

The River Doesn't Wait

A Senior Executive Guide To
Navigating the Surge of Agentic AI
And Pulling Away From Your Competitors

By Blaine Mathieu

Reading the Water Together

A Discussion and Reflection Guide for The River Doesn't Wait

A note on using this guide

This guide is built for a room of senior leaders, each of whom owns a boat of some size and is trying to navigate it down the same river. It is not a quiz on the book, and it is not designed to land everyone in the same place. The book refuses to tell you the right answer on purpose, because there isn't one that holds still. A downloadable version of this guide can be found at the book's website: <https://riverdoesntwait.com/>

A few things to set up before you start:

- **Two kinds of questions.** *Reflect* prompts are for thinking about on your own before the meeting. Some of them are personal, and they are easier to answer honestly in your own head than out loud. *Discuss* prompts are for the room, and the good ones are built to split it. Having an outside facilitator might help ensure that the strongest voices in the room don't dominate.
- **A menu, not a checklist.** Pick three or four questions per session that fit your group. Trying to get through all of them is the surest way to have a shallow conversation about all of them.
- **It ends in work, not talk.** The final session is a working session, where each of you applies the framework to an outcome you actually own. That is where a book club turns into the thing the book is asking you to do.

You can run this in three sessions, grouped as below, or chapter by chapter if your group prefers seven shorter sittings. The three-session shape is the one I'd recommend for a busy team.

Session 1 - The Setup

Foreword and Chapters 1 through 3: the curve, the choice, and the river

These chapters establish the world before any decision gets made. The Foreword names the underlying problem (the gap between what your systems are supposed to do and what they actually do, and why speed turns a tolerable gap into a structural risk). Chapter 1 makes the case that the agentic AI capability curve is real and accelerating, and that the leaders pulling ahead are conservative institutions, not just digital natives. Chapter 2 says every enterprise is already making a retrofit-or-reimagine choice, deliberately or by default, and that neither end of the spectrum is safe. Chapter 3 reframes the choice as something you navigate continuously rather than select once.

Reflect (before the session)

- Two years ago, were you asking *when* about AI? What are you asking now? The book says the market's question has shifted from "when" to "why isn't this working yet." Has yours?

- The Foreword says “you cannot delegate this.” Be honest about what you have, in practice, delegated. What AI decisions in your area are currently being made by your CIO, a vendor, an analyst firm, or a task force, that the book would say are actually yours?
- Chapter 2 argues retrofit failure is far more common than reimagine failure, but reimagine failure (Klarna) is the one that lives in executives’ heads. Which fear has actually been driving your decisions: moving too fast, or moving too slow?

Discuss (in the room)

- The book claims most enterprises are “doing this wrong, in the same wrong way, for the same wrong reasons,” and that the cost of inaction was negligible in 2024 but shows up on quarterly calls by 2027. Do you buy that timeline for *our* industry, or is it faster, slower, or wrong? Where does this room disagree?
- Pick a real, current AI initiative in our organization. Is it honestly a retrofit or a reimagine? The book warns that “you can spend a billion dollars on AI without reimagining anything.” Would everyone in this room classify our initiative the same way?
- Chapter 3 says the leaders aren’t ahead because they picked a better spot, but because they noticed faster when the spot moved. Do we have any mechanism today for noticing when our position has gone stale, or do we only reassess on the annual planning calendar?
- Dimon’s line is “if you put your head in the sand, you will lose.” The book’s sharper question is whether your nearest competitor’s next strategy deck still looks like yours. Name the competitor whose AI-related move would actually scare you. Are we watching them, or just watching the technology?

Anchor case. JPMorgan is the book’s load-bearing example precisely because it is huge, regulated, and risk-averse, the opposite of a YOLO startup. Does that make the example more relevant to us, or easier to dismiss as “they have a \$19.8 billion budget and we don’t”? What is our honest version of “we are not Klarna”?

Session 2 - The Three Tensions

Chapters 4 through 6: workflow and workforce, context, and governance

Inside the boat, every piece of work is governed by three linked settings. Tension 1 (Chapter 4) is workflow and workforce, and the claim is that they are one decision most enterprises wrongly treat as two. Tension 2 (Chapter 5) is context, and the claim is that the scope of context you can set equals the scope of organizational reach you actually hold. Tension 3 (Chapter 6) is supervision, autonomy, and governance, and the claim is that you either build the rules into how the work runs or you bolt a lookout on afterward. The book’s foundational argument is that these three move together, not one at a time.

Reflect (before the session)

- In your own area, are workflow and workforce one decision or two? Who actually owns the combined call, given that Ops, HR, and IT usually each own a piece on a different schedule?
- Chapter 5 says the hard part of context is making explicit what you have kept implicit in your own head for years. What is one thing you “just know” about how your function really works that you have never written down or said out loud?

- Chapter 6 says if your firm ends up bolt-on, it will be because the default carried you there. Which way is the default currently carrying your area?

Discuss (in the room)

- Chapter 4's bluntest claim: slotting agents into existing roles is "the slow path to irrelevance." Is that true everywhere, or are there functions in our organization where putting agents into existing seats is exactly the right and sufficient move? Where do we disagree?
- Chapter 4 names two shapes of reimagine: collapsing the org chart (Block) and rewiring the work inside the boxes (JPMorgan's KYC). If we reimaged one thing, which shape would it be, and what layer in *our* organization gets squeezed as a result? Can we say that honestly in this room?
- Chapter 5's uncomfortable claim: providing rich context to a team means giving up your exclusive hold on the knowledge your influence is built on, and that is uncomfortable in ways that have nothing to do with the technology. Is that discomfort real for people here? What context do you find yourself reluctant to make explicit, and why?
- Chapter 5's three levels of context: what you own, what sits inside the firm but outside your purview, and what sits outside the firm entirely. Which of our most important pieces of work actually require reaching into level two, and have we built the peer coalitions that takes? Or are we keeping work at level one because it is easier to control?
- Chapter 6 offers a tell: structural words (foundation, architecture, infrastructure) signal built-in governance; process words (council, framework, human in the loop) signal bolt-on. Listen to how *we* talk about AI governance. Which vocabulary do we use, and what does that reveal about which one we have actually built?
- The Air Canada, Microsoft Copilot, and Deloitte cases all failed not because the agent was too autonomous, but because "the thing watching the envelopes was the wrong shape." Where in our operation is the review machinery still built for the old shape of work while the work has quietly changed underneath it?

Anchor case. The book points out that the Big Four are selling agentic transformation while running it on themselves, so their case studies are "both descriptions and sales documents." When a vendor or advisor brings us their transformation story, how good are we at separating the description from the pitch?

Session 3 - The Practice and the Personal

Chapter 7 and the Epilogue: the plan you keep remaking, and you

Chapter 7 reframes everything: you are not managing AI, you are achieving outcomes you already own through a medium that now moves under you. The three Tensions collapse into one steering judgment (how far toward reimagine does this work need to go), outcomes get triaged by how fast the water is moving, and the plan is provisional by design. You sense at the speed of the river and commit at the speed your organization can actually absorb. The Epilogue puts the advisor's hat down and asks what all of this does to you, the person in the seat.

Reflect (before the session)

- Chapter 7 calls the snapshot mindset (deploy, then reassess on a schedule) “the one habit worth breaking on purpose.” Be honest: is your current AI planning a snapshot wearing work clothes? When did you last re-answer the watching questions for an outcome you own?
- The Epilogue’s central question: “Who am I, when the thing I was is done by something other than me?” The author says peace with it is still a way off for him. Where are you with that question?
- The author admits he catches himself wanting to automate things simply because he can, treating human friction as a defect to engineer out. Do you notice that pull in yourself?

Discuss (in the room)

- Chapter 7’s reframe: “you are not managing AI; you are achieving the outcomes you already own.” Do we talk about AI as an initiative (task force, center of excellence, portfolio of pilots) or as a faster way to move the numbers we are measured on? Which mode are we actually in, and what would it take to switch?
- The book says the failure mode of large enterprises is not that they change too slowly, but that they *sense* too slowly. Do we agree that is our potential failure mode? Where are we reacting to a river that moved two channels over while we weren’t looking?
- Chapter 7 insists the watch cannot be handed off to a market-intelligence team, a consultancy, an analyst subscription, or a research agent, because only the people steering an outcome can tell whether the water around it has changed. Are we already outsourcing our sensing, and to whom?
- The Epilogue argues that every decision about this technology is “a small moral choice in operational clothing”: which work goes to a machine, who you stand behind when a role dissolves, whether the people behind you inherit a ladder worth climbing. Is that the author overreaching, or is it the truest thing in the book? Where does the room split?

Bridge to the working session. The author says the single highest-value thing to do the week after finishing the book is to get the two or three executives who share your most important fast-water outcome into a room to read the water together. You are now in that room. Which shared outcome should this group take on first?

The Working Session

“Read the Water Together”

This is where the book club earns its keep. Run it as a working session, not a discussion. The goal is not consensus. It is to surface where this room genuinely disagrees about the water, because that disagreement is the most useful thing you will find.

Part A - on your own (before or at the start)

Each leader picks one outcome they personally own. Not “our AI strategy.” A real number on your scorecard: cycle time, cost to serve, retention, risk reduction. For that outcome, answer:

1. **The water.** Is this outcome in fast water or slow water right now, and what’s the evidence? Document-heavy work, coordination across many systems, high-volume exception handling,

and fragmented knowledge tend to sit in fast water. Work dominated by error costs, human trust, or the texture of a relationship tends to sit in slower water.

2. **The steering judgment.** How far toward reimagine does this work need to go this year?
3. **The three Tension settings that follow from that judgment.**
 - *Workflow and workforce:* agents into existing seats, or roles and work redrawn together? What gets squeezed?
 - *Context:* what does the team need to see that it can't today, and is that at level one, two, or three?
 - *Supervision, autonomy, governance:* how much runs without a human checking each step, and is the protection built into how the work runs or just inspected afterward?
4. **The help you need, and from whom.** From the CEO: air cover and the mandate to redesign rather than optimize. From the CIO or CTO: the platform and the built-in controls that let a team run wider without running reckless. From peer executives: the parts of the context that cross your boundary into theirs.
5. **The watch.** Which three questions will you keep re-answering for this outcome, and how often?
 - Has an AI-related capability crossed from demo to deployable against this actual work?
 - Has someone reset what “good” looks like for this outcome?
 - Has something you judged could wait started moving faster than you expected?

Part B - together (live)

Pick the one shared fast-water outcome that crosses two or three people in this room. Read the water together: where does it sit, how far toward reimagine does it go, where do the three Tension settings land, and what is the first move and who owns it? Spend most of your time on where you disagree about the water itself, not on the wording of the plan.

Before you close. This is not a one-time exercise. Put a date on the calendar to run it again, and remember the one rule that matters: the sensing has to outpace the moving. You watch at the speed of the river, and you commit at the speed the organization can absorb.