costs you a shortfall every week it survives.*

6. Forecast it by period

Once you are measuring honestly, forecast it the way you forecast anything else — by building it up from its components and projecting each one forward, to produce a figure per period rather than one number for the year.

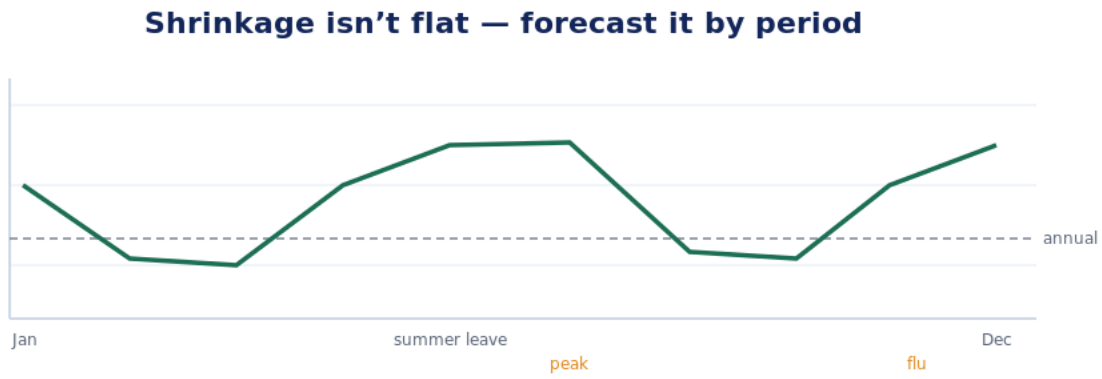

Project planned shrinkage from the leave and training plans; project unplanned from seasonal history. A single annual figure hides the weeks you’ll be short.

Planned shrinkage you can project almost directly from the plans themselves: the leave calendar, the training schedule, the meeting rhythm. Unplanned shrinkage you project from seasonal history — sickness curves, absence patterns, attrition timing — the same trend-and-seasonality thinking you apply to volume. Add them and you get a shrinkage profile that is higher in leave-heavy and flu season and lower in the quiet months, which is the only version that staffs the year correctly. A flat annual figure carried across twelve months systematically over-staffs the calm periods and under-staffs the hard ones — the worst of both.

7. Reducing the part you control

Forecasting shrinkage honestly tells you the truth; it does not by itself reduce it. Some of it is immovable and should be — breaks and leave are not waste, and training is an investment. But a meaningful slice is addressable, and the plan should separate “lost time we accept” from “lost time we can recover.”

On the planned side, the lever is placement and discipline: steer training and projects into the troughs, hold the meeting load down in peak weeks, and stop planned shrinkage from quietly inflating through scope creep. On the unplanned side, the lever is the operation’s health: sickness and attrition shrinkage are downstream of how the place is run — chronic high occupancy, poor scheduling and weak management drive both, so the cheapest shrinkage reduction is often a better-run floor. But be honest about the limit: you will not target your way to a sickness rate the operation has never achieved, and planning on the reduction before you have delivered it just re-imports the optimism you worked to remove.

8. An operating discipline

Shrinkage done honestly is a habit, not a one-off. A short operating discipline:

- **Define what is in it** — a written list of categories, so your number is comparable to itself over time.
- **Split planned from unplanned** — two lines, two different levers, never one blended percentage.
- **Measure from actuals**, not policy targets — reconcile bottom-up categories against the contracted-minus-contact gap.
- **Forecast by period** — a weekly or monthly profile, planned projected from the plans, unplanned from seasonal history.
- **Place planned shrinkage deliberately** — into the troughs, away from the peaks.
- **Buffer for the unplanned** — and lean on standby or overtime so an ordinary sick day is not a service incident.
- **Plan the rate you have**, not the one you wish for — bank reductions only once delivered.
- **Re-baseline regularly** against the latest actuals, because the profile drifts.

Do those and shrinkage stops being the quiet hole in the bottom of the plan and becomes what it should be: a forecast input, made honestly, that the rest of the staffing calculation can trust.

Appendix. Shrinkage, in formulas

The shrinkage gross-up, in one place:

- **Shrinkage % = lost paid hours ÷ total paid hours** (across all categories, planned and unplanned).
- **Productive hours per agent = contracted hours × (1 − shrinkage).**
- **Required contracted FTE = required productive agent-hours ÷ productive hours per agent.**
- **Equivalently, required FTE = productive-heads requirement ÷ (1 − shrinkage).**

Worked: a 100 productive-head requirement at 30% shrinkage needs 100 ÷ 0.70 ≈ 143 rostered FTE; at 40% it needs 100 ÷ 0.60 ≈ 167. The ten-point difference is 24 heads — which is why shrinkage is worth forecasting as carefully as the volume it sits behind.

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.