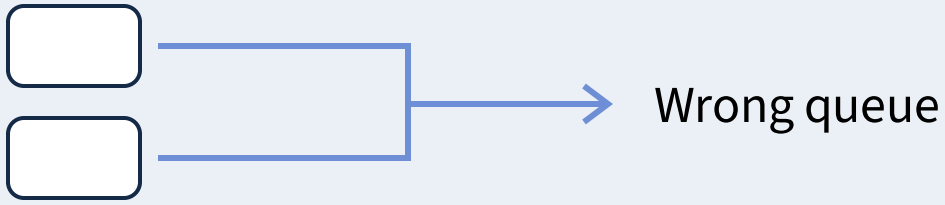
AI risk in support.
A field guide to customer outcomes.



Familiar harms.
Less predictable failures.

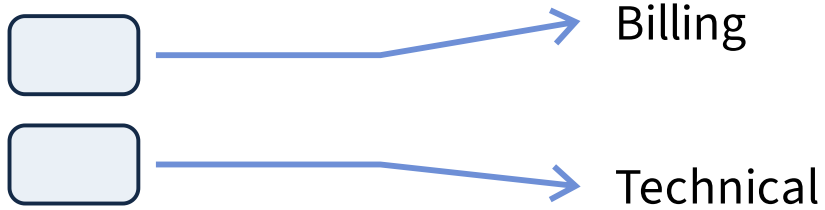
Fixed routing rule

Faulty conditions repeat errors.



AI classifier

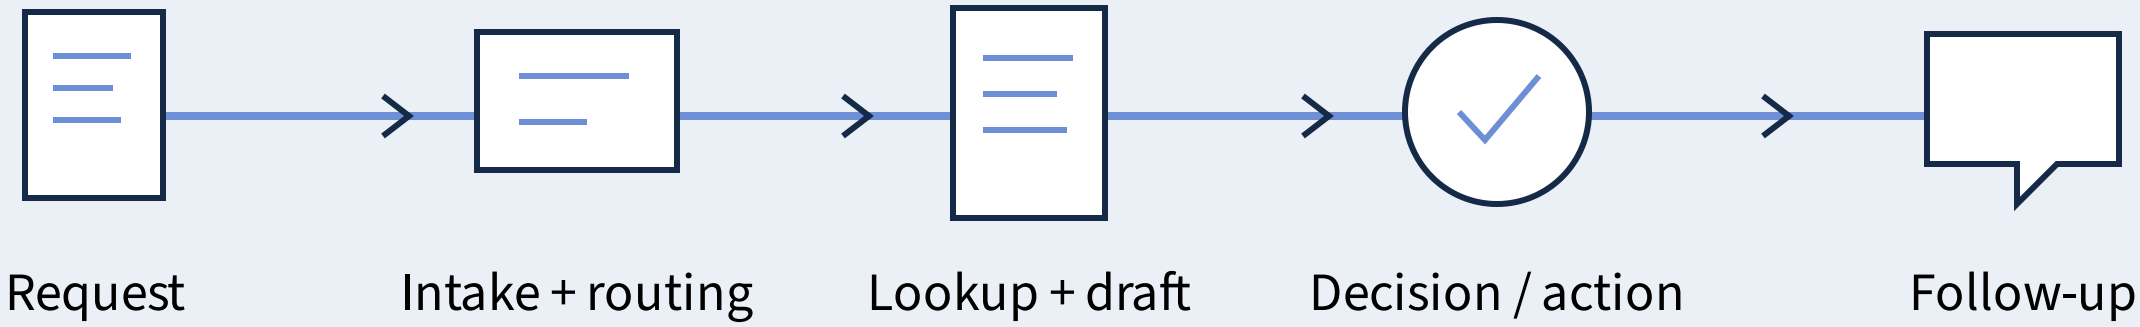
Unfamiliar wording shifts routing.



AI adds uncertain outputs, data dependence and human overreliance.

Both delay service. Reviewers can compound errors.

Follow the request.
Track the authority.



AI advises →

Acts after approval →

Acts autonomously

At each handoff: data? permissions? reviewer? human access?

Where a draft becomes a promise.

Illustrative refund case

Unchecked draft → sent promise

“You qualify for a full refund. It will arrive within five business days.”

Safer response

“I’ll send your request to billing for an eligibility review. Your refund is not yet approved.”

Company rule: Billing reviews eligibility before support promises approval.

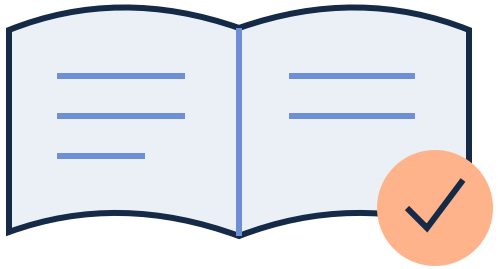
Turn a failure into a risk decision

Failure	Unapproved refund promised
Impact	Customer disputes; corrective work
Likelihood	Test realistic requests and exceptions
Exposure	Which channels and decisions?
Uncertainty	What remains untested?

Prioritize likelihood and impact. Unknown is not low risk.

Restrict commitments until controls are tested.

Safeguard the consequential handoffs



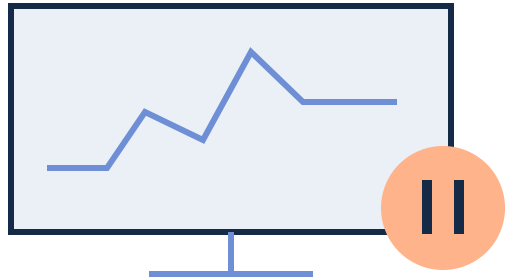
Before drafting

Current policy
Minimal customer data
Restricted access



Before sending / acting

Check facts and authority
Escalate eligibility
Block unauthorized actions



During operation

Test; monitor failures
Human handoff
Disable automation

Document the risk that remains.

Keep the workflow under review

GOVERN Owner, authority, escalation throughout

MAP

Task, data, customers and potential harms

MEASURE

Realistic tests, errors, feedback and overrides

MANAGE

Prioritize fixes, restrict uses, recover



Continuous functions, not a one-way checklist.

Start bounded.
Expand with evidence.

One workflow. One owner. Clear launch and pause criteria.

