Data Science Case Studies: Patterns and Tradeoffs

Course goal: evaluate decisions across the full project lifecycle, not just model accuracy

Patterns: recurring decision situations and their common solution shapes

Tradeoffs: cost vs. speed, interpretability vs. accuracy, exploration vs. exploitation, leakage risk vs. robustness

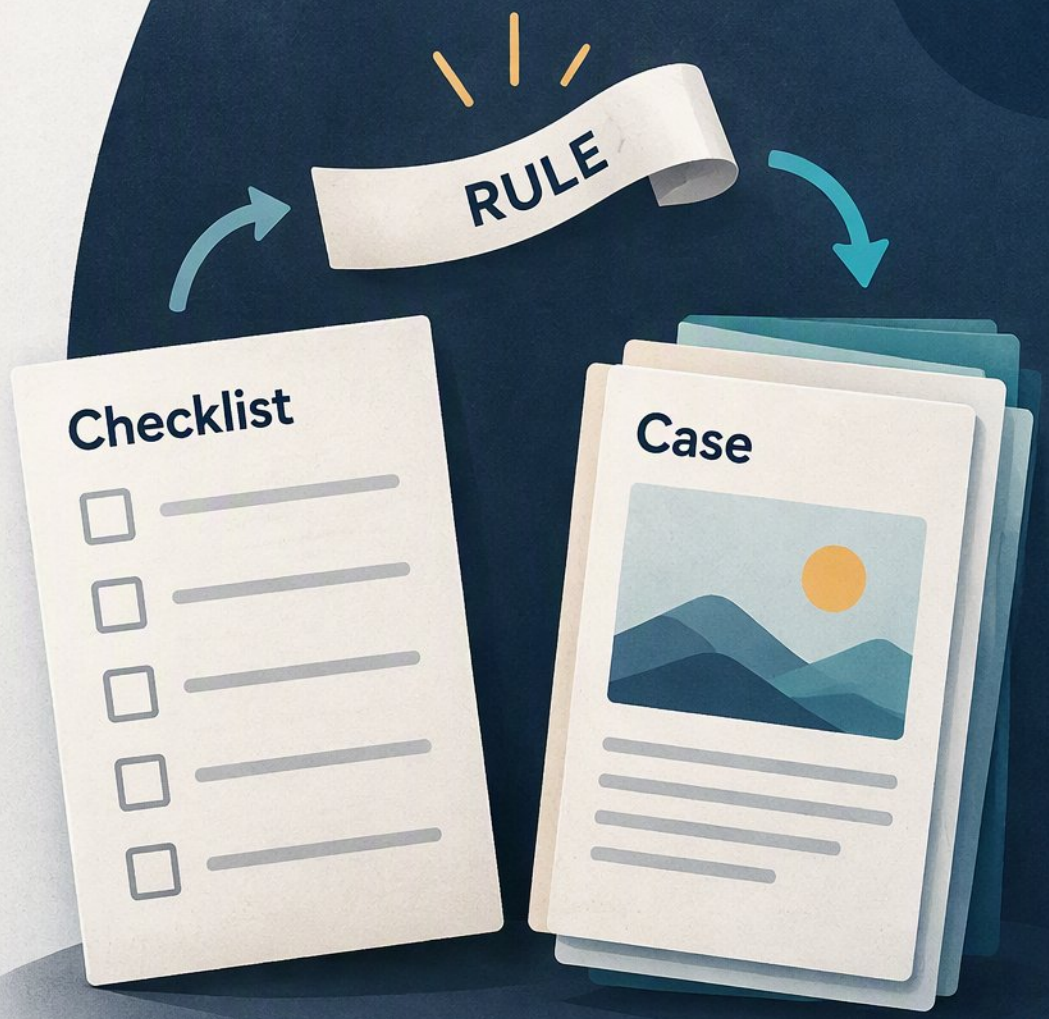
Structure: one case per lifecycle stage, each with a short decision review

Outcome: diagnose the central tradeoff, justify a defensible choice, communicate residual risk



Why Cases Beat Checklists

- Checklists encode what to do; cases teach when a rule should bend.
- Same technique, different context: aggressive imputation helps surveys, harms time series.
- Reusable case anatomy: context, constraint, decision point, options, evidence, outcome.
- Anti-pattern: reading cases as success stories, not decision records.
- Preserve null results and failures; extract patterns without overgeneralizing.



The Lifecycle Decision Map

SAME QUESTIONS
A STRONGER TOMORROW



- Practical loop: scope, acquire data, build features, train, evaluate, deploy, monitor, retrain
- Lifecycles are iterative; returning from validation to data acquisition is normal, not failure
- Each stage has one or two dominant tradeoffs worth naming early
- Match decision-review questions to each stage: framing, sourcing, modeling, validation, deployment, monitoring
- Shared stage vocabulary helps case discussions travel across teams and roles

BETTER DECISIONS
BRIGHTER POSSIBILITIES

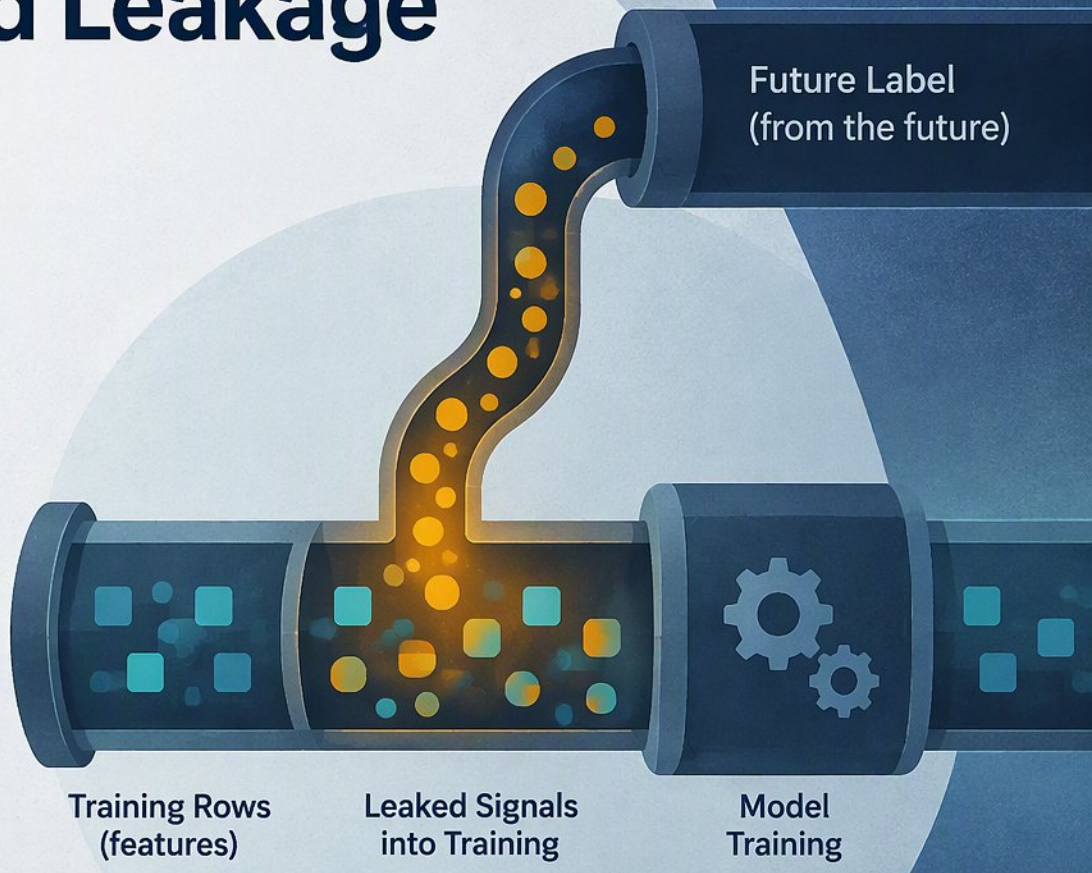
Framing the Problem and the Metric

- Separate business question, modeling question, and measurable metric—they diverge often
- Pair an optimization metric with a guardrail, e.g., recall at fixed review capacity
- Case: fraud request reframed from 'maximize accuracy' to capacity-constrained recall
- Tradeoff: simplicity for stakeholder alignment vs. fidelity to real cost structure
- Risk: proxy metrics drift from the business objective; define metrics that survive handoffs



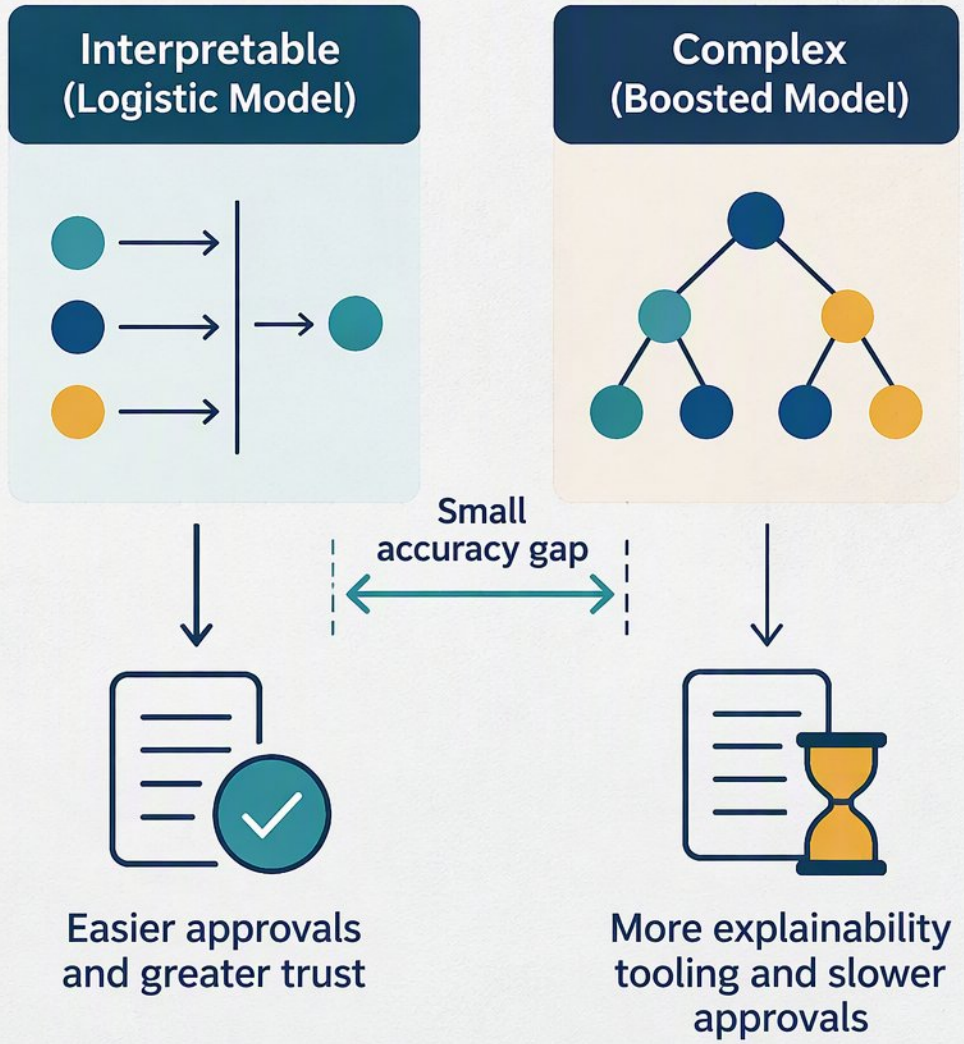
Data Sourcing and Leakage

- Four leakage shapes: target, temporal, group, and pipeline (fitting) leakage
- Case pattern: churn model with near-perfect offline AUC collapsed after deployment
- Validate at deployment timing: point-in-time splits, embargo gaps, group-aware splits
- Checks: feature availability at prediction time, lineage review, shuffled-target test
- Tradeoff: richer feature sets vs. maintainability, privacy, and leakage risk



Interpretability vs. Accuracy

- Simpler models win when audits, trust, or adoption matter
- Small accuracy gains rarely justify explainability tooling and slower approvals
- On tabular data, interpretable models often land within a small margin
- SHAP explains the model's behavior, not causality or business action
- Raise complexity only when lift changes a downstream decision



Validation, Backtesting, and Errors That Lied

BETTER VALIDATION. MORE RELIABLE MODELS.

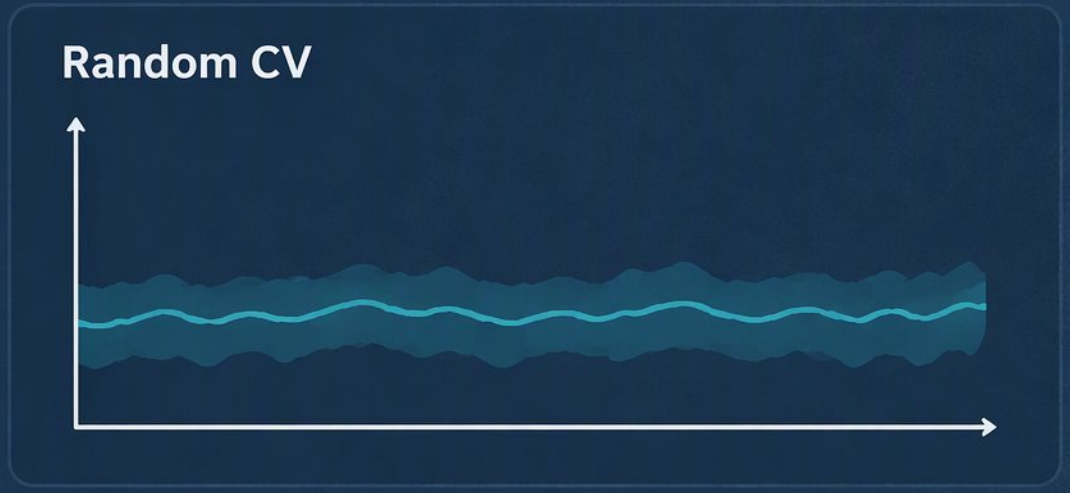
Case pattern:
random CV passed, chronological backtesting failed

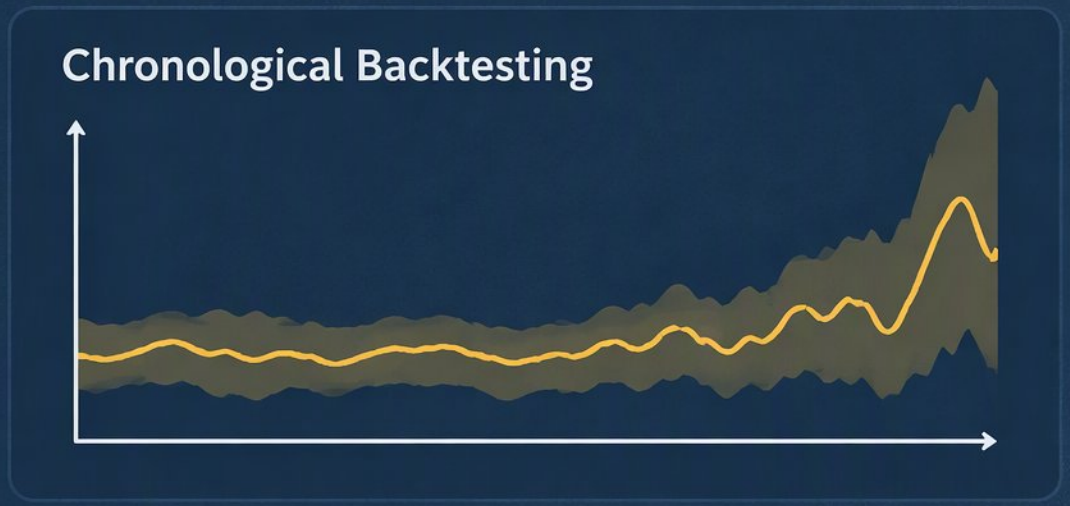
Split choice:
random, temporal, blocked/expanding-window, grouped, nested CV

Backtest overfitting:
test-set peeking and metric shopping inflate reported performance

Report uncertainty:
confidence intervals and error decomposition, not one point estimate

Tradeoff:
validation rigor vs. speed and training data; rigor scales with deployment cost

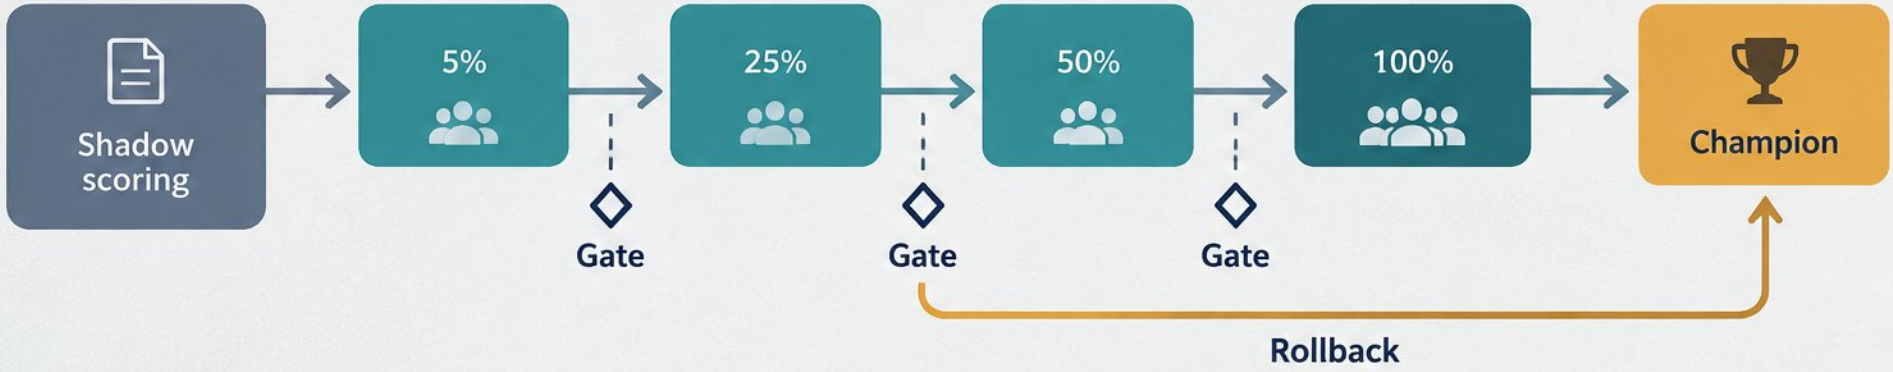




SAME DATA. A CLEARER STORY.

Deployment Patterns: Shadow, Canary, Rollback

Safer changes. More confidence. Better outcomes.

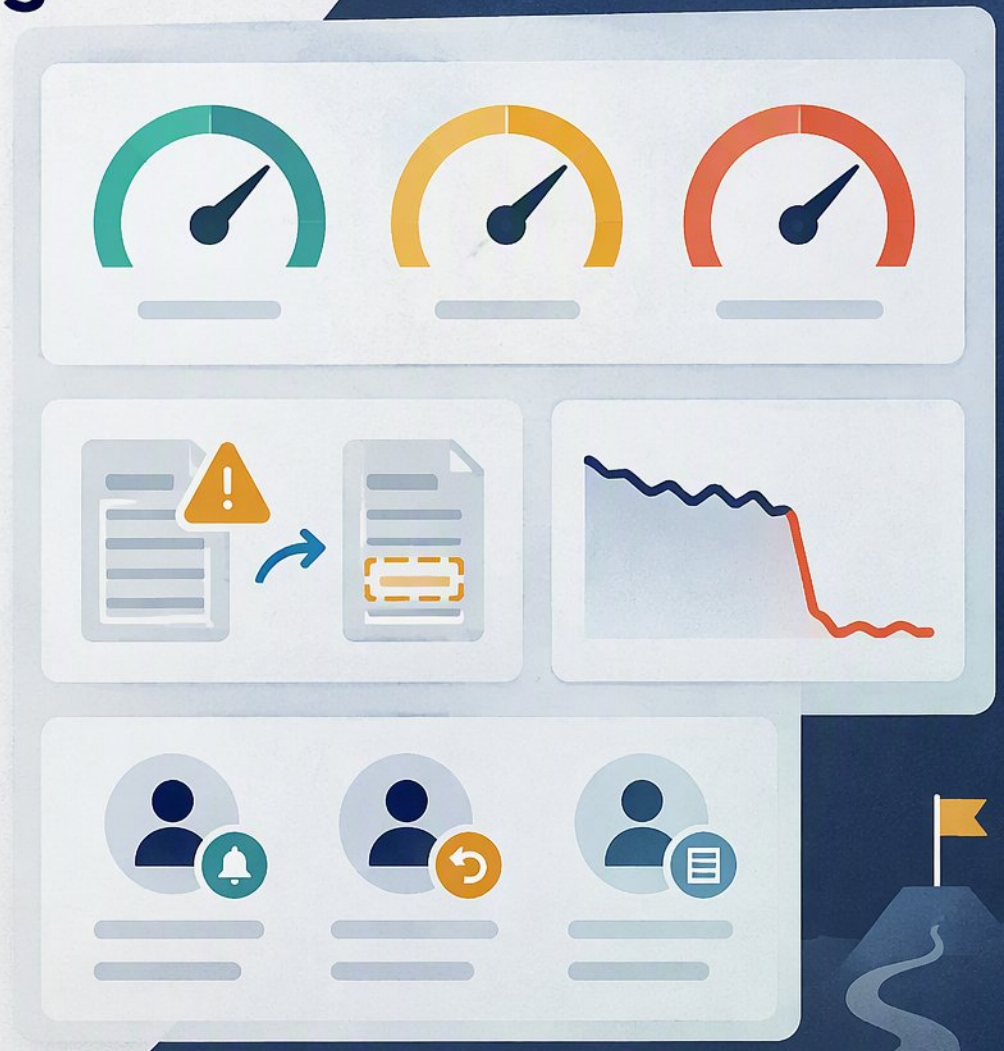


- **Shadow scoring**: challenger predicts live traffic, logged but not served
- **Canary rollout:** 5% → 25% → 50% → 100%, with a gate at each step
- **Candidate** must beat champion by a margin on a forward-in-time holdout
- **Rollback** is a one-step revert to a warm, versioned champion
- **Case:** silent recommender degradation after a product change — schema contracts would have caught it
- **Tradeoff:** fast shipping vs. rollout safety, weighed against the cost of a wrong deployment

DEPLOY THOUGHTFULLY
BUILD TOMORROW

Monitoring and Retraining Tradeoffs

- Monitor data quality, feature drift, prediction drift, outcomes, and business guardrails
- Combine three retraining triggers: drift alarms, scheduled retrains, and manual retrains
- Incident case: a renamed payment field produced silent zeros and recall collapse
- Cadence tradeoff: staleness vs. stability vs. operational cost; don't retrain on noise
- Define ownership: who is alerted, who rolls back, what counts as a model incident



OBSERVE LEARN IMPROVE

Fairness, Privacy, and Risk

RESPONSIBLE AI
FOR A FAIRER TOMORROW

- Accuracy can hide unequal error rates across subgroups.
- Fairness notions conflict: calibration vs. equalized odds is a policy choice.
- Differential privacy can widen accuracy gaps for minority groups.
- Show the Pareto frontier; do not bury tradeoffs in one tuned metric.
- Document residual risk and name a sign-off owner.



SAFER MODELS STRONGER PEOPLE BRIGHTER OUTCOMES

Communicating Decisions and Residual Risk

Clear context. Shared understanding. Confident decisions.

SAME RECOMMENDATION. DIFFERENT DEPTH.

1. Decision memo: context, options, evidence, recommendation, residual risk, reversal triggers
2. "97% accurate" becomes "3 in 100 flagged customers are wrong — here is review cost"
3. Tailor depth for executives, engineers, and domain experts; same recommendation
4. Annotate charts, compare to naive baselines, ground metrics in concrete examples
5. Make tradeoffs falsifiable: name owner, expiry, and re-evaluation triggers

Decision Memo

- Context
- Options
- Evidence
- Recommendation
- Residual risk
- Reversal triggers

Tailor the depth, keep the message aligned.

Executives



Engineers



Domain experts



BETTER CONVERSATIONS LEAD TO MORE RESILIENT OUTCOMES.

Case Clinic: Frame, Split, and Measure

ANALYZE / DISCUSS / DECIDE

BETTER QUESTIONS
LEAD TO CLEARER CHOICES

Exercise A
rewrite fraud accuracy as capacity-constrained recall plus a guardrail.

Exercise B
find two churn leakage paths; propose point-in-time validation with embargo.

Deliverable
half-page decision record — problem, options, evidence, recommendation, reversal trigger.

Peer review
Is the tradeoff explicit?
Is the evidence deployment-valid?
Is residual risk named?

Debrief
compare team choices to surface context dependence and reasonable disagreement.

DIFFERENT PERSPECTIVES
STRONGER DECISIONS

Decision Record (Half-Page Template)

Problem	
Options	
Evidence	
Recommendation	
Reversal trigger	

Peer Review Rubric

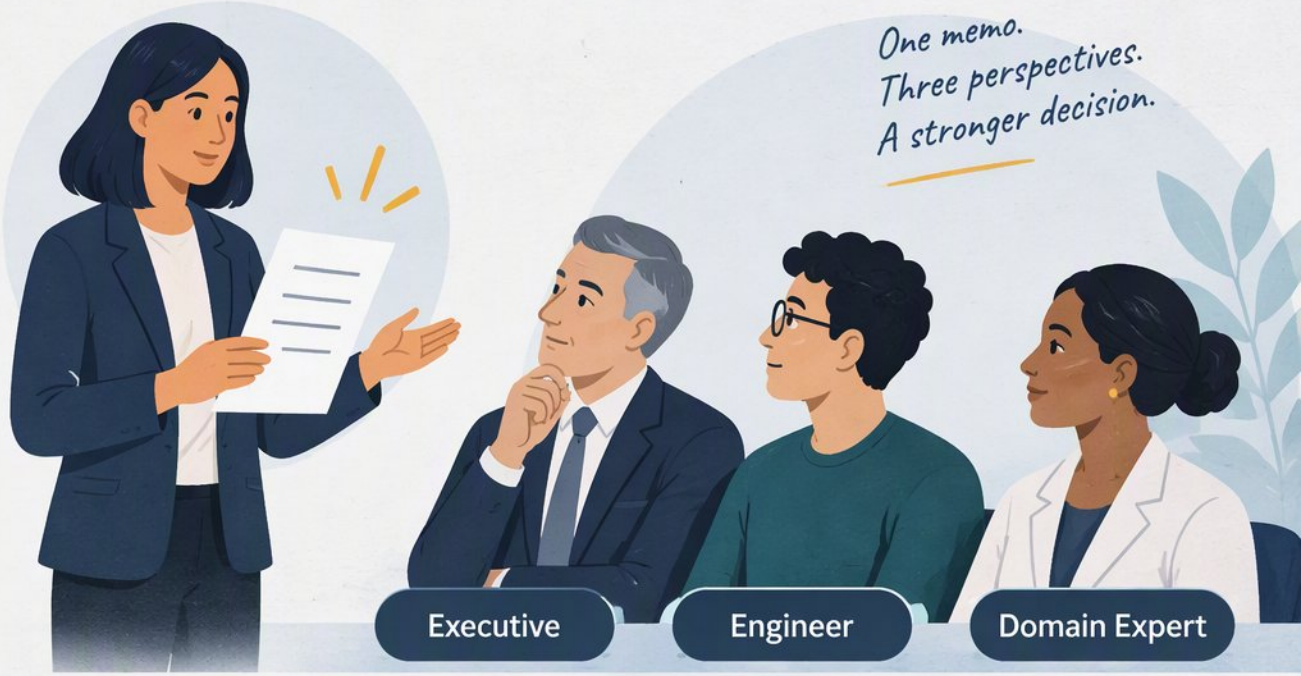
- Is the tradeoff explicit?
- Is the evidence deployment-valid?
- Is residual risk named?

SAME PROBLEMS. BRIGHTER THINKING.

Workshop: Defend a Tradeoff

SAME RECOMMENDATION
DIFFERENT PERSPECTIVES

1. Pick one case packet: monitoring incident, fairness conflict, or interpretability decision.
2. Identify the decision point, then write a one-page decision memo from the course template.
3. Each memo must name one guardrail metric and one residual-risk statement.
4. Role-play: present the same recommendation to an executive, an engineer, and a domain expert.
5. Scoring rubric: tradeoff named, evidence deployment-valid, risk owned, reversal triggers stated.
6. Close with a personal checklist of three tradeoffs to watch next quarter.



Scoring Rubric

- tradeoff named
- evidence deployment-valid
- risk owned
- reversal triggers stated

A sharper next quarter starts with your checklist.

Takeaways and Your Decision Checklist

CLEARER DECISIONS. STRONGER OUTCOMES.

PEOPLE
IDEAS
BETTER
DECISIONS

Five recurring tradeoffs

- cost/speed
- interpretability/accuracy
- exploration/exploitation
- leakage/robustness
- performance/compliance

Four questions per project

- What is the decision?
- What metric reflects cost?
- Is evidence deployment-valid?
- Who owns residual risk?

Every case is a decision record; keep failures and null results in the team library

Write tradeoffs explicitly and falsifiably so they survive handoffs and audits

Adopt the decision memo and review rubric as standing team practice

PersonWise.

PersonWise • AI digital-human course platform

Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/68723bcb/public