

PRODUCT MANAGEMENT + AI PRODUCT MANAGEMENT

THE PRODUCT LEADER'S YEAR

AUTHORITY EDITION

A comprehensive field guide to product judgment, strategy, discovery, delivery, data, AI, leadership, and portfolio proof

ANUSTUP MUKHERJEE



THE PRODUCT LEADER'S YEAR

Authority Edition

Product management and AI product management from first principles
to portfolio-grade practice

THE PRODUCT LEADER'S YEAR

AUTHORITY EDITION

Anustup Mukherjee

First authority edition, 2026

Copyright © 2026 Anustup Mukherjee. All rights reserved.

The explanations, diagrams, exercises, solutions, and portfolio studios in this edition are original. Established thinkers are attributed where their ideas form part of the intellectual lineage. No attribution implies endorsement.

The legal, regulatory, security, financial, clinical, and certification material is educational. Revalidate current primary sources and obtain qualified review for consequential use.

Editorial promise

A rigorous book for commercially experienced readers building global product and AI product capability.

This edition is designed to teach judgment rather than recite terminology. Every chapter contains an original decision model, explicit intellectual lineage, critical reading, a worked example, failure modes, applied problems, worked reference solutions, and a portfolio artefact. The ten studios then combine those skills into substantial projects that can be built with synthetic or permissioned data.

Examples are designed for readers moving from marketing, operations, consulting, entrepreneurship, or adjacent roles into product and AI product work. Claims remain evidence-bounded: a synthetic demonstration is not a live result; a benchmark is not commercial impact; and a pilot is not proof of scale.

PRIVACY AND DISCLOSURE STANDARD

All people, clients, organisations, metrics, and operational situations in teaching examples are fictionalised composites unless a public source is cited. No example discloses confidential employer or client information. Brand references describe public publishing or illustrative service contexts; they do not claim client endorsement or measured commercial results.


HOW TO READ AUTHORITATIVE IDEAS

Learn the durable decision behind each author's framework. Credit the source. Test the idea against the situation. Keep the evidence and consequences visible. Never substitute reputation for judgment.

How to use the authority edition

Study one chapter, make one decision, build one artefact, and invite one critique.

- Career transition: complete the first chapter of every part, then build Cases 2, 4, 6, and 9 as flagship work.
- Workflow-automation services: convert a case studio into a paid diagnostic and bounded pilot only after adapting it to permissioned evidence and an explicit client boundary.
- Product leadership: use the methods as review systems for strategy, discovery, delivery, metrics, AI quality, and organisational design.
- Teaching: assign the problems before showing the worked solutions; assess reasoning, evidence, and limitation rather than framework vocabulary.
- Revalidation: check dated regulations, standards, vendor capabilities, and platform documentation before consequential use.



Contents

I. PRODUCT JUDGMENT AND THE ROLE	9
Product management is accountable problem selection	11
Translate operating experience into product evidence	17
Manage the four product risks	23
Design ethics, accessibility, and trust into the product	29
Work through a product operating model	35
Improve decision quality under uncertainty	41
Define product principles and decision rights	47
Build a personal product apprenticeship	53
II. CUSTOMER DISCOVERY	59
Discover problems before designing solutions	61
Interview for behaviour and causality	67
Observe workflows and invisible work	73
Use Jobs to Be Done without turning it into a slogan	79
Segment by situation, behaviour, and constraint	85
Map journeys and service blueprints	91
Build an opportunity solution tree	97
Map assumptions and test the riskiest first	103
Synthesize research into defensible insight	109
III. STRATEGY AND BUSINESS DESIGN	115
Write a strategy kernel	117
Make an integrated set of strategic choices	123
Connect vision, strategy, roadmap, and backlog	129
Position the product deliberately	135
Analyse competition and substitutes	141
Design the business model as a system	147
Choose pricing and packaging	153
Manage portfolios and cannibalisation	159
Turn strategy into outcomes and roadmaps	165
IV. EXPERIENCE, DELIVERY, AND SYSTEMS	171
Frame human-centred product experiences	173
Prototype to answer a question	179
Run usability tests that change decisions	185
Write requirements that preserve the problem	191
Use story maps to design coherent releases	197
Apply agile and Scrum with judgment	203
Shape work before committing capacity	209
Develop technical fluency for product decisions	215
Design platforms, APIs, and ecosystems	221
Manage reliability, quality, and delivery performance	227

V. DATA, GROWTH, AND COMMERCIAL PRACTICE	233
Design a metric system that explains value	235
Analyse funnels, cohorts, and retention	241
Run experiments and reason about causality	247
Diagnose churn and build durable retention	253
Design growth loops and network effects	259
Plan go-to-market and launch	265
Build enterprise products across the buying system	271
Model unit economics and product investment	277
VI. AI PRODUCT MANAGEMENT	283
Understand AI, machine learning, and language-model systems	285
Select AI use cases with evidence	291
Build data readiness and provenance	297
Design prompts, context, and structured outputs	303
Build retrieval-augmented generation systems	309
Choose workflows, tools, and agents	315
Create evaluations that govern release	321
Design human-AI interaction for control	327
Govern AI as a living management system	333
Navigate privacy, regulation, and responsible deployment	339
Operate AI for reliability, cost, and change	345
Manage AI incidents and vendor dependency	351
VII. LEADERSHIP, OPERATIONS, AND CAREER	357
Lead decisions across functions	359
Handle conflict, negotiation, and executive communication	365
Build product operations without creating bureaucracy	371
Design empowered teams and healthy product organisations	377
Turn service work into a product learning system	383
Create productised automation offers that win projects	389
Build a defensible product and AI product portfolio	395
Prepare for product interviews and case presentations	401
Run a three-year learning and authority plan	407
VIII. MARKET-READY PRODUCT CAREER	413
Read the product job market by role archetype	415
Build product sense, discovery, and experience judgment	423
Use analytics, SQL, experimentation, and data storytelling	429
Develop technical product fluency across APIs, cloud, data, and security	435
Build AI-product fluency in retrieval, agents, evaluations, and operations	443
Lead execution, communication, and ambiguity	451
Connect go-to-market, growth, pricing, domain expertise, and portfolio proof	457
IX. PRODUCT CAREER LAUNCH	463
Choose target roles and build an evidence-based gap map	465
Translate prior experience into credible product evidence	473
Build a portfolio hiring managers can inspect	479

Master product-sense and discovery interviews	485
Master execution, metrics, SQL, and experimentation interviews	493
Master strategy, technical, and AI-product interviews	499
Master behavioural, leadership, and stakeholder interviews	505
Run the job search, evaluate offers, and enter the first 90 days	513
JOB-MARKET EVIDENCE REGISTER	521
JOB-READINESS ASSESSMENT	523
X. PORTFOLIO PROJECT STUDIOS	527
Portfolio studio master scorecard	528
Campaign performance command centre	529
Lead intake, qualification, and CRM handoff	535
WhatsApp-first appointment and no-show reduction	541
AI support triage and agent copilot	547
Invoice and document extraction with human review	553
Permission-aware RAG research assistant	559
Marketplace provider onboarding and activation	565
Subscription churn diagnosis and win-back	571
Agentic campaign quality-assurance workflow	577
AI governance and evaluation control centre	583
REFERENCE HANDBOOK	589
Product analytics formulas that support decisions	589
Research and statistics for product decisions	591
Technical fluency for product leaders	592
AI evaluation laboratory	594
B2B, enterprise, and platform product practice	595
Product interview and case-presentation playbook	597
Twelve-week portfolio curriculum	598
TOOLKIT	601
INTELLECTUAL LINEAGE AND READING PATH	607
Glossary	609
REFERENCES AND REVALIDATION	611

I. PRODUCT JUDGMENT AND THE ROLE

Build the judgment, evidence discipline, ethics, and operating model needed to make consequential product decisions.


Product management is accountable problem selection

The work begins with choosing a valuable problem and ends with evidence about whether the choice improved an outcome.

A feature request names an intervention; a product problem names a harmful condition somebody experiences. Automating a weekly report may reduce preparation effort while leaving the consequential delay untouched: people still cannot decide which anomaly requires action. Start by identifying the decision, its owner, its deadline, and the consequence of missing it. Then distinguish a customer outcome, such as reaching a defensible decision sooner, from an output, such as generating a report. An output can be useful, but accountability requires a reasoned link between the output and the condition the team intends to improve.

Problem selection has an opportunity cost. The same analyst and engineering hours cannot simultaneously improve report generation, source reconciliation, and exception ownership. Compare these options using the size of the avoidable burden, the team's ability to influence it, and the uncertainty about the proposed mechanism. A large total delay is not automatically an addressable delay. If an external approval board contributes most waiting time, a faster interface may barely change the result. Decompose elapsed time into controllable and uncontrollable portions before promising improvement, and look for interventions that reduce the dominant controllable portion.

An accountable product decision includes a rule for changing course. Before building, specify the observation that would justify continuing, narrowing, or stopping. A manual exception digest can test whether traceable evidence and named owners reduce decision delay without first building integrations. The result must include the labour required to create and review the digest, because an apparent customer benefit can merely transfer work backstage. Treat a small pilot as evidence about a mechanism in a defined setting. It supports the next commitment; it does not establish durable adoption, broad causality, or savings across every workflow.

PRODUCT MANAGEMENT IS ACCOUNTABLE PROBLEM SELECTION

Outcome selection leads to affected-party evidence, an assumption test, and a continue-or-stop rule; learning can revise the selected problem.

INTELLECTUAL LINEAGE

Marty Cagan's product-risk and empowered-team writing; Melissa Perri's critique of the build trap.

Apply the method

1. Outcome

Outcome: define the elapsed time from a report-ready timestamp to a recorded anomaly decision. State the decision deadline and acceptable error burden. Exclude report preparation time from this measure so the team can see whether the proposed intervention reaches the actual bottleneck.

2. Affected people

Affected people: observe an analyst, an approver, and the person who acts on an anomaly. Capture what each needs to trust the decision. Identify who will inherit review or correction work if a digest replaces the report.

3. Critical assumption

Critical assumption: write the claim that evidence gaps and unclear ownership cause most addressable delay. Separate it from the claim that automated reporting saves effort. List the historical events that would contradict the ownership explanation before selecting the pilot.

4. Cheapest evidence

Cheapest evidence: assemble a manual digest for a bounded set of synthetic or permissioned anomalies, with provenance and named owners. Compare decision timestamps and review effort against the existing process. Keep preparation visible rather than concealing human work behind a prototype.

5. Decision rule

Decision rule: set a minimum reduction in decision delay and a maximum new weekly operating burden. Apply both conditions together. Agree which mistakes require stopping the test even when the average time improves, including unsupported escalation or lost anomaly ownership.


6. Learning loop

Learning loop: review individual delayed decisions alongside the aggregate. Determine whether the intervention changed evidence access, ownership, or approval behaviour. Refine the problem statement when another constraint dominates, and record why the next allocation of capacity follows from the evidence.

Worked example

Synthetic example: a studio observes a weekly reporting workflow that requires four analyst hours to prepare the report and twelve elapsed hours after publication to decide on anomalies. A report-generation option would save three analyst hours but leave the twelve-hour decision delay unchanged. A manually prepared decision digest would consume two additional analyst hours while reducing decision delay to seven hours in a simulated pilot. The digest produces a five-hour reduction, calculated as twelve minus seven, or 41.7% of the original delay. These are different kinds of hours: elapsed waiting cannot be added to analyst effort or treated as a labour saving.

The studio's synthetic decision rule requires at least a 30% reduction in anomaly decision delay and no more than two extra analyst hours per week during discovery. The digest meets both boundaries; report generation meets neither the outcome improvement condition nor the causal hypothesis. The defensible choice is a bounded digest pilot with named decision owners, not a claim that the digest is already a scalable product. If operating the digest later requires three additional hours, it breaches the labour boundary. The team would need to simplify preparation or revise the commercial model before expanding the pilot, even if the decision-time improvement persists.



Failures to recognise

- Adding elapsed waiting hours to staff effort and presenting the sum as financial savings.
- Defining success as reports produced when the selected problem is delayed anomaly decisions.
- Expanding after the outcome improves while ignoring the pilot's explicit operating-effort ceiling.

Practice problems

Problem 1

Using the synthetic figures, calculate the digest's absolute and percentage reduction in decision delay. State whether this equals saved analyst labour.

Problem 2

If digest operation rises to three additional analyst hours weekly while decision delay remains seven hours, does the original decision rule permit expansion?

Problem 3

Which option directly addresses the selected outcome: report generation or the decision digest? Identify the measurement that could overturn that choice.

Reference answers

Answer 1: Separate elapsed delay from labour

Decision delay falls from twelve to seven hours, so the absolute reduction is $12 - 7 = 5$ elapsed hours. Dividing five by the twelve-hour baseline gives 0.4167, or approximately 41.7%. Analyst labour instead rises by two hours in the digest pilot. The result therefore cannot be described as five saved analyst hours. It is a shorter decision interval purchased with additional operating effort. The team must retain separate measures for delay, labour, and decision quality when deciding whether the trade-off is worthwhile.

Answer 2: Apply the full decision rule

The time condition still passes: the reduction remains 41.7%, exceeding the required 30%. The labour condition fails because three additional hours exceeds the two-hour ceiling by one hour, or 50% of that ceiling. The original rule therefore does not permit expansion. Continue only within the authorised discovery boundary while testing a lower-effort preparation method, or explicitly reconsider the economic constraint. Quietly relaxing the threshold after seeing a favourable time result would remove the accountability that made the pilot useful.

Answer 3: Choose the intervention linked to the problem

The decision digest directly addresses the interval between report publication and an anomaly decision. Report generation reduces the separate preparation stage and leaves the twelve-hour outcome unchanged in this example. The choice could be overturned by repeated timestamp evidence showing that faster report availability moves the actual decision deadline earlier, or by a digest that increases decision errors despite reducing delay. Measure recorded decisions with comparable anomalies and owners. A favourable opinion about either interface would not establish the required causal link.

PORTFOLIO APPLICATION

Apply product management is accountable problem selection to one of the ten portfolio studios. Attach the reasoning, the result of one test, a rejected alternative, and the evidence that would change your recommendation.


Translate operating experience into product evidence

Marketing and operations experience becomes credible product strength when decisions, constraints, artefacts, and measured learning are visible.

Transferable operating experience becomes product evidence when a reader can inspect the reasoning behind a change. Describe the original workflow, the observed friction, the constraints, and the decision you actually controlled. Improving a campaign handoff might involve redefining a required information field, negotiating an approval rule, or testing a reusable checklist. Those actions demonstrate problem framing and cross-functional judgment without implying that the role carried formal product ownership. A credible account preserves the employment title and distinguishes personal contribution from the team's work. The objective is an accurate, decision-rich case rather than a retroactive change of professional identity.

A case should connect artefacts to claims. A before-and-after handoff map can support a claim about process structure; timestamp records can support a claim about elapsed time; a decision memo can show which alternatives were considered. None automatically proves incremental revenue or a causal treatment effect. Specify the denominator, measurement window, and changes that occurred alongside the intervention. If staffing increased during the same month, a shorter queue may have several explanations. Report the descriptive improvement and its limitations. Strong evidence is often modest and well bounded, rather than impressive but impossible for another person to verify.

Confidentiality changes how evidence is presented, not whether reasoning can be examined. Replace names and sensitive commercial figures with synthetic examples or an explicitly normalised index, and explain the transformation. Rebuild a public-facing artefact from scratch instead of exporting private operational records. Keep the hypothetical reconstruction distinct from any verified historical result. A reviewer should be able to see the candidate's question, alternatives, trade-off, and learning while remaining unable to identify a private client or colleague. If a measurement is unavailable, state that gap and describe a suitable future measurement instead of manufacturing a result.



Apply the method

1. Past situation

Past situation: choose one bounded handoff problem with an observable start and finish. Describe the actual role and responsibilities. Separate the workflow context from confidential identifiers, and record which baseline measures exist before deciding what the case can claim.

2. Product question

Product question: translate the operational complaint into a question about behaviour or constraints, such as whether missing approval information prolongs the handoff. Define the affected user and decision point. Avoid treating a successful campaign as evidence for an unrelated product hypothesis.

3. Personal decision

Personal decision: name the choice you made, the alternatives rejected, and the constraint that mattered. Distinguish proposing a checklist from owning its implementation or approving its rollout. Attach a reconstructed decision record that shows reasoning without disclosing private documents.

4. Cross-functional action

Cross-functional action: explain how planning, delivery, and approval roles contributed to the change. Show the disagreement or dependency that had to be resolved. Credit those contributions explicitly, and describe how the decision boundary became clearer after the collaboration.

5. Evidence

Evidence: compare the same handoff measure across clearly stated windows and denominators. Present descriptive calculations with source status. Use synthetic figures for public exercises, label them visibly, and avoid turning a shorter cycle time into an unmeasured revenue claim.

6. Limitation

Limitation: identify concurrent staffing, seasonality, scope changes, and missing quality data. State which claims the artefacts support and which remain uncertain. Finish with a specific next measurement or test that would improve confidence in the explanation for the observed change.

Worked example

Synthetic portfolio reconstruction: forty campaign handoffs in a baseline month took an average of five working days from draft completion to approved delivery. Forty handoffs in a later month averaged three working days after a completeness checklist and named approval owner were introduced. The descriptive reduction is two days, or $2 \div 5 = 40\%$. The candidate personally designed the checklist and facilitated agreement on ownership; another team implemented the tracking field. The case should report those contributions separately. No employment title is changed, and no actual employer, client, campaign, or private performance figure is represented by this constructed example.

The candidate then discovers that the approval team increased from two people to three in the later month. This concurrent change weakens a claim that the checklist alone caused the reduction. A defensible case says that average handoff time fell from five to three days during the combined operational changes, with incomplete records dropping from twelve of forty to four of forty. Missing-information frequency declined from 30% to 10%, an absolute decrease of twenty percentage points. That measure is more closely connected to the checklist mechanism, although it still does not isolate causality. The public artefact includes the caveat and a proposal to examine comparable approval capacity.

CASE DECISION TRACE



BOUNDED RECOMMENDATION

Preserve accurate scope while making transferable judgment inspectable.

Failures to recognise

- Changing an employment title to make a transferable workflow case sound like formal product ownership.
- Presenting a synthetic public reconstruction as an actual historical client result.
- Claiming causality from a before-and-after comparison that also changed approval capacity.

Practice problems

Problem 1

Calculate the relative decrease in average handoff time and the absolute percentage-point decrease in incomplete records.

Problem 2

If the later month contains eight incomplete records among forty handoffs, recalculate the incomplete-record rate and its change from baseline.

Problem 3

Does the staffing increase invalidate every claim, or only some? Write one supported statement and one unsupported causal statement.

Reference answers

Answer 1: Report both forms of change

Average handoff time falls by $5 - 3 = 2$ working days. Relative to five days, that is $2 \div 5 = 40\%$. Incomplete-record frequency starts at $12 \div 40 = 30\%$ and ends at $4 \div 40 = 10\%$. The absolute decrease is twenty percentage points; the relative decrease is $20 \div 30 = 66.7\%$. Use the percentage-point figure when comparing rates directly. These calculations describe the synthetic windows and do not establish that a checklist independently caused either improvement.

Answer 2: Recalculate the changed assumption

Eight incomplete records among forty handoffs gives $8 \div 40 = 20\%$. Compared with the 30% baseline, the absolute decrease is ten percentage points and the relative decrease is $10 \div 30 = 33.3\%$. The case can describe this smaller improvement without changing the five-to-three-day handoff calculation. It should not keep the original twenty-point claim after updating the denominator or numerator. Check that incomplete means the same thing in both periods, because a looser definition could create an artificial improvement.

Answer 3: Keep description separate from attribution

The staffing increase does not erase the recorded descriptive difference. A supported statement is that the synthetic later period had a three-day average handoff time versus five days in the baseline period while checklist and staffing changes both occurred. An unsupported statement is that the checklist alone reduced handoff time by 40%. The extra approver could reduce queue time independently. Stronger attribution would require comparable staffing conditions or another design that separates the effects, plus consistent handoff scope and quality measures.

PORTFOLIO APPLICATION

Apply translate operating experience into product evidence to one of the ten portfolio studios. Attach the reasoning, the result of one test, a rejected alternative, and the evidence that would change your recommendation.


Manage the four product risks

Value, usability, feasibility, and viability must all be reduced before a team scales commitment.

The four product risks ask different questions: whether people want the proposed value, can use the solution, whether the team can build it, and whether operating it works for the business. An AI proposal assistant can perform well on one question and fail another. A polished demo may show that text can be generated while leaving permission checks, correction effort, and willingness to use the draft unresolved. Keep separate evidence for each risk instead of calculating an average confidence score that lets a strong usability result compensate for a prohibited data flow or an unworkable cost structure.

Tests should target the uncertainty that could invalidate the proposal. Observe a manager checking citations to study usability; inspect permission behaviour with deliberately inaccessible documents to study feasibility and business constraints; compare verified drafting time with the current process to study value. A technical benchmark alone cannot show that managers will trust the output in a real approval workflow. Likewise, interview enthusiasm cannot establish reliable permission enforcement. Every test needs a representative task, an observable result, and a boundary on what it proves. Record missing evidence as missing rather than treating the absence of failure as a pass.

Risk reduction is staged commitment, not an assertion that all risk can disappear. A low-cost trial may justify a limited pilot when the remaining exposure is bounded and recoverable. Some conditions should remain hard gates: a proposal assistant must not disclose documents its user cannot access, even if it saves drafting time. Other conditions may be negotiated, such as a lower initial margin while review work is being simplified. State which conditions are mandatory and who may accept a trade-off. Revisit the risk register when document types, usage volume, permissions, or serving costs change.



Apply the method

1. Value

Value: compare the complete drafting-and-verification task with the current process. Include correction and citation checking. Require a meaningful reduction in verified completion time rather than a faster first draft, and note which proposal types the test actually covers.

2. Usability

Usability: ask representative managers to locate evidence, correct an unsupported sentence, and submit a reviewed draft. Observe assistance required and misunderstandings. Test failure and recovery paths, because effortless generation can hide difficulty in recognising or correcting a consequential error.

3. Feasibility

Feasibility: test access controls, retrieval quality, document updates, and peak-load behaviour with authorised fixtures. Include a document the user cannot access. Verify enforcement in the serving path; a warning in the interface does not substitute for a permission check.

4. Viability

Viability: calculate serving cost, reviewer labour, support burden, and proposed revenue per accepted proposal. Have the relevant commercial and governance owners examine the assumptions. Separate acceptable discovery spending from the economic conditions required for a sustainable managed service.

5. Evidence gaps

Evidence gaps: maintain one row per material claim with test result, sample boundary, and unresolved issue. Mark a simulated result as simulated. Rank the gaps by their ability to invalidate the pilot, without averaging a hard permission failure into softer benefits.

6. Kill or continue


Kill or continue: apply the four conditions together before increasing access or volume. Continue narrowly when recoverable gaps have owners and bounded exposure. Stop or redesign when permission enforcement fails, verified work increases, or the commercial model cannot cover ongoing review.

Worked example

Synthetic example: an assistant reduces initial drafting from sixty to twenty minutes, but managers spend twenty-five minutes checking citations and correcting unsupported text. Verified completion takes $20 + 25 = 45$ minutes, a fifteen-minute improvement over the sixty-minute manual baseline. At a synthetic reviewer labour cost of \$24 per hour, review costs $25 \div 60 \times \$24 = \10 per proposal. Model serving costs another \$2. A service price of \$15 therefore leaves \$3 contribution before fixed costs and other variable support. The team evaluates value using the forty-five-minute completed task, rather than advertising forty minutes saved from draft generation alone.

The synthetic pilot also tests ten task scenarios and an inaccessible document. Eight of ten correction tasks are completed without help, but one permission fixture leaks an inaccessible sentence. The team's predeclared requirements are at least a ten-minute verified-time reduction, at least eight of ten successful correction tasks, a minimum \$3 contribution per proposal, and zero permission leaks. The first three conditions pass at their stated boundaries; the permission condition fails. The appropriate decision is to repair access enforcement and retest before any wider pilot. Strong performance elsewhere does not offset the leak. If serving cost rises to \$4, contribution also falls below the economic requirement.

CASE DECISION TRACE



BOUNDED RECOMMENDATION

Run a different test for each material risk.

Failures to recognise

- Using a generation-speed benchmark as evidence that managers save time on verified proposals.
- Hiding a permission leak inside an average score across unrelated risk categories.
- Calculating AI unit economics without paid human review and correction effort.

Practice problems

Problem 1

Calculate verified time saved, reviewer cost, and contribution per proposal using the synthetic inputs.

Problem 2

If serving cost rises from \$2 to \$4 while review time stays twenty-five minutes, does the economic condition pass?

Problem 3

Which risk prevents expansion despite the other results? Explain why averaging test scores is inappropriate.

Reference answers

Answer 1: Count the complete task and its cost

Verified completion takes twenty minutes drafting plus twenty-five minutes review, or forty-five minutes. Relative to the sixty-minute baseline, the time saving is fifteen minutes, equivalent to 25%. Reviewer labour is $25 \div 60 \times \$24 = \10 . With \$2 serving cost, total stated variable cost is \$12 and contribution is $\$15 - \$12 = \$3$. This meets the stated contribution floor exactly. Fixed overhead, customer acquisition, and omitted support costs must still be considered before calling the overall service profitable.

Answer 2: Recheck the economic boundary

Serving cost at \$4 combines with \$10 reviewer labour to produce \$14 stated variable cost per proposal. Contribution becomes $\$15 - \$14 = \$1$, two dollars below the \$3 floor. The economic condition therefore fails even though verified completion remains forty-five minutes. The team could investigate lower review effort, a different price, or a narrower document scope, but it cannot report the original margin. A cost change is new risk evidence and should reopen the viability decision before volume is increased.

Answer 3: Treat permission enforcement as a gate

The inaccessible sentence exposes a permission failure, involving feasibility of enforcement and the business consequences of disclosure. The predeclared rule requires zero leaks, so one observed leak prevents expansion. An average across eight successful usability tasks, favourable time saving, and economics would conceal the non-negotiable condition. Repair the serving-path access checks and retest denied access, changed permissions, and document boundaries. Passing a subsequent bounded test supports a new pilot decision; it does not establish that every possible access scenario is safe.

PORTFOLIO APPLICATION

Apply manage the four product risks to one of the ten portfolio studios. Attach the reasoning, the result of one test, a rejected alternative, and the evidence that would change your recommendation.

