

The Fragility Lab

Decisions under uncertainty, four assumptions, one verdict

PERSONAL PROTOTYPE · WORKED EXAMPLE

The problem

A 4.8x recommendation looks robust on paper. Then reality intervenes, and the biggest number is usually the one whose model is most wrong in your favour.

Method

- Each of four parameters is estimated from imperfect evidence.
- The tool runs a hundred scenarios and maps the whole parameter landscape.
- The gap between the optimistic number and the realistic spread is visible before anything is committed.

What you can test

- Change the four assumptions one at a time.
- Find the one that flips the verdict from holds to breaks. It is usually the one nobody had looked at.
- Read the spread, not the headline figure.

How it was built

ROUND 1

Same brief to four models.
Two disagree.

ROUND 2

The disagreement goes
back with evidence.

ROUND 3

What survived is verified
against the source.

DECISION

Signed, with the trail of
rounds attached.

Evidence and limits. A worked global-health example with stylised parameters. It shows how a recommendation depends on its inputs; it is not an evaluation of any real programme.