

SCORECARD

Five questions on your AI model choices

These are the five questions for management from the paper "AI economics in fintech: reading card statements". For each one, tick whether your team could answer it today, and note the evidence you'd point to.

1 Which model does each of our AI use cases run on, and when did we last test that choice against the alternatives?

A good answer: a list of use cases with the model and version for each, and a test date in the last six months.

☐ Yes ☐ Partly ☐ No

Evidence we would point to

2 Do we measure accuracy in money, on a test set built from our own data with known answers?

A good answer: a set of real documents or cases with checked answers, and an error rate stated in pounds or euros.

☐ Yes ☐ Partly ☐ No

Evidence we would point to

3 How much of our own knowledge is written into our prompts, and who owns and maintains them?

A good answer: a named owner, prompts kept with version history, and a record of what changed and why.

☐ Yes ☐ Partly ☐ No

Evidence we would point to

4 Where does each model run, under what retention terms, and which companies in the chain, including any router, see our data?

A good answer: for each use case, the host, region, retention terms and every supplier that handles the data, reviewed by risk or compliance.

☐ Yes ☐ Partly ☐ No

Evidence we would point to

5 Which of our tasks need an advanced model, and what would a cheaper model that passes the same test save us at our volumes?

A good answer: monthly volumes for each task, the cost per task for the current model and for the cheapest one that passes, and the yearly difference.

☐ Yes ☐ Partly ☐ No

Evidence we would point to

If you ticked Partly or No on two or more, those questions are the place to start. Happy to talk them through. Let me know at martin@scalepointpartners.com.