CHAPTER 14 / 3 MINUTE BRIEF

Seasonal Variation

Read the full chapter → / Read this brief online →

The main idea

When someone asks for a better capacity forecast, I'd first want to understand what's wrong with the one they have. Perhaps the model handles seasonality poorly, or perhaps the organization keeps changing the conditions the model was supposed to describe. If the launch moved, the promotion expanded, or a new feature changed the kind of work each customer performs, those are different problems from failing to recognize a recurring pattern in the measurements.

I want a good model, and I want to know how well it performs, but I also want access to the people who can explain what the business intends to do. Otherwise we can spend a lot of time improving our estimate of what would have happened if the plan had stayed the same. The model may be doing an excellent job within its assumptions while the capacity team prepares for the wrong workload.

There is real mathematical work here, and we'll do some of it. A forecast earns trust when it is scored against what was known on the day it was made, in the unit of the decision it served, with the business's information kept beside the model.

Explore this section in the full chapter →

Start with a comparison we can understand

A **seasonal-naive forecast** gives us a useful starting point: it repeats the latest observation from the matching season, so next Monday's forecast is the most recently observed Monday. Hyndman and Athanasopoulos treat it as a standard benchmark in *Forecasting: Principles and Practice*, with a blunt rule: a new method that can't beat these simple alternatives isn't worth considering.[1]

Explore this section in the full chapter →

Recreate what you could have known

The evaluation has to recreate what we knew when each forecast would have been made: choose a date in the past, forecast using only the information available then, compare the result with what happened, then move the date forward and repeat. Each starting date is a **forecast origin**; the repeated comparison, which can test tomorrow's forecast or one made weeks ahead, is **rolling-origin evaluation**.[2]

Our companion starts after four weeks of history, advances the origin weekly, and compares both methods at one, seven, eight, fourteen, and forty-two days. Eight days corresponds to the constructed launch decision; forty-two days corresponds to the initial supply estimate. I'd want to see the latter before relying on the model for purchasing: predicting tomorrow well says little about a decision whose supply consequence arrives six weeks later. That horizon also has the fewest completed cases, seven, and their overlapping six-week windows make them far from independent.

Explore this section in the full chapter →

[1] Rob J. Hyndman and George Athanasopoulos, *Forecasting: Principles and Practice*, third edition, section 5.2, "Some simple forecasting methods," https://otexts.com/fpp3/simple-methods.html, accessed October 3, 2026. Defines the seasonal-naive method, including its general form for horizons beyond one season, and says a new method that does not beat these benchmarks is not worth considering. The four-week matching-weekday mean and numerical fixture are constructed for this book.

[2] Hyndman and Athanasopoulos, *Forecasting: Principles and Practice*, section 5.10, "Time series cross-validation," https://otexts.com/fpp3/tscv.html, accessed October 3, 2026. Describes evaluation on a rolling forecasting origin, including multi-step horizons. The companion's weekly origins and evidence-availability dates are this book's adaptation.

Put the conversation next to the model

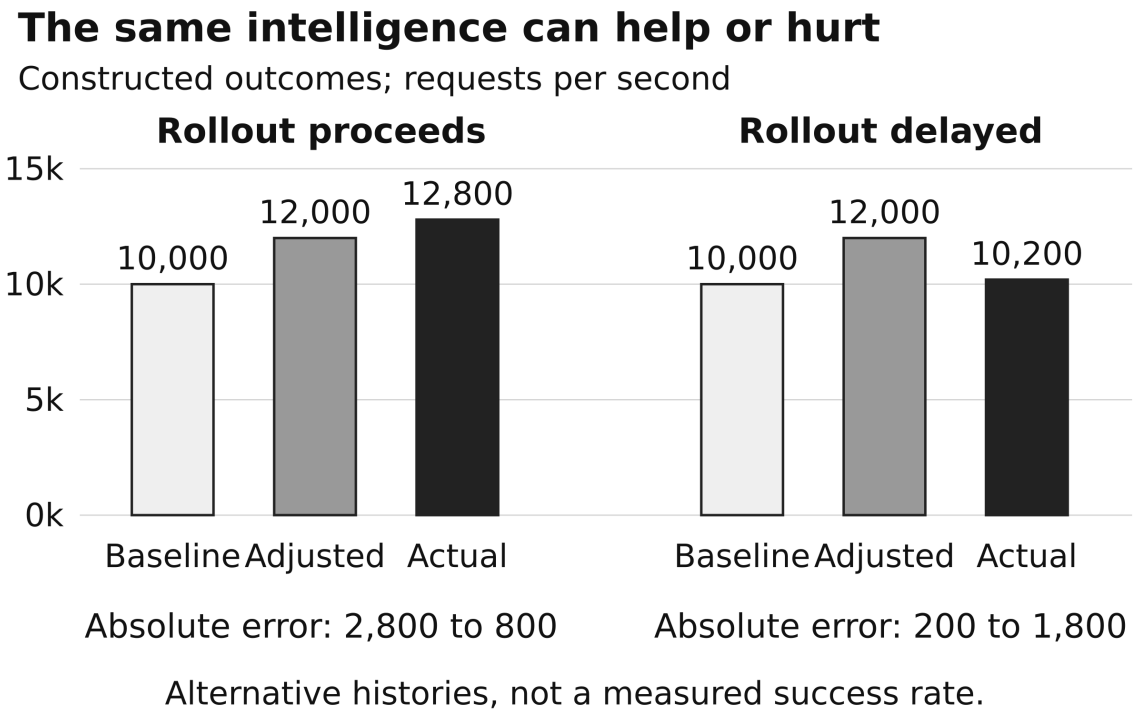


Diagram 14.1. The same event adjustment reduces error when the rollout proceeds and increases error when it is delayed. The bars show two alternative constructed outcomes.

Explore this section in the full chapter →

What the uncertainty changes

A **prediction interval** gives a range with a stated probability under a model and its assumptions. We can check whether a claimed 90 percent interval contains roughly 90 percent of later outcomes at the horizon we care about. A low-demand and a high-demand business scenario give us a different kind of range.[3] Both can help a decision, but our staged 12,000 and full-rollout 14,000 cases don't become the 50th and 95th percentiles merely because those labels would make the chart easier to present. We haven't established those probabilities.

Explore this section in the full chapter →

The discussion becomes useful when we connect those mechanisms to options. A delayed rollout may let us defer the purchase. Delayed supply may mean staging longer, if the business and batch owner can accept the consequences. A changed request mix requires someone qualified to reassess the service rate. Where we have defensible probability estimates, they can inform those choices. Where we don't, the alternatives and their conditions still deserve to be examined.

I sometimes call forecasting magic, partly because of how much judgment can disappear behind the final number. I'd rather let the reader see that judgment: why we chose the historical comparison, what we learned from the business, and where we still have to make an assumption before the decision deadline. That doesn't eliminate guessing. It gives us something specific to revisit when the world behaves differently.

Explore this section in the full chapter →

[3] Hyndman and Athanasopoulos, *Forecasting: Principles and Practice*, section 5.5, "Distributional forecasts and prediction intervals," https://otexts.com/fpp3/prediction-intervals.html, accessed October 3, 2026. Supports the distinction between point forecasts, forecast distributions, and prediction intervals; this book's rollout scenarios are not calibrated intervals.