Attribution methods for AI benefits

How to separate what the AI did from everything else that changed in the same quarter, and what each method costs.

Terence Kok with Dr. Christina Chua

Revenue rose and the AI went live in the same quarter.

So did a routing redesign, a campaign and a new team lead. A before-and-after comparison credits all of it to the model.

Attribution asks one question: what would have happened over the same weeks without the AI?

12% before and after. About 5 points net.

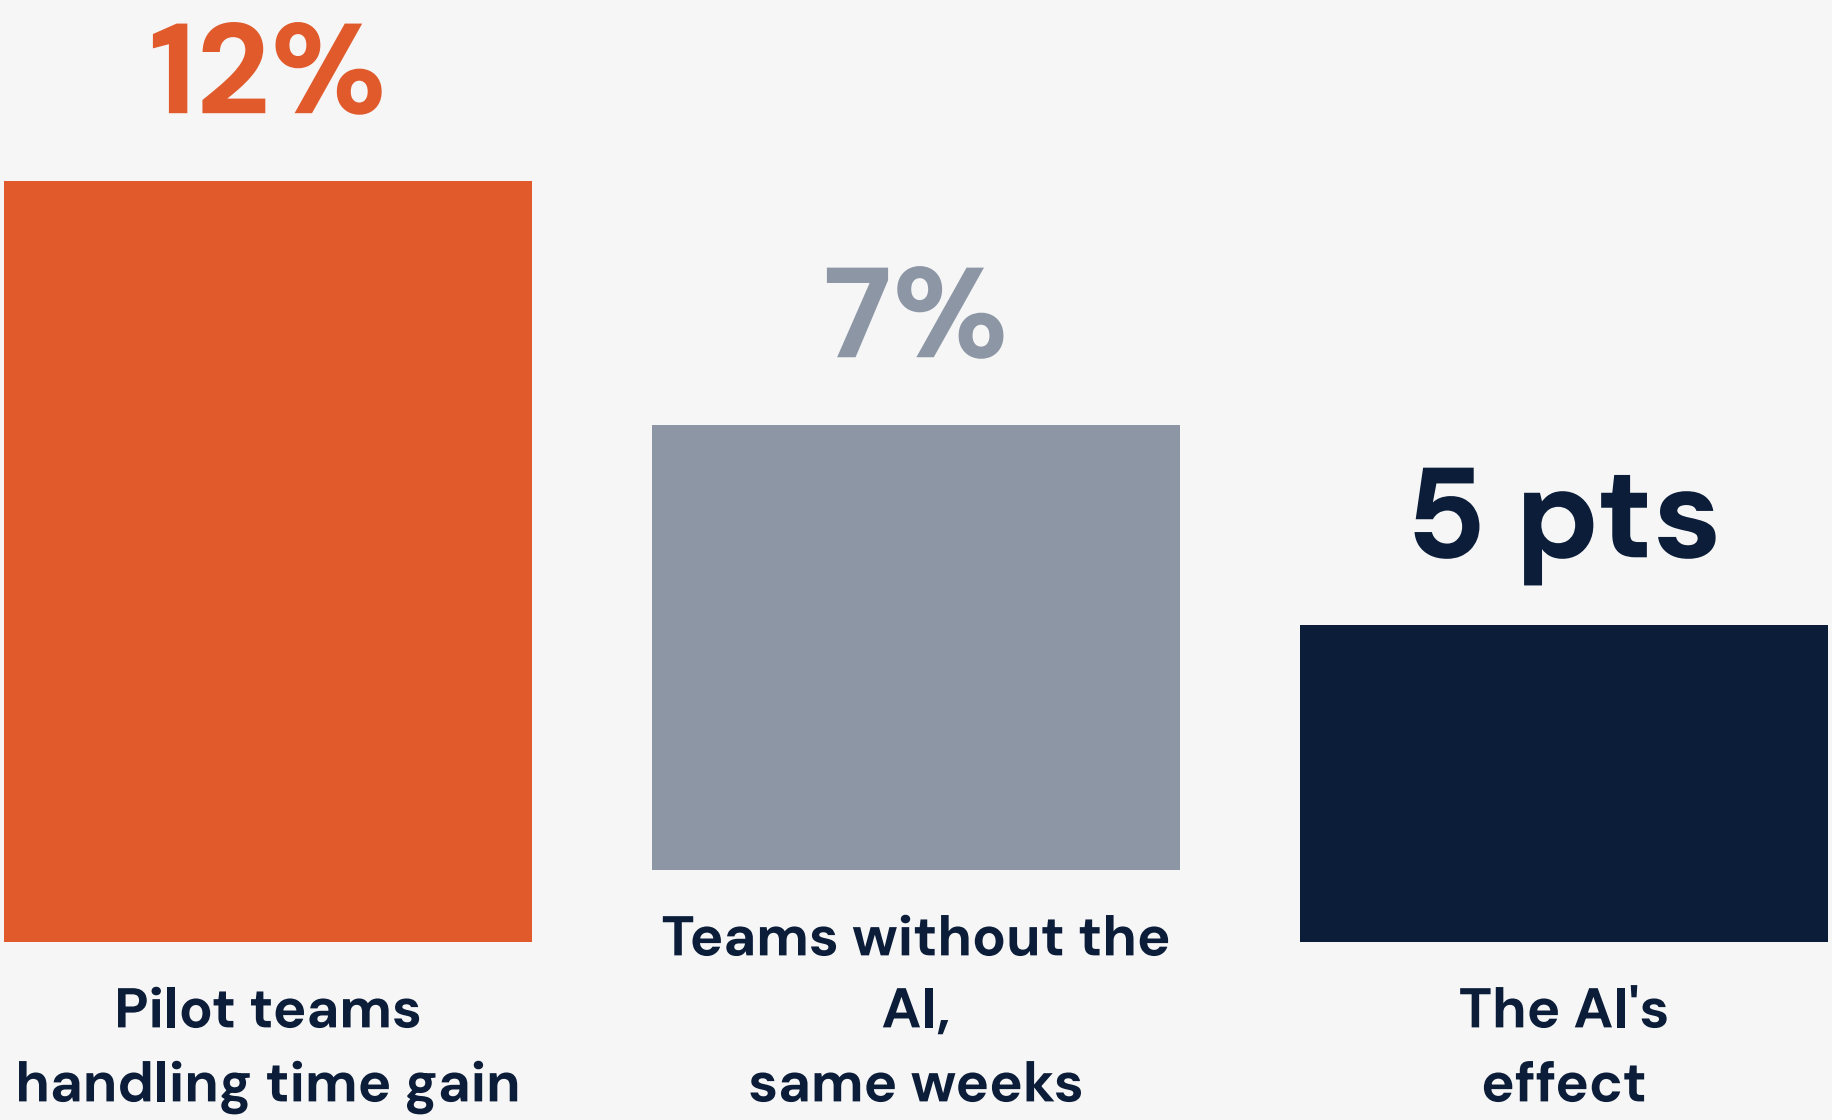

Illustrative contact-centre pilot. The other teams gained 7 points from a routing redesign.

26%

more completed tasks, with a standard error of 10.3 points.

Three randomised trials at Microsoft, Accenture and a Fortune 100 company, 4,867 developers. Each was noisy on its own. Cui et al., Management Science, 2026.

Large, well-run trials still give a wide range. A pilot with no comparison group gives no reliable estimate.

Compare each wave with the teams still waiting.

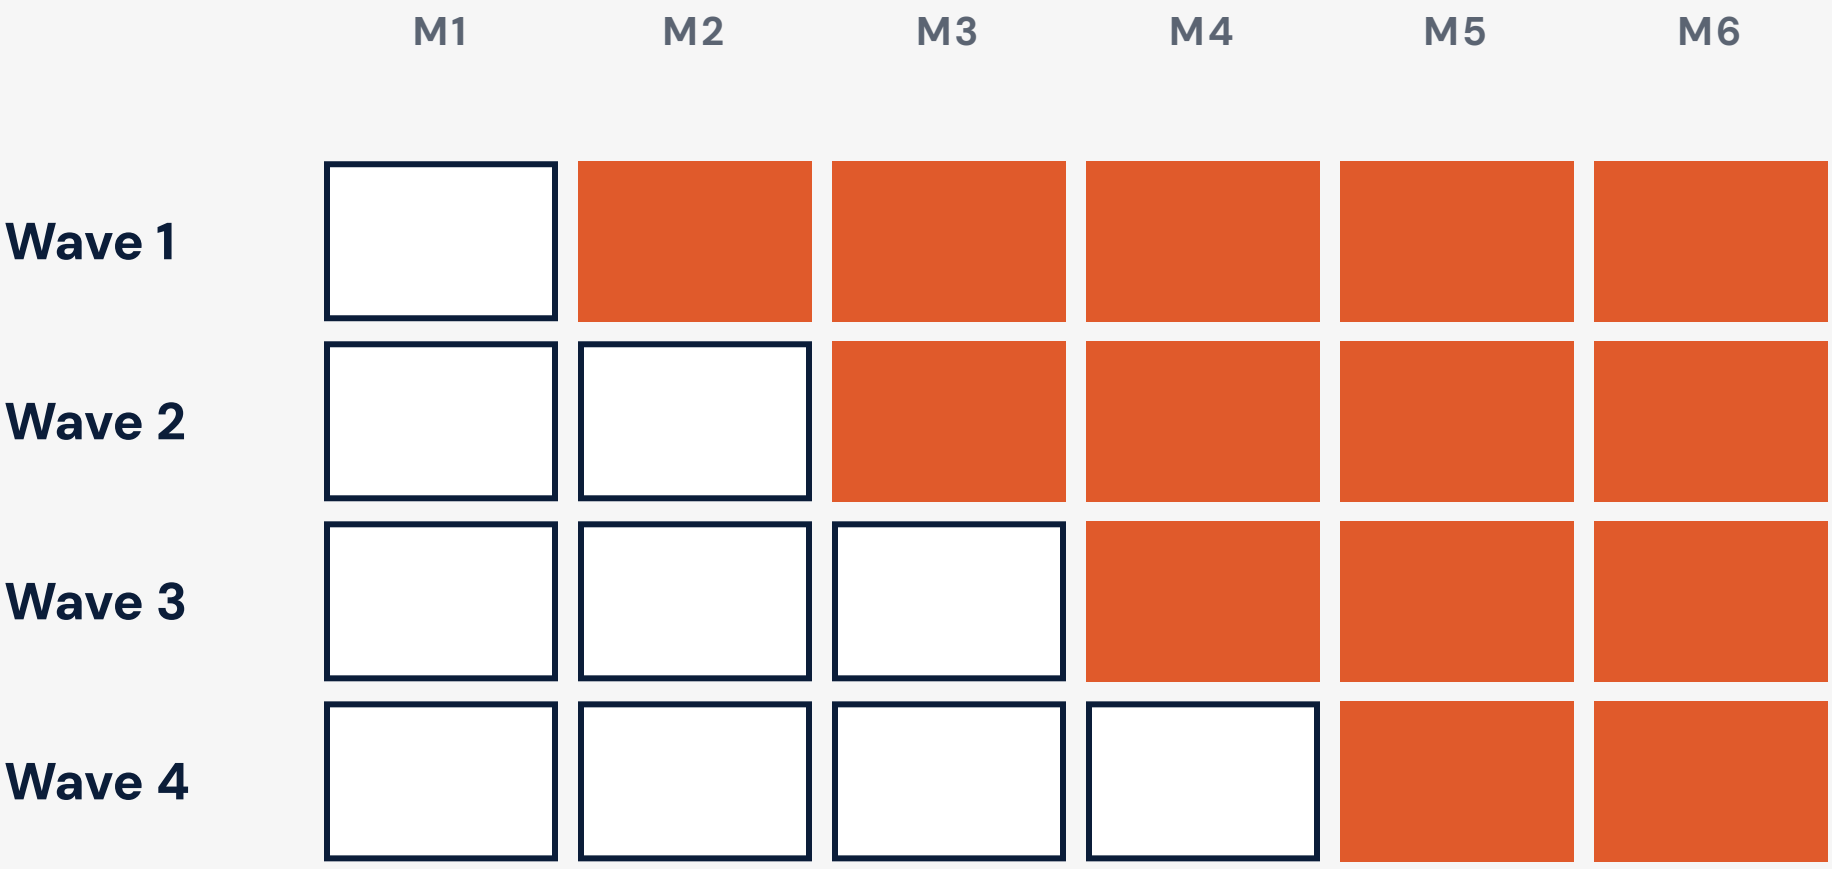

Draw the wave order by lot. Use an estimator that never treats early waves as the comparison for later ones.

80%

of workers given the tool used it.

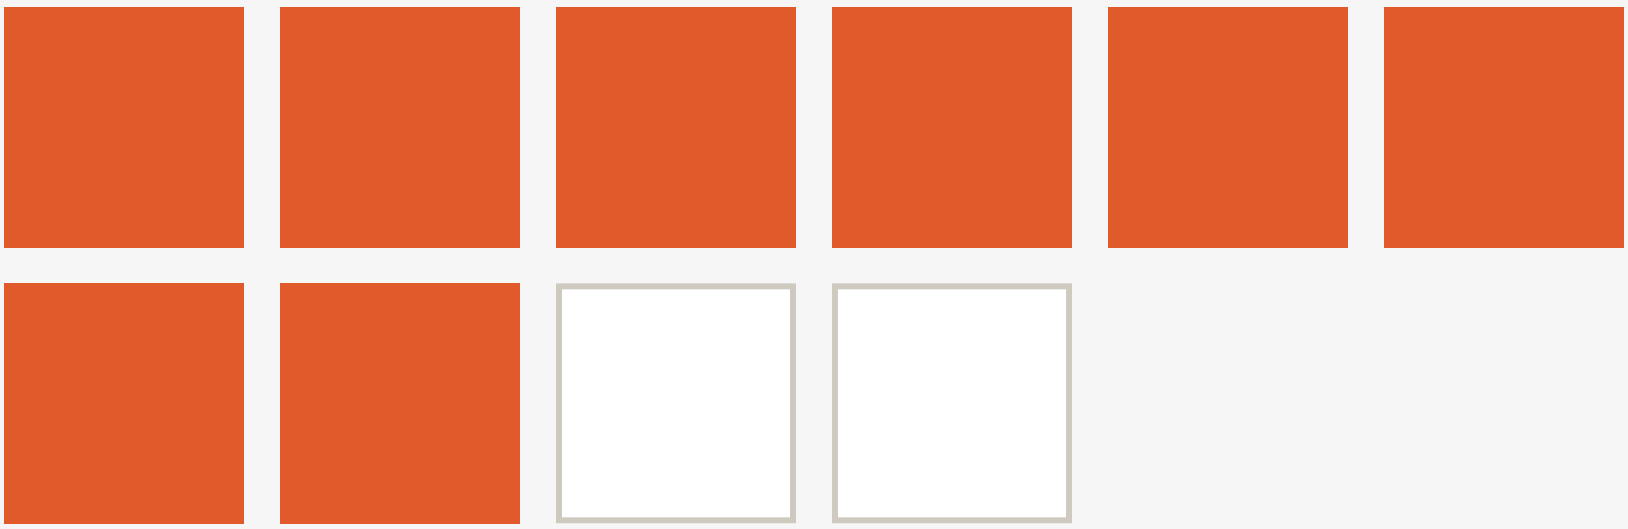

Users saved about two hours a week on email (Dillon et al., 2025, 66 firms, 7,137 workers). Across every licence that is about 1.6 hours, our arithmetic.

Six methods, strongest first

1 **Randomised holdout**
Cost: benefit forgone by the holdout

2 **Staggered rollout, order by lot**
Cost: later waves wait

3 **Switchback test**
Cost: weeks of on/off time blocks

4 **Synthetic control**
Cost: long unit-level history

5 **Matched comparison**
Cost: groups may differ unseen

6 **Trend-adjusted before and after**
Cost: anything launched the same week

When no comparison group exists

1 **Report it as capacity, not cash**

2 **Grade it low and keep it apart from measured benefits**

3 **Record that the benefit was not measurable**

Never book a heavy-user versus light-user gap as a benefit. It measures who chose to use the tool.

Five questions before accepting a benefit

1 **Which group, over which weeks, shows what would have happened?**

2 **Who chose the rollout order, and was it drawn by lot?**

3 **Is the figure per licence or per active user?**

4 **What else changed, and how was it removed?**

5 **Was the method written down before the results?**

Choose the comparison before the pilot starts.

One line in the pilot charter. It is the step that lets a CFO trust the result.

Terence Kok · AI & Innovation Strategist

SAVE · COMMENT · FOLLOW