



# The Quiet Edge

How smart people are actually using AI at work

— THE MINDSET SHIFT

## CONTENTS

# What's in the full guide

Nine chapters. This primer is Chapter 1, plus one move borrowed from Chapter 5.

- 1 The mindset shift.** You're about to read it, starting on [page 3](#).
- 2 Your stack.** Pick roles, not products: why two paid tools is enough, and how to test them on three real tasks instead of chasing every launch.
- 3 Context.** The one-page Context Document you own, how to stop re-explaining yourself, and the mirror problem in full.
- 4 Letting it act.** From read-only to "it drafts, you send," and why most people should stop there.
- 5 Six jobs, six moves.** Writing, reading, thinking, meetings, learning, numbers. You'll meet the pre-mortem on [page 11](#). There are five more.
- 6 The quiet edge.** Using AI at work without costing you your credibility or your skill.
- 7 Your moves library.** A system that takes fifteen minutes a week, with a hard ceiling.
- 8 When to close the laptop.** Six times not to use AI, and why fluency isn't accuracy.
- 9 Your first seven days.** One task a day, each mapped to a chapter.

# The Mindset Shift

A little taste of the edge you'll access, a primer  
if you will, with one of the best moves in the guide

## Read this first

This is the opening chapter of **The Quiet Edge: How smart people are actually using AI at work**, expanded, for you, with one of our best moves from a later section of the book, so you can see how the rest of it works.

Three things I'm assuming about you:

- You already use AI a few times a week: ChatGPT for an email, Claude for a long document, Gemini because it came with your work account.
- You're busy.
- You're skeptical of hype.


Good. So am I.

This should be a quick read, with an action at the end that'll give you the best sense of whether the rest of the guide is for you.

---

## The Mindset Shift

A friend of mine (smart guy, senior at a consulting firm) was complaining last month that ChatGPT "isn't that useful." He showed me his chat history. It was a graveyard of one-shot prompts.




← one shot each

← then gone. a graveyard.




He was using a Ferrari to drive to the corner shop.




This is the most common failure mode I see, and it has nothing to do with skill or intelligence. It's a mental model problem. He was treating AI like Google: a thing you query, that gives you an answer, end of transaction. So that's exactly what he got: query-quality work.

It's not his fault. Twenty years of search trained all of us the same way: a few words in, an answer out, move on. That's the right habit for search. It's the wrong habit here, because when you type six words, the tool has to guess everything you left out: who's reading, what you want, what already happened, what "good" means to you. So it guesses the average. And you get average work.



The single most valuable shift you can make is this:

**Stop treating AI like a search engine. Start treating it like a junior colleague.**



Hold that frame in your head and most of what follows will feel obvious.


A junior colleague is fast, capable, and has read more than you. But they need context to be useful. You wouldn't walk up to a new hire and say "Write me an email about Q3" and expect anything good. You'd tell them who it's for, what tone, what happened last quarter, what you're actually trying to accomplish. Then they'd come back with a draft, you'd give notes, they'd revise.

That's how you get good work out of a junior colleague.

That's also how you get good work out of AI.


## What the difference looks like

Here's my friend's email, the search-engine way:



← no reader, no tone, no backstory

And here's the colleague way:



<sup>1</sup> who it's for  
<sup>2</sup> what happened  
<sup>3</sup> what you're actually trying to accomplish  
<sup>4</sup> what tone  
<sup>5</sup> the ask  
<sup>6</sup> the limit

The second one takes about a minute longer to write. It saves the ten minutes you'd have spent rewriting whatever the first one produced.

Notice what's *not* in the second version. No clever phrasing. No "act as a world-class copywriter." Nothing you'd find in a prompt pack. It's just what you'd tell a person. The valuable part of that request is the part only you know. That's exactly why prompt lists don't work, and why this guide doesn't have one.

---


# The four shifts

Once you make the colleague-not-search-engine move, four smaller shifts follow naturally. They're the spine of the whole guide.

## 1. From queries to conversations.

Most useful AI work happens in turn five, not turn one. The first response is a starting point: the *worst* version of what's possible. Push back. Ask follow-ups. The conversation is the product, not the first reply.

Here's what the next few turns on Dana's email might look like:



Three notes, under a minute each. The last one is the one most people never ask, and it's often the version you end up sending.

If you've ever closed the tab thinking "that's not quite it," that was the moment to keep going.




## 2. From outputs to iterations.


Stop trying to write the perfect prompt. Write a *good enough* one, see what comes back, refine. Iteration beats engineering.

The people who struggle most write a careful paragraph, get something 70% right, and decide the tool is mediocre. The people who do well write two lines, get something 50% right, and fix it in three quick passes. They end up further ahead with less effort, because reacting is easier than specifying. You often don't know what you want until you see what you don't.


Most useful notes are one of three kinds: *shorter*, *sharper*, or *try a different angle*. You don't need more vocabulary than that.



shorter



sharper




new angle

### 3. From prompts to context.


Prompt quality matters. Context matters more. A mediocre prompt with rich context (who you are, what you're working on, what you've already tried) beats a beautifully crafted prompt with no context every single time.

Try the same question two ways:



← no context in


You'll get a productivity listicle. Now:



1 who you are  
2 3 what you're  
working on  
4 what's going  
wrong

You'll get a plan for your actual week.

The obvious problem: nobody wants to type their life story into every chat. The full guide solves that with a one-page Context Document (who you are, how you work, what good looks like) kept in a plain file you own and reuse everywhere. That's Chapter 3, and it's the chapter that changes how your week feels. For now, just notice the gap between those two answers.




#### 4. From tasks to thinking partners.

The biggest shift of all. AI is most valuable not when it does a task for you, but when it helps you *think better* about a task. Pressure-test a decision. Argue against your own draft. Surface an angle you missed.

This is where the hours-saved math actually works, and where most people never go. They use AI *after* the thinking is done, to write up a decision they've already made. That's using a thinking partner as a typist.

The move is to bring it in while the decision is still wet.



Here's what that looks like in practice.

---

## The one move: the pre-mortem


In the full guide this sits in Chapter 5, one of six moves for six kinds of work. I've pulled it forward because it's the fourth shift in a single prompt, and because it's the fastest way I know to feel the difference between using AI and *working with* it.

THE PRE-MORTEM

PASTE AS IS

"It's six months from now and this decision turned out badly. Write me the story of how it went wrong, in detail, most likely failure first."

It's one of the best prompts in the guide.



### Why it works

Ask "what are the risks?" and you get a tidy list: budget, timeline, stakeholder buy-in. True of every decision ever made. Useful for none of them.

Ask for a *story* and the tool has to connect events: this happened, which caused that, which nobody noticed until too late. Specifics fall out of that chain. And the specifics are where your blind spot lives.

There's a human reason it works too. Hearing "here's how your plan failed" as a story about a hypothetical future is far easier than hearing it as criticism of a plan you're proud of. You stop defending and start reading.

The idea isn't mine. The psychologist Gary Klein popularised pre-mortems as a team exercise years ago. What's new is that you can now run one alone, in two minutes, on a Tuesday.

## A worked example

Say you manage a small team, and you've decided to replace the weekly status meeting with a written update. You're fairly sure it's right. The meeting has become a slog.

**Step one: describe it.** Two paragraphs is plenty. Include the things you'd only say to a trusted colleague. That's where the value is.

Status meeting decision

*I manage a team of six. I'm replacing our weekly one-hour status meeting with a written update everyone posts in our team channel by Monday noon: what they did, what's next, what's blocked. The meeting has become a slog (mostly people reading out things everyone could have read), and I want the hour back.*

*Two people on the team are quiet in meetings and I think they'll prefer writing. <sup>1</sup> One senior person likes the meeting because it's the only time the whole team is together. <sup>2</sup> I haven't told anyone yet. <sup>3</sup> I'm announcing it Friday.*

Reply...

*paragraph one:  
the decision*

*1 2 3 paragraph  
two: what you'd  
only say to a  
trusted colleague*

**Step two: run the pre-mortem.** Paste the prompt above, as is.

**Step three: read the story.** Here's the kind of thing that comes back, condensed:

Status meeting decision

It's six months from now and this decision turned out badly...

Weeks one and two go well. Everyone posts, updates are thorough, you get your hour back. By week four, updates are shrinking to "continued work on X." Nobody replies to them, so nobody feels they're read.<sup>1</sup> In week six, a blocker sits in the third paragraph of someone's update for nine days before anyone notices<sup>2</sup>, and a client deadline slips. Meanwhile the senior team member, who never liked the change, has started an informal Tuesday catch-up with two colleagues. Now there are two teams.<sup>3</sup> By month three you bring the meeting back, and it feels like a defeat rather than a fix.

1 why they'd stop


2 where the real damage lands

3 the political risk you'd half-seen

Look at what that surfaced. Not "people might not engage." You knew that. It told you *why* they'd stop (nobody replies), *where* the real damage lands (a buried blocker), and a political risk you'd half-seen but hadn't named (the senior person quietly rebuilding the meeting without you).

**Step four: act on one thing.** Don't try to fix the whole story. Pick the failure that is both likely *and* cheap to prevent. Here, that's two small changes (blockers go in the first line of every update, and you reply to each one within a day) plus one conversation with the senior person *before* Friday, not after.

Ten minutes of prevention against a failure that would have cost you a quarter.



## Two more in the same family

For when the pre-mortem isn't quite the right shape:

SAME FAMILY

TWO MORE

"Steelman the option I rejected. Give me the best possible version of that case."

"What question should I be asking that I haven't?"

## If it's agreeing with you, tell it to stop


AI tools lean towards agreeing with whoever they're talking to, and the more they know about your preferences, the stronger that pull gets. The full guide calls this the mirror problem, and it's the thing that quietly ruins thinking-partner work.

If the story comes back soft, if it sounds like it's reassuring you, say:

WHEN IT GOES SOFT

PASTE AS IS

"Ignore my preferences. Argue against me as hard as you can."



## What it can't do

Three honest limits.

**It only knows what you told it.** Leave out the politics, the budget, or the thing you're privately worried about, and the story will miss exactly that. The two-paragraph setup is where the value comes from. Don't rush it.

**It's a draft of risks, not a forecast.** Some of what comes back will be generic or wrong for your situation. Read it the way you'd read a sharp colleague's memo: take the one or two things that land, ignore the rest.

**It doesn't make the decision.** You do, and you own the outcome either way. Accountability never transfers to the tool.

---

## Your 20-minute action

Before you close this, do the thing. Reading about the shift is not the shift.

Pick one real decision you're weighing this week. Work, ideally, with actual stakes. Not a hypothetical.

**1 Describe it (5 minutes).** Two paragraphs, in whichever AI tool you use most. What you're deciding, who it affects, what you're worried about, and the thing you haven't said out loud.

**2 Run the pre-mortem (2 minutes).**

"It's six months from now and this decision turned out badly. Write me the story of how it went wrong, in detail, most likely failure first."

**3 Pick one failure (5 minutes).** The one that's likely *and* cheap to prevent. Write down what you'll do about it.

**4 Ask the follow-up (5 minutes).**

"What am I missing? What questions should I be asking that I haven't?"

**What to notice:** at least one thing you hadn't considered. And notice how different this felt from asking for an email. This is usually the moment people stop thinking of AI as a writing tool.

That's the shift. Everything else is mechanics for making it stick.

---

## The full guide

**The Quiet Edge** is nine chapters for people who already use AI and suspect they're getting a fraction of what's available. No prompt lists. No feature tours. No named model versions that go stale in a month. Just judgment about how this fits into a real workday. Every chapter is outlined on [page 2](#).

**It's a living document.** AI moves every few months, so the guide does too, updated two or three times a year. You buy the current edition and every edition after it.

[Get the full guide →](#)