



The Framework and Decision Worksheets

This appendix brings the parts of the Context Principle together in one place so you can use them when a decision reaches your desk. The chapters explain the reasoning behind each part. Here, I've kept the reference short and left room to work.

The worksheets are starting points. Copy them, adapt the prompts to your setting, or move them into a form your organization already uses. The depth should fit the decision: a contained test with fictional material might need a few lines, while a use that could affect care or private information will need supporting work attached. A completed worksheet helps a team explain its reasoning. It doesn't approve a use or replace the clinical, privacy, security, legal, or other review a particular use requires. Resources and updates for this book are at contextprinciple.com/healthcare.

The Framework at a Glance

PART OF THE FRAMEWORK	WHAT IT HELPS US DO	RETURN TO
Context Principle	Develop enough understanding for the intervention and its possible impact.	Chapters 1 and 4
Context Cycle	Observe, establish Context, Interpret, Decide, Act, and Reflect/Refine, then observe again. Use what happens to update the next decision.	Chapter 4
Four Lenses	Bring together Evidence about the proposal, Context about the setting, Expertise to interpret what we find, and Lived Experience from people affected.	Chapter 5
Decision Triangle	Examine the cost of being wrong, the difficulty of reversing harm, and how many people may be affected and how deeply. A serious concern in one area still deserves attention when the others appear manageable.	Chapter 6
Sufficient Understanding Test	Examine lens coverage and its weakest area, proportional depth, important uncertainties, meaningful input from affected people, and investigated failure conditions.	Chapters 1 and 6
Understanding Debt	Examine the continuing consequences of earlier interventions made without enough understanding, and address the work they have left unresolved.	Chapter 7

The Sufficient Understanding Test

These are the five questions introduced in Chapter 1 and developed in Chapter 6. I use them to organize a decision meeting, not to score one.

1. Have we examined the evidence, context, expertise, and lived experience, and identified where our understanding is weakest?
2. Is the depth of our evaluation proportionate to the possible consequences?
3. Have we named the uncertainties that matter to the decision?
4. Have the people who will live with the decision had a meaningful opportunity to inform it?
5. Can we describe a plausible failure and show how we have investigated it?

An incomplete answer should lead to a specific response: a missing perspective to hear, a more suitable evaluation, additional resources, or a narrower proposal. Sometimes the answer is to stop.

WORKSHEET

From Chapter 4

Context Cycle Record

DECISION _____

OWNER _____

DATE _____

SCOPE _____

1 Observe

What was seen, heard, or measured; when and where; whose experience is represented.

2 Context

The relevant history, people, constraints, responsibilities, and connections to other work.

3 Interpret

The likely explanation, a plausible alternative where relevant, and the assumptions still needing examination.

4 Decide

The choice, alternatives considered, remaining uncertainties, and why the next step fits the possible consequences.

5 Act

The scope, owner, preparation, measures, support, and conditions for pausing or changing course.

6 Reflect/Refine

The review date; expected versus observed results; what changed in your understanding; the next decision and its owner.

WORKSHEET

From Chapters 1 and 6

Sufficient Understanding Test

DECISION _____ DATE _____

1 Lens coverage

What did we learn from the evidence, the local context, relevant experts, and the people affected?
Where is our understanding weakest, and could that gap change the decision?

2 Proportional depth

What have we done to understand the cost of being wrong, the difficulty of reversing the harm, and the scope of human impact? Have we checked that the safeguards we're counting on will work?

3 Important uncertainties

Which questions need an answer before we proceed, and which can we investigate during a limited evaluation? Who is following each one?

4 Meaningful input

What did the people affected tell us, and what did their contribution change? Whose experience is still missing?

5 Investigated failure

What plausible failure have we examined? What would prevent or detect it, and what would happen if those measures fell short?

WORKSHEET

From Chapter 6

Decision Record

DECISION _____ OWNER _____

DATE _____ REVIEW DATE _____

1 Proposed next action

What is permitted, for whom, with which information, and within what boundaries.

2 Stakes

Cost of being wrong, difficulty of reversing harm, and scope of human impact; reasoning for each.

3 Readiness

What the five questions establish, the most important remaining gap, and how it will be handled.

4 Cost of delay

What waiting would affect and what further investigation is expected to resolve.

5 Decision and responsibility

Proceed, investigate, revise, or decline; the accountable owner and conditions.

6 Review and response

What will be monitored, when the decision returns for review, and what would require an earlier pause.

WORKSHEET

From Chapter 7

Understanding Debt Record

PROCESS _____

OWNER _____

DATE OPENED _____

REVIEW DATE _____

1 Earlier decision

The change introduced, its purpose, and the understanding it required.

2 Unresolved question

What was missed or left unclear; the evidence for that account and possible alternative explanations.

3 Continuing consequences

The extra work or other effects still occurring, including who experiences them.

4 Effect on the next proposal

What depends on the unresolved process and whether it changes the proposed scope or timing.

5 Repair

The specific change, responsible person, resources, and transition arrangements.

6 Follow-up

How we will judge whether the correction helped, when we will review it, and how concerns can prompt an earlier response.

B

The Four Lenses and Context Checklist

Chapter 5 introduced four lenses for developing understanding: Evidence, Context, Expertise, and Lived Experience. In practice they overlap. A nurse describing a shift may contribute all four in one conversation, and we don't need to stop that conversation to sort it into categories. The questions below are here to help you notice which perspective is thin before the decision is made.

Choose the questions that fit the decision. A small practice might work through them in a few conversations with staff, patients, and an outside specialist. A health system considering a broad deployment may need several teams and a more formal review. Either way, the purpose is to find what could change the choice or the implementation.

Evidence: What Supports Our Understanding?

- What was tested, on which people and settings, and with which version of the product?
- Does the evidence address the benefit we're counting on, or a related task? A tool can predict something accurately without helping anyone act on the prediction.
- Which findings make the proposal less convincing? Has someone looked for the conditions under which the tool performs poorly?
- How were our own records produced? Is there a field we're about to rely on that needs explanation first?
- What important part of the expected benefit hasn't been demonstrated?

Context: What Surrounds the Decision?

- What are the staffing, resources, and working conditions this use would depend on?
- What has changed recently, and what happened the last time a similar change was introduced?
- Who maintains the information, the process, or the system the tool would rely on?
- Which interests surround the proposal? A vendor, a sponsor with a launch date, or a department under pressure to improve a measure may each draw attention to some questions more than others.
- What would need to change here for the proposal to work?

Expertise: Who Can Help Us Interpret What We Find?

- Whose knowledge does this decision need: clinical, technical, operational, analytical, or familiarity with a particular population?
- Where does each person's expertise end, and who else should be involved?
- What reasoning sits behind each judgment? What would change the recommendation?
- Are the experts involved while the plan can still change, or only reviewing a nearly finished implementation?

Lived Experience: What Is This Like for the People Affected?

These are the questions from Chapter 5 that I would begin with. Ask them about the work before asking about the tool.

- What are you trying to accomplish?
- Where does the process make that difficult?
- What happens when the usual approach doesn't work?
- What would you want us to understand before changing it?

Then consider whose experience you haven't heard. The staff member available for a midday meeting may work under different conditions from a colleague covering evenings. A patient comfortable with an online survey may face different barriers from someone who doesn't use the portal. Plan how you will tell people what their input changed.

When the Lenses Don't Agree

A report may show improvement while staff say the work has become harder. Before accepting either account, find out what each one measures. Look at the relevant records, observe the steps, and ask people to explain specific examples. Bringing the lenses together doesn't mean giving every statement equal weight, and it doesn't

require everyone to agree. It means understanding the disagreement well enough to explain the decision and what would cause you to reconsider it.

WORKSHEET

From Chapter 5

Four Lenses Record

DECISION _____

DATE _____

1 Evidence

Relevant findings, their sources, and their limits. What important part of our expected benefit has not been demonstrated?

2 Context

Conditions, resources, history, and constraints that affect this use. What would need to change here for the proposal to work?

3 Expertise

Relevant judgments, the reasoning behind them, and unresolved disagreements. What requires knowledge we haven't yet involved?

4 Lived Experience

What affected people describe and whose perspectives are missing. What problem do they think needs to be solved?

5 What changes the decision

The weakest part of our current understanding, who will investigate it, and when we will revisit it.

C

AI Vendor Question Bank

These are questions I would bring to a conversation about a new tool, an AI feature in existing software, a renewal, or a change in how we use a product. You won't need all forty for every proposal. Choose the ones that help your team understand the decision and its possible consequences. Some answers will need to come from your own organization.

Keep a record of the answers, including the date, product version or configuration, and supporting documents. If something remains unanswered, name who will follow up. Be clear about whether you're discussing a feature available today or one still being developed, and whether a claimed benefit has been measured or is an estimate. Where you need a commitment from the vendor, check how it appears in the agreement. Answering the questions helps us make a decision; it doesn't give the tool approval to be used.

What Are We Considering?

- ☐ 1. What problem is this product designed to help with? What uses are outside its scope?
- ☐ 2. What does it produce or predict, and what are we expected to do with the result?
- ☐ 3. Which features would we get now? Which require another purchase, work by our team, or a future release?
- ☐ 4. What can it read, change, send, or set in motion on its own? Where does a person have to approve an action?

- ☐ 5. What would it need from our systems, information, and staff to work as intended?

What Supports the Claims?

- ☐ 6. Can we review the evidence for each claimed benefit, including the methods, limitations, and product version tested?
- ☐ 7. Who did the evaluation and paid for it? What part did the vendor play in the study and its reporting?
- ☐ 8. What was the product compared with, and what improved? Was that accuracy, time, cost, patient outcomes, or another result?
- ☐ 9. Which people and settings were included? Where has performance been weaker or left untested, including across languages and working conditions?
- ☐ 10. What errors do you know about, and how were they found? How would our staff recognize and respond to them?

Would It Fit Our Work?

- ☐ 11. How does our intended use differ from the uses already evaluated? What do we still need to investigate?
- ☐ 12. What could we evaluate before committing, and what would that require from each team? Include the information we'd use, permission to use it, and the time, expertise, and cost involved.
- ☐ 13. Can you show us an ordinary task and one where something goes wrong? We'd want to see what happens with missing or conflicting information, using examples we can examine safely.
- ☐ 14. What would we need to fix or prepare first, including data, connections, or parts of our current process?
- ☐ 15. What would the evaluation need to show before we went further? What findings would make us stop or change the plan?

How Would We Introduce It?

- ☐ 16. What will your team do during implementation, and what will ours do? Who is responsible for each part?
- ☐ 17. How will staff help shape the plan? Where patients, caregivers, or community members are affected, how will we hear from them while changes are still possible?
- ☐ 18. What training and practice will people need for their different responsibilities? How will we prepare someone who joins after launch?
- ☐ 19. Who will check outputs and handle exceptions? What knowledge, information, and time will they need?
- ☐ 20. How will we explain the change and respond to people's concerns? What will staff do when the tool is unavailable or can't help with a task?

What Happens to Our Information?

- ☐ 21. What information comes in and goes out? Where is it processed and stored, and who can access it, including other providers you use?
- ☐ 22. How do we control and check people's access and the product's connections to our systems?
- ☐ 23. How long do you keep the information, and what can you use it for? Can it be used to improve products or models, and what choices do we have?
- ☐ 24. What can you provide for our privacy and security review, including agreements and evidence of how the protections work?
- ☐ 25. If information is exposed or accessed inappropriately, how would we be told? How would our teams investigate and address it together?

What Would It Cost, and What Have Others Learned?

- ☐ 26. What would we pay at the start and over time? What could change the bill, including usage, minimum commitments, and renewal terms?
- ☐ 27. What costs would fall to us outside your price? Include connecting systems, preparing staff, reviewing outputs, monitoring, support, and moving from our current process.
- ☐ 28. Which benefits have been measured and which are estimates? What would have to happen here for us to see the expected savings or improvements?
- ☐ 29. Can we speak with customers doing similar work with this version, including people who use it daily and teams that had difficulties?
- ☐ 30. What did those customers have to change? What work remained, and how did they resolve implementation or support problems?

What Happens After Launch?

- ☐ 31. What information will you give us to judge how the tool is working, including errors and activity records? What will we have to collect ourselves?
- ☐ 32. Who do we contact for help, and how quickly should we expect a response? What happens with an urgent problem or one outside support hours, and does any of that cost extra?
- ☐ 33. Which product or service changes will you tell us about, and how much notice will we have? What evaluation will you share for changes to a model or feature?
- ☐ 34. Can we evaluate an update before using it, delay it, limit its use, or return to the earlier version? Where do those options have limits?
- ☐ 35. Who can pause the tool, and what happens to work already underway? What would need to be checked before restarting?

What Are We Agreeing To?

- ☐ **36.** Which promises about performance, support, and our information will be in the agreement? What limits or exclusions apply?
- ☐ **37.** If the product doesn't meet the agreement, what can we do? How do the renewal and termination terms affect our options?
- ☐ **38.** If we leave, what data, outputs, records, and settings can we take with us? In what format, how soon, and at what cost?
- ☐ **39.** What would you delete, what would you keep, and how would we confirm that? How does the agreement address information already used to develop or improve models?
- ☐ **40.** What would each team need to do so we could continue the service while moving away from the product? Include access to records, staff preparation, and unfinished work.

Take unanswered questions back to the team making the decision. If an answer could change whether or how you proceed, explain that before the next step. Assign the follow-up to someone who can resolve it, bringing in the specialists the question requires.

D

Governance and Implementation Templates

These templates support the governance, evaluation, and implementation work described in Chapters 8 through 13. They are written so a small practice and a large health system can both use them. A smaller organization might complete them in a conversation between a manager, a clinician, and an outside IT provider. A larger one might route them through existing quality, technology, and service review. The arrangements should fit the organization and the consequences of the use.

Use what you already have where it does the work. If your organization has an intake form, a project charter, or an improvement template, add the missing questions to it rather than creating a second process.

Responsibilities to Write Down

Before the first request arrives, write a short description of who can make each kind of decision. Chapter 8 develops this in more detail. At a minimum, I would want the following to be clear:

- Who can approve an evaluation, authorize routine use, change the permitted scope, or pause a tool?
- Where do unresolved concerns go, and who has authority to resolve competing priorities?
- Who is the leader accountable for each initiative's intended result?

- Who manages access and technical support, and who owns how the tool is used within the service?
- Who maintains the inventory of AI uses, and how does it stay current when software changes?

Review Paths

Chapter 8 describes organizing review around the proposed use. The names and delegation rules can fit your organization.

- **A limited test** may qualify for a delegated review when its purpose, information, and consequences are well contained.
- **A proposed use in routine work** needs review of the actual workflow, including the information used, who checks the output, and how a correction would be handled.
- **A use with more consequential effects** requires deeper investigation and the relevant specialist involvement.

These paths organize the work. They don't establish universal risk categories or guarantee approval. A small pilot involving real people can still carry significant risk.

WORKSHEET

From Chapter 8

AI Use Inventory Entry

USE OR PRODUCT _____

STATUS _____

RESPONSIBLE LEADER _____

LAST UPDATED _____

1 Purpose

What problem the use addresses and what people are actually doing with the tool.

2 Information involved

What the tool receives, where it comes from, and where its output goes.

3 What the output influences

The decisions, work, or services it can affect, with or without a person's approval.

4 Version and configuration

The product version, settings, and connections reviewed, where available.

5 Approval and conditions

What has been approved, by whom, and the conditions attached. Recording a use here does not approve it.

6 Unknown details

What we still need to find out and who is responsible for finding it.

WORKSHEET

From Chapter 2

Describing the Proposed Use

PROPOSAL _____ PREPARED BY _____

DATE _____

1 Purpose

What problem are we trying to address, and what would a better result look like?

2 Information

What does the system use, where does it come from, and who maintains it?

3 Output and action

What does it produce, and what can happen next, with or without a person's approval?

4 People and work

Whose responsibilities, workload, or experience could change, including people who never use the software?

5 Preparation and oversight

What practice, review, support, and escalation will be needed, and who will provide them?

6 Evidence and follow-up

What supports this use, what remains uncertain, and how will we evaluate the result and respond to problems?

WORKSHEET

From Chapter 8

Intake Questions

REQUEST _____ REQUESTER _____

COORDINATOR _____ DATE RECEIVED _____

1 What problem are we trying to address?

The work that needs improvement, the people affected, and the result sought.

2 What are we asking permission to do?

The specific evaluation or use, including what the tool may produce or do automatically.

3 What information will be involved?

Its sources, sensitivity, access, destinations, and any unresolved questions about handling it.

4 Who is responsible?

The service owner, proposed users, technical support, and people needed for review.

5 What supports the proposal?

Relevant evidence, local conditions, input from affected people, and uncertainties.

6 What could happen if we are wrong?

Cost of error, difficulty of reversing harm, and the number of people affected and depth of impact.

7 What will implementation require?

Time, funding, training, review, support, and any earlier problems the proposal depends on.

8 How will we judge the result?

Expected benefit, checks for problems, follow-up, and arrangements for pausing or stopping.

WORKSHEET

From Chapters 8 and 10

Approval and Decision Record

USE _____

DECIDED BY _____

DATE _____

REVIEW DATE _____

1 Intended use

The task, people, information, product version, and limits of the use reviewed.

2 Basis for the decision

Relevant evidence, local evaluation, staff and patient input where appropriate, and remaining uncertainty.

3 Applicable requirements

What reviewers determined applies, the source and date, and any unresolved question.

4 Authorization and preparation

Who decided, what they permitted, and how required preparation was confirmed.

5 Responsibility during use

Who checks performance, who responds to concerns, and where the results are kept.

6 Changes and follow-up

What changed, what the team found, what it did, and who checked the response.

WORKSHEET

From Chapter 13

Experiment Plan

EXPERIMENT _____ OWNER _____

START DATE _____ REVIEW DATE _____

1 Question

What are we trying to learn, and which decision will the answer inform?

2 Boundaries

What information, people, and systems are involved? How will we prevent use outside the test?

3 Preparation

What approvals, expertise, practice, and support are needed before starting?

4 Evaluation

What will we compare, measure, and observe? What would justify continuing, changing, or stopping?

5 Pause arrangements

Which findings require an early response? Who can pause the test, and how will necessary work continue?

6 Review

Who will consider the findings, when, and what must be resolved before restarting or extending?

7 Next stage

What would change with wider use, and what would need further review?

WORKSHEET

From Chapters 11 and 12

Checking Our Commitments

IMPLEMENTATION _____ DATE _____

Reducing work

Have we accounted for review, correction, and work transferred to other people?

Supporting a better conversation

What do clinicians, patients, and caregivers say about the encounter?

Providing preparation and support

Have people had practice time, and can they reach help when they need it?

Taking concerns seriously

Are reports assessed, acted on when needed, and followed up with the people who raised them?

Explaining changes to roles or services

Do people know what is decided, what remains open, and when they will hear more?

Correcting a decision

Can a person question an outcome, reach someone responsible, and learn what was reviewed and changed?

Walking Through One Decision

Chapter 10 suggests choosing one consequential AI use and reviewing it with the person responsible. These are the questions I would want that conversation to answer:

- What are we using this tool for, and does that match the use that was reviewed?
- Who authorized it, what informed the decision, and were the required preparations completed?
- Which requirements apply, and who is responsible for keeping that assessment current?

- What have we learned during use, and what happened when someone raised a concern?
- If we needed to restrict or stop it today, who would act and how would the work continue?

If an answer takes some investigation, find out why. The information may exist in another system, a responsibility may be unclear, or the work may not have been completed. Those are different problems and need different responses.

A Ninety-Day Starting Sequence

Chapter 14 suggested using the next ninety days to put this approach into practice. This is a fuller version of that sequence. Some organizations already have much of the structure in place. Others may need outside help or more time. The dates give the work review points; they don't confer permission to launch or require work to wait.

An urgent concern needs attention when it is found. A properly reviewed activity may be ready to proceed while the wider process is still being improved. The purpose is to make progress on actual needs, with enough preparation for the consequences involved.

Days 1–30: Understand the Current Position

Begin by finding out which AI uses are proposed or active and who is responsible for them. The inventory entry above can help organize that information. Include features in existing software and uses staff have begun to explore, as well as formal purchases.

For each known use, establish what it does, which information it involves, and what its output can influence. Record uncertainty honestly and assign someone to resolve it. An inventory entry records a use; it doesn't approve it.

Alongside the inventory, talk with staff about the problems they are trying to solve. A request for an AI tool may point to a process that needs attention regardless of which product is selected. That conversation can also help identify work that has moved between teams or support that people were promised but never received.

Choose a decision for a closer review based on the concerns and opportunities you find. It might be an existing use with consequential unanswered questions or a proposal where a limited evaluation could provide useful evidence. The choice should reflect your setting, rather than a requirement to start with the easiest demonstration.

By the first review point, I would want an initial account of known uses, responsible people, significant gaps, and the next actions. It may be incomplete. The organization should still be responding to identified concerns while the work continues.

Days 31–60: Resolve the Responsibilities and Missing Work

Use what you learned to make the review process easier to follow. People need to know where to bring a proposal or concern, who coordinates the response, and who can make the relevant decision. If those arrangements already exist, examine whether they work for the AI uses you have identified.

For the selected decision, bring together the evidence, local context, expertise, and affected people's experience. Identify what still needs investigation. The result may be a clearer plan, a narrower proposal, a decision to stop, or a specific question for an experiment.

Implementation preparation belongs here too. Who needs practice with the tool? Who will review its output? What will patients or staff be told? Who will help when the usual process cannot be followed? Put the required time and support into the plan, with someone responsible for providing them.

In a smaller service, this may involve a manager, staff members, and outside technical or specialist support. A larger organization may assign the work across several teams. A whole-organization approach means connecting the responsibilities and consequences relevant to the use. It doesn't require every department to join every meeting.

Give attention to unresolved problems that would undermine the proposal. If the tool depends on current policies and nobody maintains those policies, that work needs an owner. If staff are already correcting output without enough time, the immediate plan needs to address their capacity. Adding another form would not resolve either problem.

At the next review point, the team should be able to explain the decision, the conditions attached to it, and what happens next. If important work remains, name it and explain how it affects the permitted use.

Days 61–90: Carry Out the Next Step and Review It

Proceed with the next step the review supports. That could be a bounded experiment, a change to an existing implementation, further investigation, or closing a proposal that cannot meet the need. The purpose of this period is to follow through on the decision.

If a test is appropriate, use the approach from Chapter 13: a clear question, defined boundaries, an evaluation plan, and arrangements for pausing or stopping. An existing service may instead need corrected responsibilities, better preparation, or a more workable route for concerns.

Review what happened with the people involved. Did the proposed preparation occur? Was the support available? What did staff and patients experience? What did the result establish, and which questions remain?

The review should also examine the process you used. Perhaps a request reached the wrong person, an approval condition was unclear, or a staff member noticed something the evaluation missed. Make the corresponding changes before repeating the same arrangement for another initiative.

By the end of this starting period, I would want to see decisions followed by action and review, with a record that another person can understand. The number of tools launched is not enough to judge that work. Neither is the number of proposals stopped. We need to look at whether the decisions were supported and whether the organization carried out its responsibilities.

E

Glossary

These are the terms used in this book, described the way the chapters use them. Where a term has a broader technical meaning, I've kept to what a leader needs in order to follow the decision. The chapter in parentheses is where the term is developed.

Agent. A product description suggesting that software can take actions, such as reading or changing records, sending messages, or triggering another process. Ask the vendor to demonstrate what those actions are and where a person's approval is required. (Chapter 2)

AI governance. The responsibilities and processes an organization uses to decide how AI may be used, support its implementation, and respond to what happens during use. (Chapter 8)

AI use inventory. A maintained record of AI uses in an organization, including purchased products, internally developed tools, capabilities embedded in existing software, and uses staff have begun on their own. Recording a use in the inventory does not approve it. (Chapter 8)

Business associate. Under HIPAA, a party that performs certain functions or services for a covered entity involving protected health information. Qualified reviewers should determine how the rules apply to a particular arrangement. (Chapter 10)

Confabulation. The term NIST's Generative AI Profile uses for generative AI producing inaccurate material in confident, readable language, also commonly called hallucination. (Chapter 2)

Context Cycle. The working sequence this book uses to apply the Context Principle: Observe, Context, Interpret, Decide, Act, and Reflect/Refine, then observe again. A team can return to an earlier step when new information changes its understanding. (Chapter 4)

Context Principle. My framework for decisions: understanding must precede intervention, in proportion to the potential impact of the decision. This book applies it to AI in healthcare. (Chapters 1 and 4)

Covered entity. Under HIPAA, a health plan, a healthcare clearinghouse, or a provider conducting covered electronic transactions. (Chapter 10)

Decision record. A short written account of a decision: what is permitted, the stakes, what the team established, the cost of delay, who is responsible, and when the decision returns for review. (Chapters 1 and 6)

Decision Triangle. Three considerations for judging how much understanding a decision needs: the cost of being wrong, the difficulty of reversing the harm, and the scope of human impact. A serious concern in one area deserves attention even when the other two appear manageable. (Chapter 6)

Drift. Changes in a model's input data, or in the relationships a model relies on, after it was developed or evaluated. Drift can affect performance, but it does not necessarily produce a slow, predictable decline. (Chapter 2)

External validation. Evaluating a prediction model's performance on new participant data, which might come from a later period, another organization, or a different population. Evidence of predictive performance does not, by itself, establish that using the tool improves care. (Chapters 2 and 9)

Fine-tuning. Additional training that adapts an existing model. Giving a tool a document to consult is not, by itself, the same as fine-tuning it. (Chapter 2)

Four Lenses. Four perspectives for developing understanding: Evidence, Context, Expertise, and Lived Experience. They overlap in practice. (Chapter 5)

Generative model. A model that produces content, such as text, images, or audio. (Chapter 2)

Illusion of understanding. The risk that a well-written answer is mistaken for evidence that the underlying problem has been understood. (Chapter 1)

Inference. Using a model to produce an output, as distinct from training it. Using a model does not necessarily retrain it. (Chapter 2)

Intervention. A change we introduce: a tool, a policy, a revised process, or a different assignment of responsibility. An AI feature added to existing software is an intervention even when nobody requested it and no purchase is involved. (Chapter 4)

Large language model (LLM). A type of model used for language tasks, including generating text. (Chapter 2)

Model and product. The model is the underlying system that produces an output. The product is what your organization uses, which may combine a model with instructions, document retrieval, access controls, and connections to other applications. Evidence about the model alone won't tell you how the whole product performs. (Chapters 2 and 9)

Pilot. A label for a limited trial. Calling something a pilot doesn't make its consequences small; the information involved, who sees the output, and what someone might do because of it determine the risk. (Chapters 3 and 13)

Plan-Do-Study-Act (PDSA). An improvement method the Institute for Healthcare Improvement describes for planning a test, carrying it out, studying the results, and using the learning to guide the next test. (Chapters 4 and 13)

Predictive model. A model that estimates an outcome or assigns a score from the information available to it. (Chapter 2)

Premortem. A discussion in which a team imagines that a project has failed, identifies possible reasons, and plans how to address them. (Chapter 6)

Proportional Understanding. Matching our preparation to the possible consequences of a decision. (Chapter 6)

Proxy. A measure used in place of the outcome we actually care about, such as healthcare spending used in place of health need. A proxy can be useful for one decision and unsuitable for another. (Chapters 1 and 11)

Psychological safety. In Amy Edmondson's research, whether team members feel able to take interpersonal risks, such as raising a question or admitting a mistake. (Chapter 12)

Retrieval-augmented generation (RAG). An approach that combines retrieval of source material with generation of an answer. A product might retrieve passages from an organization's policies and make them available to the language model as it responds. (Chapter 2)

Review path. A route through governance matched to the proposed use, such as a delegated review for a well-contained test or deeper specialist review for a use with more consequential effects. (Chapter 8)

Sufficient Understanding Test. Five questions for judging whether understanding is sufficient for a proposed action: lens coverage and the weakest area, proportional depth, the uncertainties that matter, meaningful input from affected people, and an investigated plausible failure. Completing it does not authorize a use. (Chapters 1 and 6)

Training. Adjusting a model's internal parameters using data. (Chapter 2)

Understanding. Making sense of information in the setting where we intend to act, including how the work actually happens and what people are expected to do with the result. (Chapter 4)

Understanding Debt. The costs that can accumulate when changes are repeatedly made without enough understanding for the stakes. Later, someone has to resolve those questions while also doing the work. (Chapter 7)

Whole-organization approach. Following the consequences of a change far enough to understand who is affected and what they need. It scales to the decision; it doesn't require every department to approve every experiment. (Chapter 2)

Workaround. An unofficial practice, such as a separate tracker or a sequence of phone calls, that compensates for something the formal process doesn't provide. Understanding what it does comes before removing it. (Chapter 7)

F

Discussion Guide

These questions are for leadership teams, governance groups, and classes working through the book together. They work best when a group applies them to an actual decision in front of it, rather than answering them in the abstract. You won't need all of them. Choose the ones that fit your setting and give people time to answer from their own experience.

A few suggestions for the conversation. Invite someone who does the work affected by AI, not only people who make decisions about it. Let disagreement surface; it often shows where the understanding is thin. And if a question leads to something your organization needs to act on, assign it before the meeting ends.

Chapter 1: The Wrong Question

1. Think of the last AI decision your organization made. Did the discussion begin with the problem or with the product?
2. Where might an AI feature already have arrived in software you use, without a formal decision?
3. When your team says a person will review the output, what would that person need in order to catch an error?

Chapter 2: What Leaders Actually Need to Know About AI

1. Choose one AI proposal. Can you describe what it produces and what can happen next, in terms a colleague outside the project would understand?

2. Who in your organization would notice first if a tool's recommendations stopped fitting the work?
3. What preparation would the people carrying out the change need, beyond learning to use the software?

Chapter 3: Hype, Fear, and the Discernment Gap

1. Where have you seen a change make one step faster while moving work somewhere else?
2. Think of someone who is enthusiastic about AI and someone who is cautious. What is each of them trying to tell you?
3. Which claims about AI in your setting have evidence behind them, and which are still expectations?

Chapter 4: The Context Principle and the Context Cycle

1. What does “seeing the room” mean in your part of healthcare? What surrounds the task that people can miss?
2. For a current problem, what have you observed, and what have you already started to explain? Are those kept separate?
3. Where does your organization already record observations, assumptions, and results? How could the Context Cycle fit into that work?

Chapter 5: The Four Lenses of Understanding

1. For a current decision, which of the four lenses is thinnest?
2. Whose experience does your organization tend not to hear, because of when or how it asks?
3. When a report and a staff account disagree, how does your team usually resolve it?

Chapter 6: How Much Understanding Is Enough?

1. Describe a decision where your team was stuck asking for more information. What would have changed the decision?
2. Choose a proposed use. What is the cost of being wrong, how hard would the harm be to reverse, and how many people would it affect?
3. What does waiting cost in your setting, and who is carrying that cost now?

Chapter 7: Understanding Debt

1. Where are people in your organization keeping a separate record or making extra calls to make a process work?
2. What would be missed if they stopped?

3. Which earlier change is a new proposal likely to inherit, and who would have to resolve it?

Chapter 8: Build Governance Before You Need It

1. If a staff member wanted to try an AI tool tomorrow, would they know where to ask and who could decide?
2. Who can pause a tool in your organization, and how would the work continue while it is paused?
3. What would make your review process easier for a requester to use without lowering its standards?

Chapter 9: Evaluating AI Tools and the Vendors Behind Them

1. What examples from your own work would you bring to a vendor demonstration, including one that doesn't go according to plan?
2. Whose time would the tool save, and whose time would it require?
3. If you needed to leave the product in a year, what would you need to take with you?

Chapter 10: Risk, Regulation, and the Audit-Ready Organization

1. Choose one AI use. Could a new leader understand why it was approved without asking the person who approved it?
2. Which approval conditions in your organization are easiest to lose track of after launch?
3. When staff find that a condition isn't working, where does that information go?

Chapter 11: Ethics as Infrastructure

1. For a prediction your organization uses or is considering, what will people actually do with it?
2. Could staff explain to a patient how AI is being used in their care and whom to contact with a concern?
3. If someone declines a tool, what happens next, and is that alternative workable?

Chapter 12: Trust Is the Deployment Environment

1. What have staff learned from earlier technology changes, and how will a new proposal account for it?
2. Which commitments has your organization made about AI, and how would you check whether they are being kept?

3. How would you know if a tool was affecting the conversation between a clinician and a patient?

Chapter 13: Innovation Without Recklessness

1. What is the smallest test that could answer a question your team has about AI?
2. Before starting, what finding would cause you to pause?
3. Think of a project that continued because of the effort already spent. What would an honest review have asked?

Chapter 14: The Enduring Leader

1. Which AI decision in front of you deserves closer attention in the next ninety days?
2. How will you stay close enough to hear what the decision meant for the people living with it?
3. Who in your organization could continue this work if you moved to another role, and what would they need from you?