

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?