

Purpose. Present the evidence from a test and secure a clear, owned scale, hold, or stop decision.

What it captures. Nine elements in a fixed order: the decision requested, the context, the evidence behind it, the assumptions not yet evidenced, the options weighed, the risks and safeguards, the recommendation, the decision (scale, hold, or stop), and the owner, next review date, and trigger.

When to use it. Bring the brief to any checkpoint where a scale, hold, or stop decision carries real cost and the people deciding need the evidence in front of them.

Common mistake. Dropping the “Do Nothing” option, which lets any proposal win by default. Project Swift’s brief stated the cost of the status quo so the recommendation competed against a real alternative.

Decision rule. Close every brief with one of three decisions: scale, hold, or stop, each carrying a named owner and a dated next review. A brief that ends without a dated decision and an owner is a discussion, not a decision.

Decision Requested

Context *including the cost of doing nothing*

Evidence *each claim traceable to an Experiment Card or discovery output*

Assumptions *anything you cannot trace to evidence*

Options Weighed *always include a Do Nothing option*

-
-
-
-

Risks and Safeguards *including sustainability*

-
-
-

Recommendation

Decision *choose one*

☐ Scale ☐ Hold ☐ Stop

Owner, Next Review Date, and Trigger *the metric or event that prompts the next review*