From Better Forecasts to Better Decisions: Data Science in the CFO Function

Kashif Zafar
Independent researcher, Dubai, United Arab Emirates
11 October 2026

Abstract

This paper argues that data science inside the chief financial officer function should be organised around consequential decisions rather than model sophistication. It sets out a six-stage operating approach: defining a named decision, owner and objective before any modelling work; reconciling model inputs to source systems and the ledger; benchmarking candidate models against existing and naive forecasts on out-of-sample data; converting probabilistic output into pre-agreed decision rules; measuring realised value against a credible counterfactual including all costs; and treating deployed models as owned, controlled and reviewed operating processes. The approach is offered as general management practice and does not constitute investment, accounting or tax advice.

Keywords: data science; management accounting; forecasting; decision analysis; model validation; business value

1. Introduction

Every finance team is being told to 'do more with data.' And what happens? More dashboards. More elaborate forecasts. More models. And nobody decides anything differently as a result.

Here's how I judge data science in the CFO function: did it improve a consequential decision? That's it. That one question shifts your attention from model sophistication to what actually matters — reliable data, credible validation, honest uncertainty, and a clear link to someone actually doing something different.

2. Pick a decision worth improving

Start with an actual problem. A named decision-maker. The actions they can take. The objective they care about. Working capital targets, collections prioritisation, pricing approvals, inventory buffers — those are good candidates.

'Improve forecasting' is not a project. 'Help the treasury team decide weekly funding needs with fewer surprises' — that's a project.

3. Get the data foundation right

Reconcile your model inputs to the source systems and the ledger. Document the definitions, the cut-off timing, the known gaps. Because if the model's revenue doesn't tie to the reported number, the whole conversation will be about that gap. Not the insight.

This work is unglamorous. It's usually the biggest chunk of the effort. And it's the only thing that makes everything else credible.

4. Set a credible baseline

Compare your fancy new model with whatever forecast the business already uses. And with the dumb-simple methods — last year plus growth, a moving average, a seasonal naive model. You'd be surprised how many sophisticated models can't consistently beat these.

If your model can't outperform a sensible baseline on out-of-sample data, its complexity is a cost. Not a feature.

5. Turn predictions into actual choices

A single number hides the risk. Show ranges and scenarios. Agree in advance how they map to action — if the downside case crosses this threshold, we draw the facility. If the upside holds, we release buffer stock.

Explicit decision rules make the model useful. And they make it obvious when someone overrides it with judgement. That's fine — as long as it's visible.

6. Measure the actual business value

Compare outcomes against something credible: the previous process, a control group, or a clearly stated counterfactual. Count everything — implementation costs, running costs, people's time.

And report honestly. Including when a project doesn't pay back. That honesty is what builds trust for the next project. Hiding failures builds nothing except a reputation for confident garbage.

7. Make analytics an operating capability, not a side project

Models degrade. The business changes, the data shifts, and suddenly your model is confidently wrong. Assign ownership for monitoring, retraining, documentation, and fallback procedures. The finance close should never depend on one analyst remembering how a Jupyter notebook works.

Treat a production model like any other critical process. Owned. Controlled. Reviewed.

8. Conclusion

The goal is not a more impressive forecast. It's a finance function that decides faster, decides better, understands its own uncertainty, and can show the value of its analytics in the same language it uses for every other investment.

References

1. Breiman, L. (2001). Statistical modeling: The two cultures. Statistical Science, 16(3), 199–231.

2. Hyndman, R. J., and Koehler, A. B. (2006). Another look at measures of forecast accuracy. International Journal of Forecasting, 22(4), 679–688.
3. Kohavi, R., Longbotham, R., Sommerfield, D., and Henne, R. M. (2009). Controlled experiments on the web: survey and practical guide. Data Mining and Knowledge Discovery, 18(1), 140–181.
4. Molnar, C. (2020). Interpretable Machine Learning: A Guide for Making Black Box Models Explainable (2nd ed.). Lulu.com.