

Critical Thinking for Adults

Michael Gannotti




Figure 1: Cover

Critical Thinking for Adults

Michael Gannotti

This illustrated edition is the teaching-manual of 2 September 2026. Endnotes run in one series. Chapter-opening images are original still-lives, not portraits of persons or children. Documents still constrain the facts.

Contents

- Welcome
- How to Use This Book This Week
- If You Only Remember Five Things
- The Work Hour
- Chapter 1 — What we mean
- Chapter 2 — Moves that scale
- Chapter 3 — Something to think about
- Chapter 4 — Two sources
- Chapter 5 — Graphs
- Chapter 6 — A short written reason
- Chapter 7 — On the job
- Chapter 8 — Approaching disagreement at work
- Chapter 9 — Records
- Chapter 10 — Resources
- A Note on Sources

Welcome

This book exists because the thinking on this job is the whole program. There is you, the work in front of you, and today's move — something you already know, a question worth waiting for, a look at one other thing, and a short written reason a colleague who missed the meeting could follow.

I wrote this for a capable, busy, willing working adult. You may be at a desk, on a shift, in a trade bay, a clinic, a warehouse, or at a retail counter. You already have a job. Many adults feel rusty when they sit down to practice thinking as thinking. That feeling is common. It is not a verdict. This book will make you fluent enough in today's move to hear an own-side "because" offered as a finished argument, a fluent recap offered as a check, and a finished-looking paragraph — including a model's paragraph — offered as a reason, and to ask a good question, without turning a stand-up into a courtroom.

You do not need to be a philosopher. You do need to hear "because we wanted it more" as a why-story, not a look at the notes. Ordinary work lives at the desk and on the floor: two vendor pages, a recap versus the notes, a spec versus a review, an invoice versus a purchase order, two handoff notes, a deck chart, posted hours versus the clock. A small set of moves, practiced on something you already know, with you holding the account.

The cover says *Critical Thinking for Adults*. The work is this desk, this bay, this shift. A job, a degree, and years of experience are not a bicycle you now ride in every domain.¹ You already think in pockets of the work you know. That is not

¹Daniel T. Willingham, "Critical Thinking: Why Is It So Hard to Teach?," *American Educator* 31, no. 2 (Summer 2007): people who have sought to teach critical thinking "have assumed that it is a skill, like riding a bicycle, and that, like other skills, once you learn it, you can apply it in any situation. Research from cognitive science shows that thinking is not that sort of skill." Same essay: critical thinking "is a type of thought that even 3-year-olds can

the same as having finished the dual job — hearing the other side, looking at a second named source, and writing a reason a colleague who missed the meeting could follow. Adulthood is not a thinking birthday.²

A 2026 National Association of Colleges and Employers member survey rated critical thinking 4.4 on a five-point scale — highly, not first. Communication and teamwork sat at 4.5.³ That survey is not a finding that this book works. The hour in front of you is the work.

What a good work hour looks like

You sit down already knowing today’s move. An object is on the table — a recap, a vendor page, a deck chart, an invoice, a handoff note, a clock beside posted hours. You look. Then you talk, or you ask one good question in a meeting — not “did we have a good meeting?” but “how do you know?” Then you wait. Three seconds feels long. It is the work. If the reason is own-side, you hear it. Then you look at one other named thing. Then you write the reason so a colleague who was not in the room could follow it. *The Work Hour* will teach this in full. Later chapters will not reinvent it.

engage in—and even trained scientists can fail in.” Access date for URLs in these notes: 2 September 2026. https://www.aft.org/sites/default/files/media/2014/Crit_Thinking.pdf. Daniel T. Willingham, “How to Teach Critical Thinking” (NSW Department of Education occasional paper, 2019): programs meant to teach general critical thinking skills have had limited success; when transfer is tested, “positive results have been absent or modest and quick to fade.” http://www.danielwillingham.com/uploads/5/0/0/7/5007325/willingham_2019_nsw_critical_thinking.pdf. A job is a domain. It is not all domains.

²Deanna Kuhn, “A Developmental Model of Critical Thinking,” *Educational Researcher* 28, no. 2 (1999): 16–25. Many adults remain multiplists, for whom critical thinking can feel irrelevant; only a minority reach an evaluative epistemology. Deanna Kuhn and Susan Pearsall, “Developmental Origins of Scientific Thinking,” *Journal of Cognition and Development* 1, no. 1 (2000): 113–29, as restated in Deanna Kuhn, “How Do People Know?,” *Psychological Science* 12, no. 1 (2001): 1–8: complex coordination of theory and evidence “may not develop fully even by adulthood.” Used here for what adulthood does not finish, not to turn this into a science book.

³National Association of Colleges and Employers, *Job Outlook 2026: Spring Update* (Bethlehem, PA: NACE, 2026). Survey conducted 12 February–17 March 2026; 185 total respondents; 142 NACE employer members (19.9 percent of eligible member respondents); 43 additional non-member responses. Importance on a 5-point scale (n = 153): Communication and Teamwork 4.5, Critical Thinking 4.4. Share rating critical thinking very or extremely important versus very or extremely proficient: 93.5 percent / 50.0 percent (proficiency n = 148). A member survey, not a trial, and not a finding that this book works. <https://naceweb.org/docs/default-source/default-document-library/2026/publication/research-report/2026-job-outlook-report-spring-update.pdf>.



Figure 2: A work-hour setup

Working adults already think, in a domain they know. A person who offers a fluent own-side “because” about a roster they run is thinking. They have not finished the dual job, and they have not written a reason a colleague could follow until they have. Adults produce argumentative discourse. They can be more strategic than teens and still fortify their own side.⁴ Critical thinking is not a bicycle a hire, a

⁴Mark Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, no. 2–3 (2001): 135–53. Thirty-one community-college young adults were more strategic than thirty-three seventh- and eighth-graders from the same inner-city population, and still own-side: they dismantled when they disagreed and fortified when they agreed. The study topic is not an exercise in this book. Keith E. Stanovich, Richard F. West, and Maggie E. Toplak, “Myside Bias, Rational Thinking, and Intelligence,” *Current Directions in Psychological Science* 22, no. 4 (2013): 259–64: myside bias “shows very little relation to intelligence,” including in university samples, and in several naturalistic between-subjects designs the association was “literally zero.”

diploma, or a title confers.⁵ We practice the moves on the work in front of us.

Critical thinking, in Robert Ennis's public sentence this book keeps, is reasonable reflective thinking focused on deciding what to believe or do.⁶ A person deciding whether a recap's "we hit the number" is true, whether to hold the line or call the vendor, or whether an email is a request, a rumor, or a decision already made is inside that sentence. "Doing justice to challenges" is the expected move: the other side is a person in the room, and the written reason has to hold a second side. "Critical" here does not mean find the fault in people.

If you used *Critical Thinking Skills for Little Thinkers*, *Critical Thinking for Young Minds*, or *Critical Thinking for Emerging Adults* at a kitchen table, one pointer each, then this desk. The five moves still work: wait; what do you mean; how do you know; what else could be true; how could we check. Young Minds added two-source, graph, counterexample, and steelman. Emerging Adults added a short written reason a stranger could follow, with a parent in the room. This book keeps that written reason and moves it onto the job: a colleague who missed the meeting could follow it. It is not a reprint of those lesson banks, and it is not a kitchen-table parent book.

Critical Thinking for Life is a different book, for a different reader. One pointer, if you need it later: Chapter 6 of that book is "Leave the page," and Chapter 8 is "When a machine finishes." Then we teach at work.

⁵Willingham, "Critical Thinking." Same bicycle constraint as note 1, applied here to hire, diploma, and title. Experts encapsulate; they do not become general critical thinkers. Huber and Kuncel's college gains are college, not the job: Christopher R. Huber and Nathan R. Kuncel, "Does College Teach Critical Thinking? A Meta-Analysis," *Review of Educational Research* 86, no. 2 (2016): 431–68.

⁶Robert H. Ennis, "What is critical thinking?," last revised 26 November 2010, <https://criticalthinking.net/what-is-critical-thinking/>: "Critical thinking is reasonable reflective thinking focused on deciding what to believe or do." Same sentence in Robert H. Ennis, *The Nature of Critical Thinking: Outlines of General Critical Thinking Dispositions and Abilities* (2013; last revised September 2015), <https://criticalthinking.net/wp-content/uploads/2024/04/The-Nature-of-Critical-Thinking.pdf>. Ennis: although "critical" is sometimes used in a negative sense, the concept presented here is not negative. "Doing justice to challenges" is named in the brief conception. The long outline is not a workplace syllabus.

What you will actually get

Each teaching chapter does seven jobs, always in the same order. You will learn why this week's idea is worth the struggle, and understand it yourself. You will get a session you can run this week: exact wording, named try-its at the desk and on the floor or shift, and a talk/meeting box. If it isn't clicking, three diagnostics and a next move. Tools, including AI, stay optional. "Done enough" lets you place by skill rather than by job title.

The eight teaching chapters follow the work, not twelve thin skill labels: what we mean; moves that scale; something to think about; two sources; graphs; a short written reason; on the job; approaching disagreement. Records and resources come last.

Today you can look at a recap and the notes — or a vendor page and a spec, or posted hours and the clock — say how you know, write the reason so a colleague who missed the meeting could follow it, and change the sentence if the look goes the other way.

What this book will not do

This book will not hand you 180 days of worksheets. Coverage is a map. Depth is the hour.

It will not treat a branded philosophy hour as a miracle. Conversation is worth ten minutes as conversation. A named school program is not a proven engine of generic critical thinking, and the large trials that exist were not this workplace.⁷ Later, a short resources chapter names common programs by fit. The file still says the job you already hold.

⁷Education Endowment Foundation, "Philosophy for Children – second trial," completed November 2021, independent evaluator NFER, <https://educationendowmentfoundation.org.uk/projects-and-evaluation/projects/philosophy-for-children-effectiveness-trial>: 198 English primary schools, 5/5 padlock security; on average, no additional progress on Key Stage 2 reading or maths for the whole cohort (0 months). English Year 6 is about age 10–11 — not this book's working-adult pond. Stephen Gorard, Nadia Siddiqui, and Beng Huat See, *Philosophy for Children: Evaluation Report and Executive Summary* (Education Endowment Foundation, 2015), ERIC ED581147: the smaller efficacy trial that did not replicate at effectiveness scale. Dialogue may still be worth ten minutes as dialogue. A branded hour is not a proven engine, and those large trials were KS2.

It will not franchise a thinking wheel, prepare a hiring inventory, or promise a score. There is no guaranteed percentile, no thinking diploma, no certificate of thinking, and no job titled Critical Thinking I. The promise is a path: the moves, on something you actually know, until you can do the next one unaided.

It is not a sequel. The homeschool books stay at their tables. *Autonomous AI and Education*, *Science for Homeschooling*, and *Computer Science for All Ages* are other books. This one is the thinking on this job.

It is not a news-literacy course for adults, and it is not a culture-war pamphlet. This book teaches *how to approach* disagreement. It does not rehearse live fights as content. We practice leaving a vendor page, a recap, a spec, an invoice, a handoff note, and posted hours. Politics, religion, identity fights, war, and medical culture-war stay off this table. Method yes. Culture-war content no.

The promise

If you only remember one sentence, remember this: a path through ordinary work is the promise. A diploma is not.

Title is not placement. Hire is not a window that closed.

Three things have to hold. You understand today's move well enough to hear an own-side "because" offered as a finished argument, a fluent recap offered as a check, and a finished-looking paragraph offered as a reason. You talk, look at a second named source, and write a short reason a colleague who missed the meeting could follow before the tool. Any helper — including an AI tool — stays a supplement you host and constrain, not a thinking partner for the first attempt, not a secret friend, and not a substitute for the unaided reason or the look at one other thing.

You can do this. You do not have to know next week's move today. You have to know this week's move well enough to sit still while you struggle, then ask one good question. Start here. Read *How to Use This Book This Week*, then the five things, then *The Work Hour*. Open the chapter that matches the skill in front of you. You will know more after one chapter than you know this morning. You will have something to try today at the desk, and something to try on the floor. We can do this.

How to Use This Book This Week

Start at the skill in front of you, not on page one because a catalog, a title, a degree, or a well-meant colleague said so.

This book is a handbook you open to this week's move, not a novel you read cover to cover. The teaching chapters run: what we mean; moves that scale; something to think about; two sources; graphs; a short written reason; on the job; approaching disagreement. Records and resources come last. Role and context bands — **desk** and **floor-or-shift** — live *inside* each teaching chapter. They tell you the grain of the same moves. They do not tell you where *this* job sits.

This book does not replace the process already on the floor. If a protocol, a roster, a spec, a handoff, or a shop floor you already use is working, keep it. Use this book to hear an own-side “because,” to run the hour, and to ask for a short written reason a colleague who missed the meeting could follow. Combine honestly. The title is still the job you already hold.

Pick the chapter by skill, not by title

Open the chapter you think is right. Skip to **What “done enough” looks like** at the end of the *previous* chapter, or to **If it isn't clicking** in the one you opened. If you can already do those checks unaided, you are too early. If the checks from two chapters back are still failing, drop back.

If you still live in own-side exposition and need the prompt “what would someone who disagrees say?”, you are still in the early band of the moves chapter, whatever your title.⁸ If you cannot yet hold a two-sentence steelman, a closed-window para-

⁸Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator* 31, no. 2 (Summer 2007): even 3-year-olds can think in a known domain; thinking is not a bicycle

graph after a tool, and a short written reason a colleague who missed the meeting could follow, you are still in a prompted-dual-focus stage, whatever the degree. A fluent recap with no notes is not “behind in philosophy.” It is still learning to leave the sentence and look at a second source.

If you can already do those checks unaided, and you are still in a “critical thinking workshop” only because the cover says so, skip ahead. Instruction should match what you actually need and already know. A publisher’s workshop label is a scope, not a licence, and not a job title.

How you use a teaching chapter

Each teaching chapter has the same seven headings, always in this order, so you always know where the explanation, the session, the named try-its, the talk/meeting box, and the tools box live.

Why this matters tells you what this idea unlocks. **Understand it yourself** gives one everyday picture, one precise picture, and three to five wrong reasons you should be able to hear — including from yourself.

How to try it this week assumes the session shape from *The Work Hour*. It will not rebuild the hour. It will give you this week’s named try-its — at least one at the desk and one on the floor or shift — and a talk/meeting box.

Practice that actually builds it names the try-its with time, materials, safety, and the move. The desk, the floor, the clinic, and the shop *are* the practice. **Tools, including AI** is optional and short. The rules live once in *The Work Hour*. **What “done enough” looks like** is how you leave: unaided work, not a perfect Tuesday.

You do not have to read the whole chapter tonight. You do have to read *this week’s* move before you put the artifact on the table.

skill. Access date for URLs in these notes: 2 September 2026. Deanna Kuhn and Susan Pearsall, “Developmental Origins of Scientific Thinking,” *Journal of Cognition and Development* 1, no. 1 (2000): 113–29, as restated in Deanna Kuhn, “How Do People Know?,” *Psychological Science* 12, no. 1 (2001): 1–8: increasing cognitive control over theory and evidence “may not develop fully even by adulthood.” Mark Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, no. 2–3 (2001): 135–53: adults produce strategic dialogue and still own-side. Used here for placement by skill, not to turn this into a science book.

The five-minute warm-up

Five minutes. Phone face down. Artifact in reach.

1. Read today's move until you can say it in one sentence.
2. Look at one object yourself. Ask, out loud, what would count as a check. Stay off any tool until you have looked.
3. Glance at the "wrong reasons you should be able to hear." Name the one you would have given last quarter.
4. Write one sentence you will actually say. Not a speech. Example: "How do you know?" Or: "What's a second place we could look?" Or: "Can you write that so a colleague who missed the meeting could follow it?"
5. Close the book to the try-it. You are ready.

If you are learning the move *while* you are stuck in the meeting, you will talk too much. Prepare first. Then sit still.

At the desk and on the floor

Every teaching chapter names try-its at the desk and try-its on the floor or shift — a handoff, posted hours versus the clock, a count, a clinic board, a shop floor, a vendor call, two quotes for the same SKU. The talk/meeting box travels. The live fight does not. If a real workplace disagreement is too hot today, pick the recap, the vendor page, the deck chart, the invoice, or the handoff note.

The five moves from ages 5–10 still work. The additions Young Minds made still work. Emerging Adults' written reason still works. This book does not reprint those lesson banks. Use this book's job objects, and the written reason a colleague who missed the meeting could follow, when the claim can bear them.

Pictures in the chapter show objects — two pages, a chart, a notebook — not a colleague to copy. Notes live at the back. "Critical" does not mean "find the fault in people."

Safety is ordinary. No public quiz of a junior in front of a vice-president. No undercover audit of a workplace. No gotcha against a colleague, a vendor, or a shop. No lecture to a stranger's customer. Noticing is not a sting. A clinic try-it uses a protocol already in use.

When to skip ahead

Skip ahead when this chapter’s “done enough” checklist is already true *unaided*. Slow down when the same own-side “because” repeats after a clear look at a second source, or when a finished-looking paragraph still has no named check. **If it isn’t clicking** will give you three likely causes and a next move.

A stranger-readable record is a dated note and a short written reason plus the check, sources named, inside the job’s real file. The title a stranger can read is still analyst, nurse, clerk, lead, tech, manager — not Critical Thinking I, and not a thinking diploma.⁹

You hold the account. A chatbot is never the only partner. Try unaided first. The full allow-and-ban list lives in *The Work Hour*.

For this week: pick the chapter by skill. Do the five-minute warm-up. Run the hour as *The Work Hour* describes it. Stop talking sooner than you want to.

⁹Work products a stranger can already understand live *inside* the job the adult holds — the project file, the audit workpaper, the nursing note, the editor’s file, the engineering calculation package, the dated note in the repo. Never Critical Thinking I. Never a thinking diploma. There is no AP Critical Thinking among College Board’s listed AP subjects.

If You Only Remember Five Things

Keep this page. The chapters will add wording, objects, and this week's questions. They will not replace these.

1. You already think in a domain you know. Critical thinking is deciding what to believe or do, with reasons, on something you actually know — a recap, a vendor page, a spreadsheet tab, a roster, a SKU, a process document, a deck chart, two handoff notes — and doing justice to challenges.¹⁰ It is not emptiness waiting for a kit, and it is not a bicycle a job, a degree, or years of experience confers. A person who offers a fluent own-side “because” about a roster they run is thinking. They have not finished the dual job. Start now. Give yourself something you already somewhat know.

2. Moves that scale, plus a short written reason a colleague who missed the meeting could follow. Wait. What do you mean? How do you know? What else could be true? How could we check? Those five still work. Two-source, graph, counterexample, and steelman still work. This book keeps one object and moves it onto the job: a short written reason a colleague who missed the meeting could follow. *I claim X because Y. We checked by Z.* Name the move. Practice it next Tuesday on a similar case. Then write the sentence. One good question beats five.¹¹

¹⁰Robert H. Ennis, “What is critical thinking?,” last revised 26 November 2010, <https://criticalthinking.net/what-is-critical-thinking/>: “Critical thinking is reasonable reflective thinking focused on deciding what to believe or do.” Access date for URLs in these notes: 2 September 2026. Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator* 31, no. 2 (Summer 2007): even 3-year-olds can think in a known domain; thinking is not a bicycle skill. Mark Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, no. 2–3 (2001): 135–53: adults produce argumentative discourse and still own-side.

¹¹Willingham, “Critical Thinking”: people have assumed critical thinking “is a skill, like riding a

3. You talk and look at a second named source before the tool. You ask one good question and wait. Three seconds is a convention, not a sacrament.¹² The silence is the work. After they stop — or after you stop — wait again. Struggle before rescue: try unaided first. Ask, wait, hint, then look. You do not finish the reason. You do not grab the answer. Then you look at one other named thing. A helper, including an AI tool, is a supplement you host. It does not sit in the chair during the first attempt. You hold the account. A chatbot is never the only partner.

4. A finished-looking paragraph is not a check. A smooth “because we wanted it more,” “because the recap said we hit the number,” “because the graph looks huge,” or a tidy paragraph with no named source — including a model’s paragraph — can be a starting point. It is not the look.¹³ Leave this sentence and look at one other thing: the notes, the spec sheet, the purchase order, the other handoff, the axis and the numbers, the clock beside the posted hours. If the look goes the other way, you may change your mind — and you correct the written sentence. That is the work. It is not losing.

bicycle,” and that once learned it applies in any situation; “Research from cognitive science shows that thinking is not that sort of skill.” Daniel T. Willingham, “How to Teach Critical Thinking” (NSW Department of Education occasional paper, 2019): programs meant to teach general critical thinking skills have had limited success. The five moves scale from *Critical Thinking Skills for Little Thinkers*; *Critical Thinking for Young Minds* added two-source, graph, counterexample, and steelman; *Critical Thinking for Emerging Adults* added a short written reason a stranger could follow. This book keeps that written reason and moves it onto the job. Deanna Kuhn and Amanda Crowell, “Dialogic Argumentation as a Bridge to Argumentative Writing,” *Written Communication* 28, no. 3 (2011): extra essay writing in the comparison group did not produce dual perspective; the missing interlocutor has to have been in the room.

¹²Robert J. Stahl, “Using ‘Think-Time’ and ‘Wait-Time’ Skillfully in the Classroom,” ERIC Digest ED370885 (1994), <https://files.eric.ed.gov/fulltext/ED370885.pdf>: think-time; about three seconds as a convention, not a sacrament; written for classroom teachers. Mary Budd Rowe, “Wait-Time and Rewards as Instructional Variables,” NARST paper, ERIC ED061103 (1972), <https://files.eric.ed.gov/fulltext/ED061103.pdf>, is elementary science class, not a workplace trial. This book uses the pause so already-present reasoning can be heard.

¹³Deanna Kuhn, “How Do People Know?,” *Psychological Science* 12, no. 1 (2001): 1–8: people mash an explanation of why a claim makes sense with evidence that the claim is true; substantial variation remains among adults. Sam Wineburg and Sarah McGrew, “Lateral Reading and the Nature of Expertise,” *Teachers College Record* 121 (2019): fact checkers leave an unfamiliar page; Ph.D. historians often stayed. At work the same *direction of attention* is leave this vendor page, this recap, this invoice, this handoff, these posted hours, and look at one other named thing. Joel Breakstone et al., “Students’ Civic Online Reasoning: A National Portrait,” *Educational Researcher* 50, no. 8 (2021): 505–15, is why the habit is taught (high-school sample); this book does not assign those items.

5. Ordinary work stakes. You hold the account. Practice on two vendor pages, a recap versus notes, a spec versus a review, an invoice versus a PO, two handoff notes, a product page versus a data sheet, a deck chart, posted hours versus the clock, a shift rumor about who works when, “we always do X” versus last quarter’s process document. This book teaches *how to approach* disagreement. It does not rehearse live fights as content. For tools: COPPA’s under-13 line does not apply here. That is not a thinking birthday and not a licence for an unsupervised chatbot as the only partner.¹⁴ A path through ordinary work is the promise. A diploma is not. Title is not placement. There is no window that closed when you were hired.

You talk, look, and write before the tool. You hold the account. Five things. Then sit down and try.

¹⁴Children’s Online Privacy Protection Act, 15 U.S.C. §§ 6501 *et seq.*; 16 C.F.R. Part 312. Federal Register 90 FR 16918 (22 April 2025): amended Rule effective 23 June 2025; general compliance date 22 April 2026. Bright line: children under 13. A working adult is not under 13. That is an account fact, not a licence for unsupervised chat. The full allow-and-ban list lives in *The Work Hour*. There is no job titled Critical Thinking I. There is no AP Critical Thinking.

The Work Hour

The hour has a shape. Learn it once. Later chapters will give you today's move. They will not rebuild this hour. When a chapter says "run the session," it means this.

You do not need a school bell. You need a warm-up, a short model, a real attempt, one good question then wait, a look at a second named source, a short written reason a colleague who missed the meeting could follow, and an unaided exit. Fifteen to twenty-five minutes on a named artifact. The shape does not change. The object on the table does.

Sit down having already done the five-minute warm-up. Ten to twenty-five minutes of a named move, a few times a week, on actual job artifacts, is the right ambition — not under-running of a proven workplace programme, because there isn't one.¹⁵ The honest pitch is *often enough that the question is familiar*.

¹⁵Education Endowment Foundation, "Philosophy for Children – second trial," completed November 2021, independent evaluator NFER, <https://educationendowmentfoundation.org.uk/projects-and-evaluation/projects/philosophy-for-children-effectiveness-trial>: 198 English primary schools, 5/5 padlock security; on average, no additional progress on Key Stage 2 reading or maths for the whole cohort (0 months). English Year 6 is about age 10–11 — **not this book's working-adult pond**. Stephen Gorard, Nadia Siddiqui, and Beng Huat See, *Philosophy for Children: Evaluation Report and Executive Summary* (Education Endowment Foundation, 2015), ERIC ED581147: the smaller efficacy trial that did not replicate at effectiveness scale. Dialogue may still be worth ten minutes as dialogue. Access date for URLs in these notes: 2 September 2026. What Works Clearinghouse, reviews search re-opened 2 September 2026: no named adult or workplace critical-thinking intervention report. Session minutes in this chapter are pedagogical design, not a national thinking-minutes table.

1. Warm-up

Two to four minutes. An object you already know. Say: “We’re going to start with something I already know.” Keep this short. A maxim about “seeing both sides” is noise until it attaches to something you already know.¹⁶

2. Short model

Three to six minutes. One new move, visible. Then you stop. When the idea is new, look at one incorrect example — a why-story offered as evidence, a fluent recap with no notes, a steelman that is actually a straw, a finished-looking paragraph with no named check. If you are still talking at minute eight, close the model.

3. Your attempt

This is the center of the hour. Five to ten minutes. The first talk and the first look are yours. If you stall: ask, wait, hint, then look. Not the reverse. Keep your hands off the finished sentence. A colleague or a tool does not sit in the chair during the first attempt.

4. One good question, then wait

An authentic question is one for which you have not already written the answer. “Did we have a good meeting?” is recitation. “How do you know we hit the number?” is the work.

Then wait. A slow three after you ask. A slow three again after they stop. Robert Stahl called that silence *think-time*; that label is classroom-origin, not a stand-up trial. About three seconds is a convention, not a sacrament.¹⁷ Mary Budd Rowe

¹⁶Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator* 31, no. 2 (Summer 2007): you can teach maxims about how people ought to think, “but without background knowledge and practice, they probably will not be able to implement the advice they memorize.”

¹⁷Robert J. Stahl, “Using ‘Think-Time’ and ‘Wait-Time’ Skillfully in the Classroom,” ERIC Digest ED370885 (1994), <https://files.eric.ed.gov/fulltext/ED370885.pdf>: eight categories of classroom silence; about three seconds as a convention. Classroom-origin. Not a finding that wait-time teaches critical thinking, and not a stand-up randomized trial.

found the same pause in elementary science class. That was science class.¹⁸ If they are mid-reason, let them finish. If the silence is hard, look at the notes, the graph, the roster, or the clock, not at the person. At work the shrug often looks like refusal; leave it. After “write the sentence,” that is task-completion time. Wait in private when the power difference is steep. One question. Then stop.

5. Hear an own-side “because”; look at a second source

Someone says “because we wanted it more.” That is a why-story. Hear it. A better question: “That’s why you hope so. What would show it?” Then look: the notes, the spec sheet, the purchase order, the axis and the numbers, the posted roster, the clock.

A fluent sentence is not a check.¹⁹ A finished-looking paragraph is not a check. Leave this sentence and look at one other named thing. If the look goes the other way, you may change your mind. That is the work. It is not losing.

6. Write the sentence, then unaided exit

Two to eight minutes. Talk first. Then write. Then look again if the sentence cannot survive the check. Book closed for the exit. No hints. No chatbot. Early on: a spoken sentence plus one written line that names the claim and the check.

¹⁸Mary Budd Rowe, “Wait-Time and Rewards as Instructional Variables: Their Influence on Language, Logic, and Fate Control,” NARST paper, ERIC ED061103 (1972), <https://files.eric.ed.gov/fulltext/ED061103.pdf>; elementary science classes; teachers typically waited about one second after a question and 0.9 seconds after a student answer; extending both to three to five seconds changed discourse. Rowe is science class. This book does not treat Rowe 1972 as a workplace randomized trial.

¹⁹Deanna Kuhn, “How Do People Know?,” *Psychological Science* 12, no. 1 (2001): 1–8. Mark Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, no. 2–3 (2001): 135–53. Sam Wineburg and Sarah McGrew, “Lateral Reading and the Nature of Expertise,” *Teachers College Record* 121 (2019): professional fact checkers leave an unfamiliar page; Ph.D. historians often stayed. At work the same direction of attention scales to a vendor page, a recap, an invoice, a handoff, posted hours, and a process claim of “always” — not a civic-web course. Joel Breakstone et al., “Students’ Civic Online Reasoning: A National Portrait,” *Educational Researcher* 50, no. 8 (2021): 505–15, is why the habit is taught (high-school sample); this book does not assign those items.

Later: a short written reason plus the check, sometimes a two-sentence steelman, sometimes a closed-window paragraph after any tool. Fade until the check does not need a prompt.

The talk/meeting-box shape, once

Opening question. Authentic. About what they mean, how they know, what a second source would show, what a chart actually shows, the strongest fair version of the other side, or whether a colleague who missed the meeting could follow the written reason. Not “did we have a good meeting.” Use one: What do you mean? How do you know? What else could be true? How could we check? What’s a second place we could look? What does this chart actually show? Can you write that so a colleague who missed the meeting could follow it?

Three to five follow-ups. Pick; you do not need all five every Tuesday. Meaning: What do you mean, *always*? Reason: How do you know? Alternative: Always? Check: What’s a second place we could look? Write or steelman: You may change your mind.

How to wait. As in section 4. Wait in private when the power difference is steep. Interrogating a junior colleague in front of a vice-president is not this book’s wait.

What a stuck silence usually means. The question was vague; wait-time after the question was zero; they are guessing what you want; they offered a why-story where a look was asked for; they think “critical” means finding fault; they think steelman means giving up; they think a finished-looking paragraph — or a model’s paragraph — is already a check. Debug, once: shrug → a smaller question and wait. Talks forever → “write one sentence.” Session becomes a fight → stop. Pose, wait, and point back at the object.

When two people do not agree, the same box holds. You do not have to win. You have to listen, then ask, then check.

1. Listen. Let them finish.
2. Ask what they mean. “What do you mean, *leading*?” “What do you mean, *always*?” Meaning first.
3. Ask for a reason. “How do you know?” A feeling, or a fluent paragraph, or a model’s paragraph, can be named. It is not yet a check.
4. Two sources. Leave this page and look at one other named thing: the spec,

the notes, the purchase order, the other handoff, the data sheet, the clock, last quarter's process document.

5. Wait. A fast answer is not better because it was first.
6. You may change your mind. That is the work. It is not losing.
7. A fluent sentence is not a check. A finished-looking paragraph is not a check.
8. Steelman. Say their claim in the strongest fair way. Then answer *that*. Inventing a better argument they did not make, then knocking it down, is not steelman. Asking a machine to write the steelman so you can paste it is write-my-reason.
9. A graph can mislead by scale. Before you trust a picture of "more," ask where the axis starts.

If a real workplace disagreement is too hot today, pick the vendor page, the deck number, the shift rumor, or the invoice.²⁰

The desk try-it shape, once

Time. Fifteen to twenty-five minutes, then stop.

Materials. A recap already in the inbox, a deck already built, a spec already written, a spreadsheet tab, a dated notebook.

Safety. Tone: curious, not courtroom. No public quiz of a junior in front of a vice-president.

The move. How do you know. Two-source. Graph. Steelman. Written reason plus the check.

Types you will meet, not a lesson bank: one-question stand-up (3–5 min); talk/meeting-box after a claim (2–8 min); wait drill, invisible; two-page comparison (8–15 min); recap versus notes (5–10 min); spec versus review (8–15 min); graph on a slide (5–10 min); paper argument map (8–15 min); steelman in speech

²⁰Willingham, "Critical Thinking": the processes of thinking are intertwined with the content of thought; even trained scientists can fail outside a domain they know. Stephan Lewandowsky, John Cook, Ulrich Ecker, and 19 co-authors, *The Debunking Handbook 2020* (2020): backfire effects occur only occasionally; the continued-influence effect is the sticky problem; detailed refutation outperforms a bare denial. Translation at work: repeat the check, not the rumor, on a vendor page, a deck number, or a shift rumor about who works when. This is a translation, not a workplace trial of the handbook.

(5–10 min); short written reason plus the check (10–20 min); last thoughts (1–2 min).

The floor/shift try-it shape, once

“Floor” means the handoff, the posted hours versus the clock, the count, the clinic board, the shop floor, the vendor call, two quotes for the same SKU.

Types you will meet: two handoff notes (3–10 min after, not during, a live procedure); posted hours versus clock (2–5 min); two quotes for the same SKU (5–10 min); invoice versus purchase order (5–10 min); product page versus data sheet (5–10 min); “what do you notice?” on a bay or a board (3–8 min).

Safety, once. No public quiz, no undercover audit, no gotcha, no lecture to a stranger’s customer. Noticing is not a sting. A clinic try-it uses a protocol already in use.

These are not the hour: a “gotcha” audit; a live-controversy field trip; civic-web test items; a five-paragraph essay night.

The written-reason shape, once

Chapter 6 owns this move. Later chapters fill the claim.

After talk: **I claim X because Y. We checked by Z.** Optionally a second sentence: **Someone who disagrees would say...** Dated. Stimulus named. Sources named. A colleague who was not in the meeting could follow it.

Talk first. Write. Then look. If the sentence cannot survive the check, correct the sentence. That is the work. It is not “never correct.”

It is not a five-paragraph essay, a hiring inventory, or a college thinking test. Extra essay writing without someone in the room is not the intervention: in one classroom comparison, the group that wrote more essays stayed near one-third dual perspective.²¹ The missing interlocutor has to have been in the meeting, or on the

²¹Deanna Kuhn and Amanda Crowell, “Dialogic Argumentation as a Bridge to Argumentative Writing,” *Written Communication* 28, no. 3 (2011): after dense dialogic practice, dual-perspective on a longitudinal essay reached 79 percent experimental versus about 29 percent comparison; extra essay writing in the comparison group did not help. Ages 11–12 at entry. Copy the shape (a reason,

floor. History-class studies licence the *shape* — more than one document, a named writing move — not a packet from another course.²² Ten to twenty minutes. One sentence plus the check is enough.

The AI rules, once

Later chapters will point here. You talk to the page, to a named colleague, and to a second source. A chatbot is never the only partner. An agent is not the thinking partner for the first attempt. The unaided reason is still yours. You hold the account. COPPA's under-13 line does not apply to a working adult. That is not a thinking birthday and not a licence for unsupervised chat.

After you have tried, you **may**:

- Ask a tool to **explain this page to yourself** from a **named** source (title, URL, page number, contract clause, spreadsheet tab). You still have to understand it well enough to hear a wrong turn.
- Ask for **extra questions**, then **vet** them. Throw out any question that imports a live controversy, a secret, a pretend expert, a colleague's private performance, or a civic-web test item.
- Use a labelled **SCRIPT** only **after you tried** — a short spoken or written sequence for *you* to say, not a paragraph to copy as your reason, your email, or your steelman.
- Ask for a **hint** after an attempt. A hint is not the finished reason and not the finished email.
- Ask a tool to **diagnose a reason you already gave**. You hear the diagnosis and decide. You still produce the next sentence.

You **may not**:

a counter, a second look, sometimes the other person's strongest case, then a short written reason). The paper's later topics are not exercises in this book.

²²Jeffery D. Nokes, Janice A. Dole, and Douglas J. Hacker, "Teaching High School Students to Use Heuristics While Reading Historical Texts," *Journal of Educational Psychology* 99, no. 3 (2007): 492–504; ages 16–17; multiple texts beat textbook on sourcing and corroboration. Susan De La Paz and Mark K. Felton, "Reading and Writing from Multiple Source Documents in History: Effects of Strategy Instruction with Low to Average High School Writers," *Contemporary Educational Psychology* 35, no. 3 (2010): 174–92; grade 11. Avishag Reisman, "Reading Like a Historian: A Document-Based History Curriculum Intervention in Urban High Schools," *Cognition and Instruction* 30, no. 1 (2012): 86–112; 236 eleventh-graders; historical thinking. Labelled history-class evidence, not a generic critical-thinking gym and not a workplace trial.

- Park an **unsupervised chatbot** as your only thinking partner.
- Treat “I am an adult / I hold the account / COPPA is off” as a permission slip.
- Ask the tool to **write your reason, email, or steelman**, then paste that text as yours.
- Ask the tool to **invent an expert identity** (“you are a procurement professor”; “you are our general counsel”) and then treat that voice as a person.
- **Grade** a colleague’s thinking — or your own — from a **detector** score.
- Print a **certificate of thinking**.
- Ask the tool to **invent facts** (what the spec lists, what the roster says, what the spreadsheet totals, what the contract clause is, what the client emailed).
- Hand yourself an **agent** (Hermes Agent, Grok Bot, or any cloud-computer worker) as a thinking partner for the first attempt.
- Paste **undisclosed AI-generated text** as yours where a named employer or journal forbids it.
- Put **confidential** workplace data into a **public** consumer tool.

On some consulting tasks inside GPT-4’s frontier, adults using the model were faster and higher-rated; on a designed outside-frontier task they were less likely to be right.²³ A fluent memo is not a check that the task was inside the frontier. This is a useful study, not a promise that every workplace will see the same result. On mid-level professional writing, ChatGPT cut time and raised grader scores and mostly substituted for effort — many treated writers submitted the first pass.²⁴ A

²³Fabrizio Dell’Acqua, Edward McFowland III, Ethan Mollick, Hila Lifshitz-Assaf, Katherine C. Kellogg, Saran Rajendran, Lisa Kraymer, François Candelon, and Karim R. Lakhani, “Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality,” Harvard Business School Working Paper 24-013 (2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4573321. Public abstract: 758 BCG consultants; on 18 inside-frontier consulting tasks, consultants using GPT-4 completed 12.2 percent more tasks, 25.1 percent faster, and more than 40 percent higher quality versus control; on a designed outside-frontier task they were 19 percentage points less likely to produce correct solutions. Rachel Layne, “Humans vs. Machines: Untangling the Tasks AI Can (and Can’t) Handle,” Harvard Business School *Working Knowledge*, 22 November 2023, <https://www.library.hbs.edu/working-knowledge/humans-vs-machines-untangling-the-tasks-ai-can-and-cant-handle>, is journalism restating those figures. A useful workplace study, not a promise that every shop will see the same result, and not a 2026-model update.

²⁴Shakked Noy and Whitney Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science* 381 (2023): 187–92, <https://doi.org/10.1126/science.adh2586>: 453 college-educated professionals; time –40 percent; quality +18 percent. MIT working paper, 2 March 2023,

faster email is performance. It is not yet your reason.

A tool that finishes the reason can make practice look better and leave you worse when the window is closed. In a high-school math trial, an unguarded chatbot raised assisted practice and then cut the unaided exam.²⁵ That is high-school math, not a workplace trial. In a university writing study, ChatGPT improved essay scores without a matching gain in knowledge or transfer; the authors warn of “metacognitive laziness.”²⁶ College writing. Not a workplace trial. Detectors mislabelled more than half of some human essays as machine-written.²⁷ A detector score is not a grade for a colleague’s email. Language models invent sources.²⁸ If you asked a tool for a paper, open the paper.

College Board does not treat AI-generated answers as the student’s work. That is school, not a workplace statute.²⁹ Named workplace analogs: Nature Portfolio

https://economics.mit.edu/sites/default/files/inline-files/Noy_Zhang_1.pdf (not peer-reviewed; draft n = 444): 68 percent of treated participants reported submitting ChatGPT’s initial output without editing; the authors describe substitution for worker effort. Use *Science* for published numbers; the working paper for mechanism, labelled draft.

²⁵Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman, “Generative AI without Guardrails Can Harm Learning: Evidence from High School Mathematics,” *Proceedings of the National Academy of Sciences* 122, no. 26 (2025): e2422633122. Field experiment, nearly 1,000 Turkish high-school math students: a ChatGPT-like tool raised assisted practice grades 48 percent relative to control, then cut unaided exam grades 17 percent. High-school math, not a workplace trial. The transferable caution is the crutch.

²⁶Yizhou Fan, Luzhen Tang, Huixiao Le, Kejie Shen, Shuyi Tan, Yueying Zhao, Yuan Shen, Xinyu Li, and Dragan Gašević, “Beware of Metacognitive Laziness: Effects of Generative Artificial Intelligence on Learning Motivation, Processes, and Performance,” *British Journal of Educational Technology* 56, no. 2 (2025): 489–530: 117 university students, English writing task; authors warn of “metacognitive laziness.” College writing. Not a workplace randomized trial.

²⁷Weixin Liang, Mert Yuksekgonul, Yining Mao, Eric Wu, and James Zou, “GPT Detectors Are Biased against Non-Native English Writers,” *Patterns* 4, no. 7 (2023): 100779: average false-positive rate on human TOEFL essays 61.22 percent.

²⁸William H. Walters and Esther Isabelle Wilder, “Fabrication and Errors in the Bibliographic Citations Generated by ChatGPT,” *Scientific Reports* 13 (2023): 14045: 55 percent of GPT-3.5 citations and 18 percent of GPT-4 citations were fabricated in that study.

²⁹College Board, “2026 AP Exam Terms and Conditions,” updated 27 April 2026, Section 4: no part of test answers, work, or submission may be generated in whole or in part using any AI tools. School exam hygiene, not a workplace statute. Nature Portfolio, “Artificial Intelligence (AI),” <https://www.nature.com/nature-portfolio/editorial-policies/ai>: human accountability is non-transferable; fabricating citations is not permitted; chatbots are not authors. International Committee of Medical Journal Editors, “Defining the Role of Authors and Contributors,” <https://www.icmje.org/recommendations/browse/roles-and-responsibilities/defining-the-role-of-authors-and-contributors.html>: disclose AI-assisted technolo-

treats fabricating citations as not permitted; the International Committee of Medical Journal Editors requires disclosure and will not list a chatbot as an author; Microsoft Copilot documentation tells users to review before send; Harvard Business School's handbook forbids representing AI output as independent work and forbids putting confidential data in public tools. Open *your* employer's acceptable-use page. This book will not write a federal workplace statute for you.

Names, in ordinary language. Hermes 4 is a chat model. Hermes Agent is a different runtime — a computer that can keep going. Grok is a chat assistant. Grok Bot is a cloud-computer agent. A public-school partnership is a third Grok. Game artist makes pictures. None of these is a critical-thinking class. A chatbot answers in a window. An agent has tools and can keep going. Neither is a thinking partner for the first attempt.

COPPA covers children under 13. It does not gate a working adult. That fact is not a chatbot licence.³⁰

Not a secret friend. A tool that talks in the first person is still a tool. It does not get a private channel, a promise to keep secrets from a manager or a colleague, or a seat in the chair during the first attempt.

You can practice this entire book with no AI. Many readers will.

gies; chatbots should not be listed as authors. Microsoft, "Data, Privacy, and Security for Microsoft Copilot," <https://learn.microsoft.com/en-us/microsoft-365/copilot/microsoft-365-copilot-privacy>: output is not guaranteed factual; review before send. Harvard Business School MBA Handbook 1.3.3, <https://www.hbs.edu/mba/handbook/standards-of-conduct/academic/chatgpt-and-ai>: do not represent AI output as independent work; do not put confidential data in public tools. Open *your* employer's acceptable-use page.

³⁰Children's Online Privacy Protection Act, 15 U.S.C. §§ 6501 *et seq.*; 16 C.F.R. Part 312. Federal Register 90 FR 16918 (22 April 2025): amended Rule effective 23 June 2025; general compliance date 22 April 2026. Bright line: children under 13. A working adult is not under 13. That is an account fact, not a licence for unsupervised chat. Hermes Agent homepage, <https://hermes-agent.org/>: "Not a chatbot. Not a copilot. An agent." xAI, "Introducing Grok Bot," 11 August 2026, <https://x.ai/news/introducing-grok-bot>, and "Grok Bot is now included with more plans," 26 August 2026, <https://x.ai/news/grok-bot-more-plans>: Grok Bot is a cloud-computer agent; game artist on the jobs list makes pictures. None of these is a critical-thinking class. There is no Grok Bot education SKU on pages opened 2 September 2026.

A life of the habit

Growing fluency is more miles on the same road, not a new faculty and not a thinking birthday. Revisit the same question next Tuesday.³¹ After a look, say “I used to think... now I think...” Ground rules — we listen; we ask what you mean; we give a reason; we can change our minds; we write a reason a colleague who missed the meeting could follow — then the next real disagreement. Posting the rules is not practice. Placement is by skill, not job title. A diploma is not. There is no window that closed when you were hired. The honest pitch is *often enough that the question is familiar*.

A Tuesday, said plainly

Here is an illustration, not a reported workplace. Someone asks, “How do you know we hit the number?” and waits. A colleague says, “because we wanted it more.” They look at the notes. Then they write the sentence. Later chapters will say: run the session as in *The Work Hour*.

³¹Willingham, “How to Teach Critical Thinking” (NSW Department of Education occasional paper, 2019): identify what is meant by critical thinking in each domain; identify the content those skills need; sequence; revisit. Robert H. Ennis, *The Nature of Critical Thinking* (2015): Disposition 9 is take a position and change it when evidence and reasons are sufficient.

Chapter 1 — What we mean

Critical Thinking for Adults

Michael Gannotti

2 September 2026



Figure 3: A still-life from a high three-quarter view: a recap page with unreadable greeked type, notes with greeked lines, a generic clock, two vendor pages with greeking, and a dated notebook with a short unreadable line. No people. No logos. No readable UI.

Why this matters

How do you know we hit the number?

The recap is still in the inbox. Someone used the word as if it were a score. You nodded, or forwarded it, or already started the next thing. The question is not a vocabulary quiz. It is the work of this book, in one breath.

Critical thinking, in the sentence this chapter keeps, is reasonable reflective thinking focused on deciding what to believe or do. A person deciding whether a recap’s “we hit the number” is true is inside that sentence. So is a person deciding whether to hold the line or call the vendor. So is a person deciding whether an email is

a request, a rumor, or a decision already made. Belief and action both count. A fluent “because” is not yet either.

This idea unlocks a week you can actually use. You do not need a philosophy degree. You need to hear “because we needed the quarter” as a story about *why* a result might have happened, not as a look at the notes. You need one good question, a wait, and something on the table you already somewhat know. Then, when the claim can bear it, a short written reason a colleague who missed the meeting could follow.

The struggle is worth it because ordinary work is full of finished-looking sentences. Recaps. Vendor pages. Shift rumors. Posted hours. A deck that already looks done. Those sentences already decide what to believe and what to do. The move is to treat them as claims, not as weather.

If you only remember one thing from this chapter: you already think in pockets of the job. Thinking is not a bicycle you now possess because you have a title, a degree, or years on the floor. Own-side talk is still the default. Dual focus is still the work. Adulthood is not a thinking birthday.

This book is not a sequel to *Critical Thinking for Life*. One pointer: that book’s Chapter 6 (“Leave the page”) and Chapter 8 (“When a machine finishes”) sit next door. Then we teach at work. It is not a reprint of *Critical Thinking Skills for Little Thinkers*, *Critical Thinking for Young Minds*, or *Critical Thinking for Emerging Adults*. Those books practiced the same family of moves on a swing, a recap at eleven, a rainfall graph, a short written reason a stranger could follow with a parent in the room. The objects here are a recap versus notes, posted hours versus the clock, two vendor pages, an invoice versus a purchase order. Dinner-table is not the spine. Meeting, handoff, vendor call, dashboard, shop floor, clinic bay are.

Understand it yourself

Here is the everyday picture.

Someone in the huddle says we hit the number. You ask, “How do you know we hit the number?” They say, “Because we needed it.” That sentence can be true as a feeling. It is not yet a check. Needing is a why-story. The notes, the system of record, last week’s actual tickets — those are looks at *whether* the claim is true.

A second everyday picture, still ordinary. The bay is running late. You have to

decide whether to hold the line or call the vendor. Volume in the room is not a criterion. The spec, the count, and the clock are. Wanting the line to keep moving is real. Wanting is not a look.

Here is the precise picture.

Robert Ennis's public sentence is the one this book keeps: critical thinking is reasonable reflective thinking focused on deciding what to believe or do.³² Reasonable: there are reasons, not only a mood. Reflective: you pause long enough to look at those reasons. Focused on deciding: the work has a point. The last two words do the job's work. Belief *and* action. A recap's verb is a belief. Holding the line is an action. Both sit inside the sentence.

The concept is not negative. "Critical" at work does not mean "find the fault in people." Doing justice to challenges is the expected move: the other side is a person in the room, and the written reason has to hold a second side.³³ The long outline of dispositions and abilities could be used as a syllabus. It is not this book's syllabus. You already have a job. Keep the sentence. Practice it on ordinary stakes.

Beside Ennis, Matthew Lipman named three features as aims: criteria, self-correction, and context.³⁴ On a recap, that means naming what you are going by (the notes, the system of record, the posted target), being willing to change the sentence when the look does not match the claim, and noticing that "hit" in this huddle is not "hit" on a different dashboard. Those are aims for the hour, not a branded programme and not a trial result. A community of inquiry is a tone: we listen, we give reasons, we can change our minds. It is not a diploma, and it is not a school trial moved onto the floor.

³²Robert H. Ennis, "What is critical thinking?," CriticalThinking.NET, last revised 11/26/10: "Critical thinking is reasonable reflective thinking focused on deciding what to believe or do." The same sentence, with "and" inserted, opens Ennis, "The Nature of Critical Thinking" (2013, last revised September 2015). Ennis, "A Logical Basis for Measuring Critical Thinking Skills," *Educational Leadership* 43, no. 2 (October 1985): 45: "reflective and reasonable thinking that is focused on deciding what to believe or do." The long outline "could be used as a syllabus." This book takes the sentence, not the syllabus.

³³Ennis, same sources: the concept is not negative and is not mere persuasion. The brief conception includes developing a reasonable position and doing justice to challenges. Disposition 9: take a position and change it when evidence and reasons are sufficient.

³⁴Matthew Lipman, "Critical Thinking: What Can It Be?," *Educational Leadership* 46, no. 1 (September 1988): 38–43: skillful, responsible thinking that facilitates good judgment because it relies upon criteria, is self-correcting, and is sensitive to context. Treated here as aims, not as trial evidence, and not as a warrant for Philosophy-for-Children pedagogy at work.

Two refusals sit beside the sentence.

First: you already think. Daniel Willingham’s line is the one that stops the waiting. Critical thinking is a type of thought that even three-year-olds can engage in — and even trained scientists can fail in.³⁵ If a three-year-old can do it in a known domain, a hired adult is not waiting for a switch to flip. Mark Felton and Deanna Kuhn compared young adolescents with young adults from the same population. The adults produced more strategic dialogue. They were not empty of reasons. They were still own-side: better at dismantling a partner they disagreed with, and better at fortifying a partner they already agreed with.³⁶ Start now, with a recap, a posted schedule, a vendor page, an email you were already answering. You do not wait for a workshop, a percentile, or a new title.

Second: thinking is not a bicycle. People who have tried to teach it have often assumed that, like riding, once you learn it you can apply it anywhere. Cognitive science says otherwise. The processes of thinking are intertwined with what is being thought about.³⁷ You can teach maxims. Without knowledge of *this* claim, and practice on *this* kind of object, the maxim will not implement. Willingham’s famous analog — a band-marching word problem that was mathematically the same as a garden problem solved minutes earlier — is a word-problem study, not a workplace trial. Few people saw the analog; a hint helped some, not most. Cite it as the constraint that still binds a fluent lead. Fluency of speech is not evidence that a general faculty has been installed. The nurse who can think with a protocol can fail on a vendor graph. The electrician who can think with a print can fail on a quarterly slide. The analyst who can think with a model can fail on a shift-handoff.

³⁵Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator* 31, no. 2 (Summer 2007): 8–19. Thinking is not a bicycle skill; processes are intertwined with domain knowledge; even 3-year-olds can engage in it in a known domain, and even trained scientists can fail. The 19 percent / 35 percent analog is a word-problem study (band / garden, least common multiple), not a workplace RCT.

³⁶Mark Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, nos. 2–3 (2001): 135–53. Thirty-one community-college young adults compared with 33 seventh- and eighth-graders from the same inner-city population; adults more strategic (Counter-C 20.42% vs. 8.51% of utterances in disagreeing dialogues) and still own-side: they dismantle when they disagree and fortify when they agree. The lab topic is not this book’s exercise type; the finding is own-side default.

³⁷Willingham, *American Educator* (Summer 2007), and Willingham, “How to Teach Critical Thinking,” NSW Department of Education occasional paper (2019): generic programs have had limited success; transfer absent or modest and quick to fade; experts encapsulate; they do not become general critical thinkers. Maxims without knowledge and practice will not implement.

A job is a domain. It is not all domains.

Own-side is still live. Felton and Kuhn's adults were more strategic than the teens, and still occupied with their own side. Dual focus is promptable, not default. Keith Stanovich, Richard West, and Maggie Toplak found that myside bias — evaluating evidence in a way that favors your own prior view — shows very little relation to intelligence, including in university samples. In several of those studies the association was literally zero.³⁸ Being “smart,” holding a degree, or sitting a Saturday test years ago does not cancel own-side. Ask the other-side prompt. Then look.

A few more limits, said once.

What develops, in Kuhn's map, is meta-knowing: thinking about one's own thought. Absolutism is common in childhood and can last. Multiplism — “that's just your opinion,” anyone's view as good as anyone else's — becomes prevalent at adolescence, and people often remain multiplists for life. Evaluativism, where reasons can be better or worse, is a minority adult achievement.³⁹ There is large individual variation. There is no Piaget stage calendar in this book. There is no critical-thinking gene. There is no window that closed when you were hired. Kuhn and Susan Pearsall are explicit that coordination of theory and evidence may not develop fully even by adulthood.⁴⁰ Kuhn later showed that people often mash a story that *makes a claim plausible* with evidence that the claim is *true*. Substantial variation remains among adults.⁴¹ “Because we needed the quarter”

³⁸Keith E. Stanovich, Richard F. West, and Maggie E. Toplak, “Myside Bias, Rational Thinking, and Intelligence,” *Current Directions in Psychological Science* 22, no. 4 (2013): 259–64. Myside bias shows very little relation to intelligence, including in university samples; in several naturalistic between-subjects designs the association was “literally zero.” High intelligence is no inoculation when no one rings the decoupling bell.

³⁹Deanna Kuhn, “A Developmental Model of Critical Thinking,” *Educational Researcher* 28, no. 2 (1999): 16–25, 46. Meta-knowing develops; many adults remain multiplists, for whom critical thinking feels irrelevant (“everyone has a right to their opinion”); only a minority progress to an evaluative epistemology; people often remain multiplists for life. Adulthood is not arrival.

⁴⁰Deanna Kuhn and Susan Pearsall, “Developmental Origins of Scientific Thinking,” *Journal of Cognition and Development* 1 (2000), as restated in Kuhn, “How Do People Know?,” *Psychological Science* 12, no. 1 (2001): 1–8. Simple theory–evidence distinction is largely in place by middle childhood and is clean in adults on a picture–sequence task; complex coordination remains incomplete; “with development, we see improvement but by no means mastery, even by adulthood.” Original 2000 pagination not independently line-checked for this chapter.

⁴¹Kuhn, “How Do People Know?,” *Psychological Science* 12, no. 1 (2001): in supporting a claim, people mash explanation of why a claim makes sense with evidence that the claim is true; substantial variation remains among adults.

is that mash, wearing work clothes.

Peter Facione’s Delphi report is expert consensus. The California Critical Thinking Skills Test built from that construct is a college-calibrated timed test: about fifty-five minutes, forty scenario items, designed for undergraduate and graduate students.⁴² A score is a score. It is not this book’s meaning. It is not the gym. Some programmes sit it. Sitting it is a sitting. Watson–Glaser is the same kind of object: an instrument some graduate schemes and employers use. Neither score is the capacity to decide what to believe or do with a recap and the notes. Paul and Elder’s apparatus is a widely circulated teaching brand. It is not a meta-analysis. Asking “what is the purpose of this email?” does not require a licensed wheel.

If a younger book in this house already named five moves for ages 5–10, a middle book named two-source, graph, counterexample, and steelman for ages 11–15, and a late-high-school book named a short written reason a stranger could follow, those moves still work. This book is not a reprint of those lesson banks. What this book *keeps* is the written reason, moved onto the job: a colleague who missed the meeting could follow it. One pointer. Then teach at the desk and on the floor.

Wrong answers you should be able to hear

These are illustrations, not portraits of real workplaces.

1. **“Because we needed the quarter.”** This is a why-story offered as a check. Hear it. Name it kindly: that is a reason the result *might* have happened. Then: how do you know we hit the number? Point at the notes, the system of record, last week’s tickets.
2. **“The recap said we hit it.”** Fluency is being offered as a look. One speaker repeated a verb. Ask what *hit* is being asked to mean, and what a second place would show.
3. **“That’s just their take.”** Multiplism, said out loud. Two accounts of the same hour can both be “takes,” and one of them can still be better supported by the notes. Holding the line is a judgment under criteria, not a truce in

⁴²Peter A. Facione, *Critical Thinking: A Statement of Expert Consensus for Purposes of Educational Assessment and Instruction* (Delphi Report), 1990: expert consensus (46 panelists), not a sampling of working adults. The California Critical Thinking Skills Test: Insight Assessment product page, accessed 2 September 2026 — designed for undergraduate and graduate students; 55 minutes; 40 scenario-based questions. A score is a score. Neither Delphi nor CCTST is this book’s meaning of critical thinking.

which no reasons apply.

4. **“Critical thinking means finding what’s wrong with people.”** The English word *critical* can sound like that. Ennis’s concept is not negative. The job is deciding what to believe or do, with reasons, and doing justice to the other side.
5. **A finished-looking paragraph with no named source, offered as already done.** Topic sentences are not a check. A model’s paragraph is not a check. Ask: could a colleague who missed the meeting follow how you know?

Five-minute warm-up

Before the try, take one claim already in the job: a recap, a posted schedule, a sentence in an email, a vendor page. Ask yourself, silently, “How do I know?” Notice whether your first answer is a feeling, a hope, a fluent verb, or a look. If it is a feeling, name a second place you could look. You are not deficient for having a first why-story. You are practicing the ear this week needs. That is what you can learn this week: Ennis’s sentence in ordinary language, and how to hear “because we needed the quarter” as a why-story, not a look at the notes.

How to try it this week

The work-hour chapter in the front of this book is the full version: session shape, meeting shape, desk try-it, floor try-it, written-reason shape, and the AI box once. Here is the shape scaled to this idea.

Session shape

Warm-up (2 minutes). Put the object on the table. A recap and, if you have them, the notes. Or posted hours and a clock. Or one email plus a second named source. Say, to yourself or to a named colleague: “We’re going to decide what to believe about this sentence, and how we know.”

Short model (2 minutes). You take the first claim. Out loud: “The recap says we hit the number. How do I know? I needed it to be true. That is a why-story. The notes are a look.” Stop. Do not finish the next person’s reason for them.

Your attempt (5–8 minutes). You take the next sentence, or a colleague takes it. Ask one good question. Wait.

One good question. The exact wording you can say this week:

“How do you know we hit the number?”

If the object is not a recap, keep the shape: “How do you know we decided that?”
 “How do you know the hours match?” “How do you know the email is a decision already made?”

Mixed practice (5 minutes). One more claim, smaller. If the first was the recap, the second can be one line of an email, or the clock against the posted card.

A written line (2–4 minutes). One spoken sentence, then, if the claim can bear it, one written line: *I claim X because Y. We checked by Z. Dated. Stimulus named.* A colleague who was not in the meeting could follow it. This book names that shape here. A later chapter owns the grain. For this week, one line is enough.

Exact wording you can say in a meeting or at a handoff

- “How do you know we hit the number?”
- “What do you mean, *hit*?”
- “Did you see, or did someone tell you, or do you hope so?”
- “What else could be true?”
- “What’s a second place we could look?”
- “You may change your mind. That is not losing.”
- “Can you write that so a colleague who missed the meeting could follow it?”

Ask one. Wait. A slow three is a convention, not a sacrament. Stahl called this think-time; the origin is classroom instruction.⁴³ Rowe’s two waits come from science class. They are not a workplace trial. After they stop, wait again. If they are mid-reason, do not cut them off. Look at the notes, the roster, or the clock, not at the person, if the silence is hard. At work the shrug often looks like refusal. Do not rescue it. After “write the sentence,” that is task-completion time, not a three-second wait. Wait in private when the power difference is steep. Interrogating a junior colleague in front of a vice-president is not think-time.

First try-it

⁴³Robert J. Stahl, “Using ‘Think-Time’ and ‘Wait-Time’ Skillfully in the Classroom,” ERIC Digest ED370885 (1994): think-time; about 3 seconds as a convention; classroom-origin. Mary Budd Rowe, “Wait-Time and Rewards as Instructional Variables,” *Journal of Research in Science Teaching* (1974; NARST 1972): science-class wait; typical pauses under a second and a half; 3–5 seconds changes discourse. Not a workplace RCT. Rowe and Stahl are two papers, not one.

Hit-the-number, desk, 5–10 minutes. A recap already in the inbox, or memory of *this* meeting plus the notes if you have them. Opening: “How do you know we hit the number?” The skill is deciding what to believe. Wanting is a why-story. The notes and the system of record are evidence *that*.

How to fade help

First time: you model the why-story and the look. Second time: you ask the question and wait. Third time: you ask yourself “how do I know?” before anyone else does. If a colleague still offers a why-story, name it once, then point at the object. If they write a fluent paragraph with no check, ask for the line a colleague who missed the meeting could follow. You may name a better reason if the claim is false. Hearing a wrong reason, then asking a better question, then naming a fact if needed, is the work. Silence in the face of a false reason is not kindness. “Never correct a colleague” is not this book’s rule.

If a tool is in the room, it waits. Try unaided first. A labelled SCRIPT after you tried is a prompt for *you* to say, not a paragraph to paste as your reason. The work-hour AI box holds the full allow and ban list. Point back. This chapter is not an AI policy paper.

When to stop talking

When a reason has been offered, one other thing has been looked at, and the claim has been kept or changed. Ten to twenty-five minutes of a named move, a few times a week, on actual job artefacts, is the right ambition. If the table is becoming a fight, stop. Postpone. Pick a smaller claim. The hour is triage.

Desk

Fluent own-side reasons in domains you know are expected, not a failure. Dual perspective *when asked* is the win for the week. A spoken sentence plus one written line that names the claim and the check is done enough. Recap versus notes. An email versus the calendar or the purchase order. A vendor page versus the spec — named, not yet the whole two-source chapter. You prompt “what would someone who disagrees say?” You will probably have to keep prompting. Intelligence will not do it for you. A degree will not do it for you.

Floor-or-shift

Same move, different jacket. Posted hours versus the clock. Hold the line or call the vendor, with criteria (spec, count, clock), not volume. Two handoff notes

after, not during, a live procedure. Safety is ordinary: no public quiz in front of a supervisor; no undercover audit; no gotcha against a colleague, a vendor, or a shop. Noticing is not a sting. A clinic try-it uses a protocol already in use — not a diagnosis of a named patient in this book. Dual focus is still the work. Myside is still live. Evaluativism is still not a group norm. Same moves. Different grain. Neither band is “too practical to reason.” Neither band is “ready to skip practice because they have the job.”

Practice that actually builds it

Desk, floor, clinic, shop *are* the practice. They do not replace the named move. Each try-it below names time, materials, safety, and the move. At least one lives at the desk. At least one lives on the floor or shift.



Figure 4: A recap page beside notes, both with unreadable type, a generic clock at the edge, no logos, no faces.

Try-it: Hit the number (desk)

Time. 5–10 minutes.

Materials. A recap already in the inbox, or memory of *this* meeting plus the notes if you have them. A dated notebook or a line in the project folder.

Safety. No humiliation of a colleague. Curious, not a courtroom. Skip live-controversy forwards.

The move. Decide what to believe. How do you know we hit the number? Wanting is a why-story. The notes / the system of record are evidence *that*.

Illustration, not a reported workplace: a recap says we hit 12,400. The notes say 11,900, or the system of record is still loading last week. Someone says, “We needed the quarter.” You wait. Then: “How do you know we hit the number?” They look. They change the sentence, or they keep it and name a different check. Either is thinking if the look happened.

Try-it: Recap versus notes (desk)

Time. 5–10 minutes.

Materials. The meeting recap and the notes from the same hour.

Safety. Ordinary reading. You are not putting a colleague on trial for their adjectives.

The move. Two records of the same hour. The recap is often a theory of why the meeting went as it did. The notes are closer to evidence of what was said. “How do you know we decided X?” answered with “because it fits the client” is a why-story. Point at the notes.

This is blocked practice: the move is named, the two pages are already in your hand. Mixed practice comes later, when you have to notice that *this* recap is like last Tuesday’s recap.

Try-it: A claim in an email (desk)

Time. 8–15 minutes with a message you were already answering.

Materials. One email plus a second named source: the notes, the calendar, the purchase order, the original client message.

Safety. Ordinary reading. Skip live-controversy forwards. Do not paste confidential text into a public consumer tool.

The move. Ennis believe-or-do in ink. Is this a request, a rumor, or a decision already made? A fluent paragraph can mash those three. Kuhn's mash wearing a finished sentence.

Try-it: Hold the line or call the vendor (floor-or-shift)

Time. 5–10 minutes, when it actually happens, or a calm replay of *this* bay / *this* SKU.

Materials. The actual decision. The spec, the count, or the clock — whichever criterion is already in use.

Safety. No public quiz. No gotcha against a vendor on the floor. Wait in private if the power difference is steep.

The move. Decide what to do, with criteria, not volume. Prompt: what would someone who disagrees say. Then, if ready: write four sentences that hold that point. Claim, because, check, and the other side.

Try-it: Posted hours versus the clock (floor-or-shift)

Time. 3–8 minutes of a shift you were already working.

Materials. The posted schedule, then the actual clock.

Safety. Ordinary looking. This is not a labor-politics seminar and not a sting. You are comparing two records of the same interval.

The move. Criterion versus observation. The card says 7–3. The clock says something else. How do you know the hours match? Wanting the card to be right is a why-story. The clock is a look.

Talk / meeting box

Opening question. “How do you know we hit the number?”

Not: “What is critical thinking?” Not: “Did we have a good meeting?”

Follow-ups (pick three to five, not all at once).

1. What do you mean, *hit*?
2. Did you see, or did someone tell you, or do you hope so?
3. What else could be true?
4. What’s a second place we could look?
5. Can you write that so a colleague who missed the meeting could follow it?
You thought X; we looked; you can say the new thing.

How to wait. After you ask, a slow three. Stahl, labelled classroom-origin. After they stop, wait again. If they are mid-reason, let them finish. Look at the notes, the graph, the roster, or the clock, not at the person, if the silence is hard. At work the shrug often looks like refusal; leave it. After “write the sentence,” that is task-completion time. Wait in private when the power difference is steep. Rowe 1972 is science class, not a stand-up RCT.

What a stuck silence usually means. They offered a why-story where a look was asked for. Wait-time 1 was zero. They are searching for a word. They are guessing what you want. The question was a recitation item in disguise. They think “critical” means finding fault. They think listening means giving up the starting spot. They think a finished paragraph is already a check.

Your job is to pose, wait, and point back at the object — not to fill the silence with the right answer, and not to turn the meeting into a fight.

One incorrect example to diagnose. A fluent recap with no notes, offered as a check. “We hit it. The recap is done. I already wrote it up.” What it usually means: testimony is doing the work of a look, and length is doing the work of a reason. What to say: “How do you know we hit the number?” Then open the notes. If the notes go the other way, correct the sentence. That is the work. It is not losing.

If it isn’t clicking

Three diagnostics. No shame. The next move is small.

1. Every answer is a why-story. “Because we needed it.” “Because the client was happy.” “Because that’s how we always close.” You asked how they know it happened. They told you why it would make sense. Slow down. Name the

mash once, in ordinary language: that is a story about *why*. I am asking *whether*. Point at one object — notes, clock, count — and wait. If they still mash, model on a smaller claim: posted hours versus the clock. “I hoped we started at seven. I looked. The card says seven. The clock said 7:12.” Then return to the recap. Going ahead while the mash is still invisible just writes the mash in ink.

2. They (or you) treat the job, the degree, or fluency as arrival. “I’ve been here eleven years.” “I already think this way — it’s how I am.” “College covered this.” Hear it. The next move is not a lecture about intelligence. It is a look. Stanovich’s finding is the adult one: myside barely moves with how smart you are. Felton and Kuhn’s adults were more strategic *and* still own-side. Willingham’s bicycle is the other half: years in *this* domain are not a bicycle for *that* artifact. Ask one question on an artifact they already have. If the first “because” is still own-side, you have not failed. You have found this week’s work. Prompt: “What would someone who disagrees say?” Then look. If a human colleague who was in the room can name the notes, ask them — not as a quiz, as a second speaker.

3. “Critical” still means finding fault, or “that’s just their take” has made checking feel rude. Slow down on meaning. Ennis’s concept is not negative. Kuhn’s multiplist adult is the person for whom evaluation feels pointless: everyone has a right to their opinion, so reasons do not apply. The next move is a cheap look, not a philosophy class. Two records of the same interval. Posted hours and the clock. Recap and notes. If they disagree, you do not have to pick a team tonight. You do have to notice that one can be better supported. If the claim is false, ask a better question, then name the fact. Correct the reason without crushing the talk. Ask a human colleague when the object is theirs and you are guessing. A tool can diagnose a reason you already gave; you hear the diagnosis and decide. A tool does not get the first attempt.

When to go ahead: you can hear a why-story, ask “How do you know?”, look at one other thing, and change the sentence if the look goes the other way — at the desk and on the floor. When to stay: the first answer is still volume, fluency, or “because we needed it,” and no look has happened.

Tools, including AI

Optional. Keep this short. The work-hour AI box is the full list.

You talk to the page, to a named colleague, and to a second source. You hold the

account. A chatbot is never the only partner. COPPA’s under-13 line does not apply to this reader. That is not a thinking birthday and not a licence for unsupervised chat.

After you have tried, you may: explain this recap *to yourself* from a named source (the notes, the spreadsheet tab, the clause); ask for extra questions, then vet them; use a labelled SCRIPT after you tried — a short sequence for *you* to say, not a paragraph to copy as your reason; ask for a hint after an attempt; ask a tool to diagnose a reason you already gave.

You may not: park an unsupervised chatbot as your only partner; ask the tool to write your reason, email, or steelman and paste that text as yours; invent an expert identity and treat that voice as a person; grade a colleague from a detector score; print a certificate of thinking; ask the tool to invent what the notes say; hand yourself an agent as a thinking partner for the first attempt; put confidential workplace data into a public consumer tool.

A fluent reason from a machine is performance, not your check. On some consulting tasks inside one model’s frontier, adults using the model were faster and higher-rated; on a designed outside-frontier task they were less likely to be right.⁴⁴ A fluent memo is not a check that the task was inside the frontier. Open *your* employer’s acceptable-use page. Label it. Then return to the recap and the notes.

What “done enough” looks like

Before you move on, you should be able to hear a why-story and ask “How do you know we hit the number?” without turning a stand-up into a courtroom. You should be able to take a recap, posted hours, an email claim, or a floor decision, say how you know, look at one other thing, and change the sentence if the look goes the other way.

Checklist

- You can say Ennis’s sentence in ordinary language: deciding what to believe or do, with reasons, doing justice to challenges.

⁴⁴Fabrizio Dell’Acqua et al., “Navigating the Jagged Technological Frontier,” Harvard Business School working paper (2023): 758 BCG consultants; inside-frontier GPT-4 faster and higher-rated; outside-frontier 19 percentage points less likely correct. Full 58-page working paper not line-checked this stream. One workplace-help sentence, then teach. Point back to the work-hour AI box.

- You can hear “because we needed the quarter” as a why-story, not a look at the notes.
- You can answer “How do you know we hit the number?” with something other than volume.
- There has been at least one desk try (recap versus notes, or an email plus a named source) and at least one floor-or-shift try (posted hours versus the clock, or hold-the-line with a criterion).
- A spoken sentence plus one written line that names the claim and the check, when asked for the other side.
- No one is treating a job, a degree, years of experience, or a workshop certificate as a thinking diploma.
- A “critical thinking workshop certificate” is a receipt, not the work. Placement is by skill, not job title. If you still live in own-side exposition and need the prompt “what would someone who disagrees say?”, you are still in this chapter’s early band of the move, whatever your title.

The title of the year’s work, if a title is needed at all, is still the job you hold — analyst, nurse, clerk, lead, tech, manager — not Critical Thinking I. That course does not exist as a job.

You understand this. You know what to say tomorrow. You know what an own-side “because” sounds like when it is not yet a check. You know what a finished-looking paragraph without a named source looks like. You can do this at the desk and on the floor.

Chapter 2 — Moves that scale

Critical Thinking for Adults

Michael Gannotti

2 September 2026



Figure 5: A still-life from a high three-quarter view: a paper map of three boxes with unreadable labels, a recap page and notes with greeking, a generic clock, a short dated notebook line. No people. No logos. No readable UI.

Why this matters

What do you mean, *always*?

A colleague said it. A recap said it. Last quarter’s process, the starting SKU, the closing time, the “industry-leading” on the front of the page. *Always* is a small word doing a large job. If you let it pass, you get a universal where a habit would have been enough. If you ask what it means, the claim has to come into the room as a picture you can check.

This chapter names the teachable object. Not a general faculty. Not a reprint of the 5–10 picture-bank, the 11–15 two-source week, or the 15–18 recap week. A small set of social-cognitive moves, practiced on material you already somewhat

know.

Five that scale from the younger books: **wait**; **what do you mean**; **how do you know**; **what else could be true**; **how could we check**. Four the middle book added: **two-source** (“what’s a second place we could look?”); **graph** (“what does this chart actually show?”); **counterexample**, which lives inside “what else / always?” — not a ninth move to reprint as a bank; **steelman** (“what’s the strongest version of what they mean?”). One this book keeps, named out loud here, owned later: **write the reason so a colleague who missed the meeting could follow it**.

The idea unlocks a week in which you close your mouth. One good question beats five. Implicit hope that “we value thinking” will not do. Philip Abrami’s 2008 review put it plainly: improvement cannot be a matter of implicit expectation.⁴⁵ That paper mixed ages. The principle still disciplines a stand-up. Name the move. Practice it. Ask one. Wait.

The struggle is worth it because ordinary work is fluent. Fluency hides the mash. A list of three “reasons” can still be three why-stories. A steelman that is actually a straw is still a straw. A finished-looking sentence can still have no check. The named moves make the mash audible before it hardens into a paragraph, a recap, or a slide.

If you only remember one thing from this chapter: skip the wait, skip the thought. At work the shrug looks like refusal. Leave it.

Understand it yourself

Here is the everyday picture.

Handoff. Someone says, “We always start with X.” You could argue. You could list last week’s tickets. You could lecture about universals. Or you could ask, “What do you mean, *always*?” and wait. Two pictures of *always* will usually

⁴⁵Philip C. Abrami, Robert M. Bernard, Evgueni Borokhovski, Anne Wade, Michael A. Surkes, Rana Tamim, and Dai Zhang, “Instructional Interventions Affecting Critical Thinking Skills and Dispositions: A Stage 1 Meta-Analysis,” *Review of Educational Research* 78, no. 4 (2008): 1102–34. “Improvement in students’ CT skills and dispositions cannot be a matter of implicit expectation.” Mixed educational levels. Named here as the principle, not as a stand-up result. Abrami et al., “Strategies for Teaching Students to Think Critically: A Meta-Analysis,” *Review of Educational Research* 85, no. 2 (2015): overall $g+ = 0.30$; graduate and adult $k = 17$, $g = 0.21$, confidence interval includes zero; combination $g = 0.57$ is 19 mixed-age studies, not a workplace cell.

appear: every single ticket this month, or often enough that it feels like every ticket. Those are different claims. Only one of them dies when you name Tuesday.

A second everyday picture, still ordinary. Stand-up. You ask how they know we decided X. Three fluent reasons arrive, all why-stories. A list is not a look. One good question, then a wait, then the notes.

Here is the precise picture.

The five openings are requests, not a costume. “What do you mean?” is Ennis’s clarifying question, the one you can hear: Why? What do you mean by...? What would be an example? What would *not* be an example?⁴⁶ “How do you know?” keeps Kuhn’s two jobs from collapsing: evidence *that* it happened, versus a story about *why* it might have.⁴⁷ “What else could be true?” is Lipman’s alertness to alternatives, said without a brand. “How could we check?” is the cheap look. Wait is the silence that lets any of those land.

Mary Budd Rowe named wait-time in science class: the pause after you ask, and the pause after they finish. Typical classrooms were under a second and a half. When the pauses lasted about three to five seconds, talk changed.⁴⁸ Robert Stahl later preferred *think-time*: a distinct period of uninterrupted silence so that both people can finish processing. Stahl listed eight categories of silence, including the pause after a question, the pause after a student stops, and the pause inside a response. About three seconds is a convention, not a sacrament.⁴⁹ Rowe is science class. Stahl is classroom instruction. They are two papers, not one. They are not a stand-up trial. They still scale as a labelled borrow. Your urge to fill the silence is

⁴⁶Robert H. Ennis, “The Nature of Critical Thinking” (2013, last revised September 2015), and “A Logical Basis for Measuring Critical Thinking Skills,” *Educational Leadership* 43, no. 2 (October 1985): 44–48. Clarifying questions a working adult can hear: Why? What do you mean by...? What would be an example? What would not be an example? The long outline is not this book’s syllabus.

⁴⁷Deanna Kuhn, “How Do People Know?,” *Psychological Science* 12, no. 1 (2001): 1–8. People mash explanation of why a claim makes sense with evidence that the claim is true; substantial variation remains among adults.

⁴⁸Mary Budd Rowe, “Wait-Time and Rewards as Instructional Variables, Their Influence on Language, Logic, and Fate Control,” *Journal of Research in Science Teaching* 11, no. 2 (1974): 81–94 (NARST 1972). Science-class wait; typical pauses under about 1.5 seconds; 3–5 seconds changes discourse. Not a workplace RCT.

⁴⁹Robert J. Stahl, “Using ‘Think-Time’ and ‘Wait-Time’ Skillfully in the Classroom,” ERIC Digest ED370885 (1994). Think-time; eight categories of silence; about 3 seconds as a convention; classroom-origin. Rowe and Stahl are two papers, not one.

often stronger at work, because a shrug looks like refusal and a supervisor is not a science teacher. Wait in private when the power difference is steep. Look at the object if the silence is hard.

The additions are the same family, a later grain, moved onto the job.

Two-source: leave this speaker. A recap and the notes are two speakers. Two quotes for the same SKU are two speakers. Two tabs from the same vendor site are often one speaker repeating themselves. Chapter 4 owns the fuller leave-this-page. This week you name the question: “What’s a second place we could look?”

Graph: a picture is a claim made with ink. “What does this chart actually show?” Title, axes, units, what is missing. Chapter 5 owns the three readings. This week you name the question so it is in the set.

Counterexample lives inside “what else / always?” A universal is a claim that one case can wound. “Always? Can we think of a time it wasn’t?” Upgrade to *usually* if that is what they meant. That is not relativism. That is precision. Stay on ordinary stakes: the process, the SKU, the closing time, last quarter’s document. Live controversy does not enter as “practice for alternatives.”

Steelman: say the strongest fair version of what they mean, then answer *that*. Dual-perspective in speech. There is no trial titled Steelman. There is no What Works Clearinghouse intervention by that name. Inventing a better argument they did not make, then knocking it down, is not steelman. Asking a machine to write the steelman so you can paste it is write-my-reason. Own-side “because we needed the quarter” is myside named in the room. Felton and Kuhn’s adults were already more strategic than teens, and still own-side.⁵⁰ Stanovich’s myside finding still holds: intelligence is no inoculation.⁵¹

The written reason, named here: **I claim X because Y. We checked by Z.** Optionally: **Someone who disagrees would say...** Dated. Stimulus named. Sources named. A colleague who was not in the meeting could follow it. Chapter 6 owns the grain. Kuhn and Crowell already showed that extra essay writing without an interlocutor is not the intervention — ages 11–12 at entry, labelled, copy the shape,

⁵⁰Mark Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, nos. 2–3 (2001): 135–53. Adults more strategic than seventh- and eighth-graders, still own-side. Lab topic is not this book’s exercise type.

⁵¹Keith E. Stanovich, Richard F. West, and Maggie E. Toplak, “Myside Bias, Rational Thinking, and Intelligence,” *Current Directions in Psychological Science* 22, no. 4 (2013): 259–64. Myside barely related to intelligence, including university samples.

not the topics.⁵² Talk first. Write. Then look. If the sentence cannot survive the check, correct the sentence.

One good question beats five. Stack five questions, get a recitation. Perform Socrates, get a performance. Ask one. Wait. A Socratic costume is tempting at work because everyone can already talk. Talking is not the same as a named move. Five questions in a row, each sharper than the last, trains people to hunt for the answer you are withholding. One authentic question about meaning, knowledge, a second source, a chart, a time it wasn't, a steelman, or a written check trains them to think about the object. This book takes the second. You are sitting beside a colleague. You are not performing a seminar.

The younger books in this house already practiced wait, meaning, how-do-you-know, what-else, and how-could-we-check on a swing, a cereal box, a park sign. The middle book already practiced two-source, graph, counterexample, and steelman on a recap, a spec sheet, two storm pages, a rainfall graph at 11–15. The late-high-school book already named a short written reason a stranger could follow, with a parent in the room. Those objects still work if they walk through the door. They are not this week's centre. This week's centre is the same family of moves on *always*, a process claim, two quotes for a SKU, a recap versus notes, and one sentence a colleague who missed the meeting could follow. Naming the family is not reprinting the bank.

Hearing a wrong reason, then asking a better question, then naming a fact if needed, is the work. "Never correct a colleague" is not this book's rule. Kuhn's classroom work evaluated reasons, generated counters, and dropped weak cards. That is correction. Correct the reason. Stay accountable to the person.

Wrong answers you should be able to hear

These are illustrations, not portraits.

1. **A shrug, fast, that you fill.** Wait-time 1 was zero. The shrug often looks like refusal. It is often search. Wait.
2. **"Always means always."** Meaning has not been asked. Two pictures have not been separated. Ask what they mean before you judge.

⁵²Deanna Kuhn and Amanda Crowell, "Dialogic Argumentation as a Vehicle for Developing Young Adolescents' Thinking," *Psychological Science* 22, no. 4 (2011): 545–52. Ages 11–12 at entry; extra essay writing in the comparison group did not produce dual-perspective transfer; dialogic practice did. Copy shape, not topics.

3. **Three fluent reasons, all why-stories.** “How do you know?” was heard as “give me a list.” A list is not a look.
4. **A steelman that is actually a straw.** “So you’re saying we *never* start with anything else, which is ridiculous.” That is not the strongest fair version. Say it so they would nod.
5. **A finished-looking paragraph, or a chatbot paragraph, offered as the reason.** Length is not the job. A machine’s fluency is performance. Ask for the line a colleague who missed the meeting could follow, written after the try, window closed.

Five-minute warm-up

Tonight, pick one small word already in the job: *always*, *hit*, *leading*, *jumped*, *nights*. Ask yourself what two pictures it might name. Then practice a slow three on a question you already ask — a stand-up, a handoff, a vendor call — and do not fill the shrug. That is what you can learn this week: the five openings plus two-source, graph, steelman, and “write it so a colleague who missed the meeting could follow it,” and how to close your mouth for a slow three when the shrug looks like refusal.

How to try it this week

The work hour already licensed session shape, meeting shape, desk try-it, and floor try-it. This chapter fills them with named moves. Practice one. Naming nine at once is a poster.

Session shape

Warm-up (1–2 minutes). Invisible wait drill. You will not name it to them as a drill. You will close your mouth.

Short model (2–3 minutes). You take one small word already in the room. Out loud: “We always start with X. What do I mean, *always*? Every ticket this month, or often enough that it feels like every ticket? Last week’s tickets are a look.” Stop. Do not finish their reason.

Your attempt (5–10 minutes). One claim. One question. Wait.

One good question. The exact wording you can say this week:

“What do you mean, *always*?”

If the object is not a process claim, keep the shape: “What do you mean, *hit*?”
 “What do you mean, *leading*?” “What do you mean, *jumped*?”

Mixed practice (5 minutes). A cousin of the same word, or a second look: recap versus notes, two quotes for the same SKU.

A written line (2–5 minutes). Claim, because, check. Dated. A colleague who missed the meeting could follow it.

If the claim can bear a picture, three boxes on paper are enough: *claim* / *because* / *unless*. That is a method, not an app. University argument-mapping courses exist; there is no large randomised trial of them, in the developer’s own words.⁵³ Three boxes in a meeting are the type. A semester course is not this week’s hour. Paper before any login.

Exact wording you can say in a meeting or at a handoff

- “What do you mean, *always*?”
- “How do you know?”
- “Always — can we think of a time it wasn’t?”
- “What’s a second place we could look?”
- “What’s the strongest version of what they mean?”
- “Can you write that so a colleague who missed the meeting could follow it?”

Ask one. Wait. A slow three is a convention, not a sacrament. Stahl, labelled. Rowe is science class. After they stop, wait again. If they are mid-reason, do not cut them off. Look at last week’s tickets, the process document, or the clock, not at the person, if the silence is hard. At work the shrug often looks like refusal. Leave it. After “write the sentence,” that is task-completion time. Wait in private when the power difference is steep.

First try-it

Meaning of *always*, floor-or-shift, 3–8 minutes. A colleague’s “we always start with X,” last week’s actual tickets or last quarter’s process document. Opening:

⁵³Tim van Gelder, “Using Argument Mapping to Improve Critical Thinking Skills,” in *The Palgrave Handbook of Critical Thinking in Higher Education*, ed. Martin Davies and Ronald Barnett (New York: Palgrave Macmillan, 2015). University argument mapping; “Unfortunately, no such study has ever been conducted” — no large RCT, his words. Three boxes on paper before any app. Type, not the course, not CCTST.

“What do you mean, *always*?” A counterexample is a time it wasn’t.

How to fade help

First time: you model two pictures of the word, then the look. Second time: you ask the question and wait. Third time: they ask themselves what the word is being asked to mean. If they still offer a universal, name one case that wasn’t, then wait. If they steelman a straw, ask them to say it so the other person would nod. If they write a fluent paragraph with no check, ask for the line a colleague who missed the meeting could follow. You may name a better reason if the claim is false. Hearing a wrong reason, then asking a better question, then naming a fact if needed, is the work.

If a tool is in the room, it waits. Try unaided first. A labelled SCRIPT after you tried is a prompt for *you* to say, not a steelman to paste. Point back to the work-hour AI box.

When to stop talking

When a meaning has been asked, a wait has happened, and either a look or a written line has followed. Ten to twenty-five minutes of a named move is the right ambition. If the table is becoming a fight, stop. Postpone. Pick a smaller claim. Teaching all nine at once is the wrong ambition. Name them. Practice one.

Desk

Wait drill, invisible, inside a meeting you were already in. How-do-you-know on recap versus notes. Steelman, then write one sentence. A paper map of three boxes if the thread is already long. Safety: wait in private when the power difference is steep; do not interrogate a junior in front of a vice-president and call it Rowe. Dual perspective *when asked* is the win for the week. You will probably have to keep prompting. Fluency at a keyboard is not evidence the wait happened. A recap you wrote in ninety seconds can still be three why-stories.

Cue similarity next time the genre recurs. Last Tuesday’s recap does not automatically produce this Tuesday’s look. Near is not far. You will have to say, out loud, “this is like the recap.” Willingham’s analog still binds: people miss a structure they have already solved when the surface is new.⁵⁴

⁵⁴Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator* 31, no. 2 (Summer 2007): 8–19. The 19 percent / 35 percent analog is a word-problem study, not a workplace RCT. Cite it as the constraint that still binds a fluent adult when the surface is new. Cue similarity; near is not far.

Floor-or-shift

Meaning of *always* on a real process claim. Two-source: two quotes for the same SKU, or two handoff notes, after, not during, a live procedure. Same family of moves. Different jacket. Safety is ordinary: no public quiz, no humiliation, no gotcha against a colleague, a vendor, or a shop. The wait still applies on the floor. A fast answer at a handoff is still an answer. It is not better because it was first. If the bay is live, ask after, not during. Noticing is not a sting.

Practice that actually builds it

Desk, floor, clinic, shop *are* the practice. They do not replace the named move. Each try-it below names time, materials, safety, and the move. At least one lives at the desk. At least one lives on the floor or shift.



Figure 6: A paper map of three boxes beside a recap and notes with unreadable type, a generic clock at the edge. No faces. No logos.

Try-it: Wait drill (desk, meeting)

Time. Invisible. 3–8 minutes of a meeting you were already in.

Materials. One question. The object already on the table — a recap, a deck chart, a process claim.

Safety. Wait in private when the power difference is steep. Do not interrogate a junior colleague in front of a vice-president. Curious, not a courtroom.

The move. Wait-1 after you ask. Wait-2 after they stop. A slow three is a convention, not a sacrament. Stahl, labelled classroom-origin. Rowe is science class.

This is blocked practice: you already know the question; the only new move is the silence. Mixed practice is later, when you have to notice that *this* shrug is search, not refusal.

Try-it: How do you know, recap versus notes (desk)

Time. 5–10 minutes.

Materials. The recap email and the notes from the same hour. A dated notebook.

Safety. Ordinary reading. You are not putting a colleague on trial for their adjectives.

The move. How do you know we decided X? The recap is often a why-story of the hour. The notes are closer to a look. Three fluent reasons are still three why-stories if none of them is a look.

Try-it: Three boxes on paper (desk)

Time. 8–15 minutes.

Materials. A disagreement already in the thread, or a recap that already has a “because.” A pad. Three boxes: claim, because, unless.

Safety. Ordinary. You are mapping a sentence, not grading a person. The map stays inside the job’s real file, not under a course name.

The move. Write the claim in one box. Write the because in the next. Write the strongest fair objection in the third. Then look. If the because cannot survive the unless, correct the because. This is blocked practice for the written-reason shape.

Mixed practice is later, when you have to notice that *this* thread needs a map and *that* handoff only needs a line.

Try-it: Steelman, then write one sentence (desk)

Time. 8–15 minutes.

Materials. A disagreement already in the thread. A dated notebook or a line in the project folder.

Safety. Say theirs so they would nod. Inventing a better argument they did not make, then knocking it down, is not steelman. Pasting a machine's steelman as yours is write-my-reason.

The move. Dual-perspective in speech, then claim + because + check in ink. A colleague who missed the meeting could follow it. Optional second sentence: someone who disagrees would say...

Illustration, not a reported workplace: a thread says we should switch vendors because last month's tickets were slow. You say, "You mean last month's tickets were slow, and that is a reason to look at the contract clause." They nod. Then you write: I claim we should open the clause because last month's tickets ran past the named window. We checked by the ticket log and clause 4.2. Someone who disagrees would say the window was always informal. Dated. Sources named.

Try-it: Meaning of *always* (floor-or-shift)

Time. 3–8 minutes.

Materials. A colleague's "we always start with X." Last week's actual tickets, or last quarter's process document.

Safety. No humiliation. You are asking what the word means, not catching someone out.

The move. What do you mean, *always*? A counterexample is a time it wasn't. Upgrade to *usually* if that is what they meant.

Try-it: Two-source on the floor (floor-or-shift)

Time. 5–10 minutes.

Materials. Two quotes for the same SKU, or two handoff notes after, not during, a live procedure.

Safety. No public quiz. No undercover audit. No lecture to a stranger's customer. A clinic try-it uses a protocol already in use — not a diagnosis of a named patient in this book.

The move. What's a second place we could look? Two speakers, not one speaker twice. Chapter 4 owns the fuller leave. This week you name the question and do the look.

Talk / meeting box

Opening question. “What do you mean, *always*?”

Not: “Did we have a good meeting?” Not: “What is the vocabulary word?”

Follow-ups (pick three to five, not all at once).

1. How do you know?
2. Always — can we think of a time it wasn't?
3. What's a second place we could look?
4. What's the strongest version of what they mean?
5. Can you write that so a colleague who missed the meeting could follow it?

How to wait. After you ask, a slow three. Stahl, labelled. After they stop, wait again. If they are mid-reason, let them finish. Look at last week's tickets, the process document, the notes, or the clock, not at the person, if the silence is hard. After “write the sentence,” that is task-completion time. Wait in private when the power difference is steep.

What a stuck silence usually means. They offered a why-story where a look was asked for. Wait-time 1 was zero. They are searching for a word. They are guessing what you want. The question was a recitation item in disguise. They think “critical” means finding fault. They think steelman means giving up. They think a finished-looking paragraph is already a check.

Your job is to pose, wait, and point back at the object — not to fill the silence with the right answer, and not to turn the meeting into a fight.

One incorrect example to diagnose. A steelman that is actually a straw. “So you're saying we *never* follow the process, which is ridiculous.” What it usually

means: dual-perspective was heard as a fight, and the weakest version was the one that got answered. What to say: “What’s the strongest version of what they mean? Say it so they would nod.” Then look at last quarter’s process document. If the sentence cannot survive the look, correct the sentence.

If it isn’t clicking

Three diagnostics. No shame. The next move is small.

1. Wait-time 1 is still zero. You ask. You fill. The shrug never had a chance to become a sentence. Slow down. Practice the wait drill on a question you already ask, in a meeting you were already in, without naming it as a drill. Look at the object if the silence is hard. If the power difference is steep, wait in private. Going ahead with five sharper questions trains recitation. One question, then a slow three, is the next move. Ask a human colleague to watch you once — not to score you, to tell you whether you filled.

2. Every “always” is still a universal, or every steelman is still a straw. Meaning has not been asked, or dual-perspective has been heard as giving up. Slow down on the word. Two pictures: every ticket, or often enough that it feels like every ticket. One case that wasn’t is enough to upgrade to *usually*. For steelman: say theirs so they would nod, *then* answer that. If they think steelman means surrender, name it: you are making their claim as strong as it can fairly be, so the look is fair. You are not awarding the vendor. You are not giving up the starting spot.

3. Three fluent reasons, or a finished paragraph, offered as the check. “How do you know?” was heard as “give me a list.” Length is doing the work of a look. Slow down. Point at one object. Recap versus notes. Two quotes for the same SKU. Last quarter’s process document. If they paste a model’s paragraph as theirs, the next move is the closed window: one line they write, claim + because + check, after the try. A tool can diagnose a reason they already gave; they hear the diagnosis and decide. A tool does not get the first attempt. If the claim is false, ask a better question, then name the fact. Correct the reason without crushing the talk.

When to go ahead: you can ask one good question, wait, hear a why-story as a why-story, and either look or write a line a colleague who missed the meeting could follow. When to stay: the shrug is still being rescued, *always* is still passing, or the steelman is still a straw.

Tools, including AI

Optional. Keep this short. Point back to the work-hour AI box.

After you have tried, you may: explain this process claim *to yourself* from a named source (last quarter’s document, the ticket log, the clause); ask for extra questions, then vet them; use a labelled SCRIPT after you tried — a short sequence for *you* to say, not a steelman to paste; ask for a hint after an attempt; ask a tool to diagnose a reason you already gave.

You may not: park an unsupervised chatbot as your only partner; ask the tool to write your reason, email, or steelman and paste that text as yours; invent an expert identity; grade a colleague from a detector; print a certificate of thinking; ask the tool to invent last week’s tickets; hand yourself an agent as a thinking partner for the first attempt; put confidential workplace data into a public consumer tool.

A fluent reason from a machine is performance, not your check. On mid-level professional writing, a widely used chatbot cut time and raised grader scores and mostly substituted for effort — many treated writers submitted the first pass.⁵⁵ A faster email is performance. It is not yet your reason. Open *your* employer’s acceptable-use page. Then return to *always* and last week’s tickets.

What “done enough” looks like

Before you move on, you should be able to ask one good question, wait, and hear an own-side “because” as not yet a check. You should be able to take *always*, a recap, two quotes, or a disagreement already in the thread, ask what they mean, look at one other thing, and write a line a colleague who missed the meeting could follow.

Checklist

- You can name the five openings (wait; what do you mean; how do you know; what else; how could we check) plus two-source, graph, steelman, and the written reason — and you practiced one this week, not nine.

⁵⁵Shakked Noy and Whitney Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science* 381, no. 6654 (2023): 187–92. 453 college-educated professionals; time about –40%; quality about +18%. MIT working paper (labelled draft): many treated writers submitted the first pass unedited. One workplace-writing sentence, then teach. Point back to the work-hour AI box.

- You can hear “always means always” as meaning not yet asked.
- You can let a shrug sit for a slow three without filling it.
- You can say theirs so they would nod, then answer *that*.
- There has been at least one desk try (wait drill, recap versus notes, or steelman-then-write) and at least one floor-or-shift try (*always* on a process claim, or two quotes / two handoff notes).
- Rowe is science class in your ear. Stahl is think-time, labelled. Neither is a workplace RCT you are replicating.
- No one is treating a workshop poster of nine moves as the work. Placement is by skill, not job title. If you still fill the shrug, or still answer the weakest version of theirs, you are still in this chapter’s early band of the move, whatever your title.

A “critical thinking workshop certificate” is a receipt, not the work. The title of the year’s work, if a title is needed at all, is still the job you hold.

You understand this. You know what to say tomorrow. You know what a straw wearing steelman’s clothes sounds like. You know what a finished-looking paragraph without a named source looks like. You can do this at the desk and on the floor.

Chapter 3 — Something to think about

Critical Thinking for Adults
Michael Gannotti
2 September 2026



Figure 7: A still-life from a high three-quarter view: a vendor page with greeking, a spreadsheet tab with greeked numbers, a generic SKU label, a process document with unreadable type. No people. No logos. No readable UI.

Why this matters

How do you know that happened?

Why might it have happened?

Those are two questions, asked separately. Work mashes them. A recap says we aligned because the client wanted it. A handoff says the line stopped because the part was wrong. A dashboard says the number is good because the tile is green. Wanting, blaming, and a color are stories about *why*. Whether it happened is a different job. The notes, the roster, the SKU line, last quarter's process document — those are looks at *that*.

This chapter is knowledge first. You cannot evaluate a vendor claim in a domain you do not know. You can teach a maxim. Without something you actually know, the maxim will not implement.⁵⁶ The something is *your* vendor, *your* spreadsheet, *your* roster, *your* SKU, *your* process document. Not a civic-media packet. Not a generic bias workshop.

The idea unlocks a week in which you stop pretending a fluent page settled a domain you have not yet learned enough to have sides. Two-source does not suspend this. Graphs do not suspend this. A short written reason does not suspend this. “Because the dashboard looked good” is not a because a colleague can check if “good” is undefined.

The struggle is worth it because ordinary work rewards speed, and speed will answer the nearer question. “Why might it have happened?” is nearer. It feels like thinking. It can be true as a story. It is not yet a look. Kuhn’s finding is the adult one: people mash a story that makes a claim plausible with evidence that the claim is true. Substantial variation remains among adults.⁵⁷ Hearing the mash, then asking the other question, is this week’s work.

If you only remember one thing from this chapter: a maxim without something to think about will not implement. Name what you would need to know. Then look, or wait.

You do not need to be a philosopher. You do need to hear “because the part was wrong” as a why-story, and “because the dashboard looked good” as a sentence that still owes a definition. You do need to treat “I don’t know this pocket yet” as the next move, not as a confession.

This is not a science book. Kuhn’s scientific-thinking papers were opened for the mash, not for a lab coat. This is not a sequel to *Critical Thinking for Life* and not a reprint of the younger lesson banks. One pointer. Then teach on the artifact you already hold.

⁵⁶Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator* 31, no. 2 (Summer 2007): 8–19. “You can teach students maxims about how they ought to think, but without background knowledge and practice, they probably will not be able to implement the advice they memorize.” Thinking is not a bicycle skill; processes are intertwined with domain knowledge.

⁵⁷Deanna Kuhn, “How Do People Know?,” *Psychological Science* 12, no. 1 (2001): 1–8. People mash explanation of why a claim makes sense with evidence that the claim is true; substantial variation remains among adults. Epistemic understanding progresses, but few adults achieve understanding of the complementary strengths and weaknesses of evidence and explanation in argument.

Understand it yourself

Here is the everyday picture.

Handoff. Someone says the line stopped because the part was wrong. You could nod. You could add a better story. Or you could ask, first, “How do you know that happened?” — that the line stopped, that the part was the part, that “wrong” means out of spec. Then, separately, “Why might it have happened?” Those are two questions. The first wants a look: the count, the label, the protocol already in use. The second wants a theory. Mixing them is how a theory of the part swallows the look.

A second everyday picture, still ordinary. A vendor page says industry-leading. You do not yet know what the product does in *your* workflow. Sides have not arrived. The maxim “consider both sides” has nowhere to sit. Name what you would need to know. Then either learn it or do not pretend the page settled it.

Here is the precise picture.

Willingham’s constraint belongs in the same sitting as Ennis’s sentence, or the sentence becomes a poster. The processes of thinking are intertwined with the content of thought. Even three-year-olds can think in a known domain. Even trained scientists can fail outside theirs.⁵⁸ You already think in pockets of the job. A job is many micro-domains: receiving is not cycle count is not vendor dispute. Years in *this* pocket are not a bicycle for *that* artifact. Experts encapsulate. They do not become general critical thinkers.⁵⁹

Maxims without knowledge will not implement. “Consider both sides.” “Check the source.” “Look at the axis.” Each maxim can get you to the precipice. Implementing “consider both sides” on two vendor pages requires knowing what the product does. Implementing “check the axis” requires knowing what the quantity is. A half-day of maxims is a half-day of precipice.

Kuhn’s two questions are the other half of the precise picture. “How do you know

⁵⁸Willingham, *American Educator* (Summer 2007). Even 3-year-olds can engage in critical thinking in a known domain; even trained scientists can fail outside theirs. The 19 percent / 35 percent analog is a word-problem study, not a workplace RCT.

⁵⁹Daniel T. Willingham, “How to Teach Critical Thinking,” NSW Department of Education occasional paper (2019). Generic programs have had limited success; transfer absent or modest and quick to fade; experts encapsulate; they do not become general critical thinkers. A job is a domain. It is not all domains.

that happened?” asks for evidence *that*. “Why might it have happened?” asks for a theory. Young children mash them. Adults mash them on complex, multi-feature problems, sometimes denying they ever thought otherwise.⁶⁰ The simple distinction — did the event occur versus why would it — is largely in place by middle childhood, and adults make no errors on a picture-sequence task built for that simple distinction. That is not mastery. Complex coordination remains incomplete even by adulthood.⁶¹ The workplace analog is not a race and a trophy. It is two artifacts that *feel* like the same knowing: a recap and the notes; a spec and a review; two handoff notes; a number and its caption.

A third precise mark, so this chapter does not become a bias workshop. In medicine, diagnostic error is not primarily uncorrected fast-thinking bias. Training people to recognise named biases has not been an effective educational strategy. Reorganising knowledge has small, consistent benefits.⁶² That is medicine, labelled. The workplace translation is not “become a doctor.” It is: do not run a critical-thinking move on a vendor claim in a domain you have not yet learned enough to have sides. A profession that requires “critical thinking” on paper has not, in one college review, outperformed other majors on generic tests.⁶³ Field-specific expertise does not readily transfer. A finite set of reasoning skills as measured by those tests is not a bicycle.⁶⁴

⁶⁰Deanna Kuhn, “A Developmental Model of Critical Thinking,” *Educational Researcher* 28, no. 2 (1999): 16–25, 46. Like preschoolers, many older individuals blur theory-based and evidence-based sources of belief; evidence serves merely to illustrate what one knows to be true; improvement but by no means mastery, even by adulthood.

⁶¹Deanna Kuhn and Susan Pearsall, “Developmental Origins of Scientific Thinking,” *Journal of Cognition and Development* 1 (2000), as restated in Kuhn 2001. Simple theory–evidence distinction is largely in place by middle childhood and is clean in adults on a picture-sequence task; complex coordination remains incomplete. Original 2000 pagination not independently line-checked for this chapter. This book copies the mash, not a science curriculum.

⁶²Geoffrey Norman, Jonathan Sherbino, and colleagues, “The Causes of Errors in Clinical Reasoning: Cognitive Biases, Knowledge Deficits, and Dual Process Thinking,” *Academic Medicine* 92, no. 1 (2017). Diagnostic error is not primarily uncorrected System-1 bias; educational strategies directed at recognising biases are ineffective; knowledge reorganisation has small consistent benefits. Medicine, labelled. Not a white-collar RCT. The workplace translation is knowledge first, not “become a doctor.”

⁶³Christopher R. Huber and Nathan R. Kuncel, “Does College Teach Critical Thinking? A Meta-Analysis,” *Review of Educational Research* 86, no. 2 (2016): 431–68. College, not the job. Nursing (NLN explicit-CT requirement) did not outperform other majors on generic tests. Four-year mixed-design skill gain about 0.59 SD is college inventory change, not a job result.

⁶⁴Nathan R. Kuncel, as restated in Huber and Kuncel (2016): field-specific expertise does not readily transfer; a finite set of specific reasoning skills as measured by critical-thinking tests is not

Kahneman’s dual-process language is a map of fast impression versus slower monitoring. It is not a synonym for critical thinking. System 2 is effortful, capacity-limited, and often lazy. Running it is not the same as deciding well what to believe or do. You can spend slow effort generating own-side arguments. You can spend it on the wrong criterion. You can spend it inside the wrong domain. Intuition in a known domain is generally skilled. The opened sources do not give a workplace rule for when to spend the slower system. This book will not invent one. Attribute substitution is the mechanism worth keeping: the mind will answer a nearby easy question. “Is this vendor good?” becomes “do I already like them?” “Did the count change?” becomes “does this recap feel like what I remember?” That is Ennis’s problem in a different jacket: you decided what to believe by answering a different question.

Gigerenzer’s natural frequencies, if they appear at all this week, are risk communication at work — incident rates, “cuts defects in half,” a relative percent without a denominator — labelled, not a theory of critical thinking. Ask what the number is a number *of*. That already requires knowledge of the quantity.

Two-source, graph, and a short written reason do not suspend knowledge first. Leaving a vendor page for a spec sheet still requires knowing what the spec is a spec *of*. Reading a deck chart still requires knowing the unit. Writing “I claim X because Y; we checked by Z” still requires that Y and Z are about the same thing. “Because the dashboard looked good” is not a because a colleague can check if “good” is undefined.

Wrong answers you should be able to hear

These are illustrations, not portraits.

1. **“The line stopped because the part was wrong.”** A why-story offered where a look was asked for. What it usually means: theory arrived first. What to say: “How do you know that happened?” Then, separately, “Why might it have happened?”
2. **“I don’t know this domain, so I can’t think.”** What it usually means: knowledge-first was heard as a character flaw. It is the next move. Name what you would need to know. Then learn it, or wait, or ask a human colleague who holds that pocket. Pretending the page settled it is the miss.
3. **A bias-word offered as the diagnosis.** “That’s just confirmation bias.”

a bicycle. Named here as the constraint, not as a workplace trial.

“That’s System 1.” What it usually means: a workshop vocabulary arrived without the case. Named-bias training is the wrong intervention in the medical literature this chapter borrows, labelled. Return to the vendor, the SKU, the roster. What would you need to know to have sides?

4. **“Because the dashboard looked good.”** “Good” is undefined. A color is not a criterion. What is this a number *of*? Who is included? Average of what?
5. **A maxim offered as a check.** “We considered both sides.” “We thought critically.” What it usually means: the precipice was treated as the landing. Name the something. Then look.

Five-minute warm-up

Before the try, take one artifact you already use: a vendor one-pager, a spreadsheet tab, a posted roster, a SKU line, last quarter’s process document. Write two columns: *How do I know that happened?* and *Why might it have happened?* Put one look under the first. Put one theory under the second. If you cannot name what you would need to know to have sides, write that sentence instead: “I do not yet know enough about X to have sides.” That is thinking. It is not a defect. That is what you can learn this week: name what you would need to know to evaluate this page, then look or wait; and, on a real handoff, ask the two questions separately.

How to try it this week

The work hour already licensed session shape. This chapter fills it with something to think about.

Session shape

Warm-up (2 minutes). Put *your* artifact on the table. Not a civic packet. Not a list of biases. A vendor page you actually use, a tab you already own, a roster already posted, a SKU already in the bay, a process document already on the wall.

Short model (2–3 minutes). You take the first claim. Out loud: “The line stopped because the part was wrong. How do I know that happened? I have a theory about the part. That is a why-story. The count and the label are a look. Why might it have happened is a separate question.” Stop. Do not finish their reason.

Your attempt (5–12 minutes). One artifact. Two questions, asked separately. Wait between them.

One good question, then the other. The exact wording you can say this week:

“How do you know that happened?”

Then, separately:

“Why might it have happened?”

If the object is a vendor page, keep the shape: “What would I need to know about this product to have sides?” If the object is a spreadsheet tab: “What is this a number *of*?”

Mixed practice (5 minutes). A cousin in the same job: if the first was the vendor page, the second can be the SKU line, or the roster.

A written line (2–4 minutes). Claim, because, check. If “good” or “wrong” or “aligned” is still undefined, the line is not yet a line. Name the something.

Exact wording you can say in a meeting or at a handoff

- “How do you know that happened?”
- “Why might it have happened?” — asked second, not as a substitute.
- “What do you mean by that verb?”
- “Did you see it, or is this a story of why it would make sense?”
- “What would we have to know about this vendor / this SKU / this roster to have sides?”
- “What’s a second place we could look?”
- “Can you write that so a colleague who missed the meeting could follow it?”

Ask one. Wait. A slow three. Stahl, labelled. After they stop, wait again. Look at the roster, the SKU, the process document, or the clock, not at the person, if the silence is hard. At work the shrug often looks like refusal. Leave it. After “write the sentence,” that is task-completion time. Wait in private when the power difference is steep.

First try-it

Your vendor page, desk, 8–15 minutes. One vendor one-pager you actually use. Opening: name what you would need to know to have sides. Then either learn it or do not pretend the page settled it.

How to fade help

First time: you model the two questions separately, then the look. Second time: you ask “how do you know that happened?” and wait. If they offer a why-story, name it once: that is a theory. I am still asking whether. Then ask the second question, later. Third time: they separate the two without the prompt. If they still mash, point at one object you can hold. If they treat “I don’t know this domain” as a defect, name it as the next move. You may name a better reason if the claim is false. Hearing a wrong reason, then asking a better question, then naming a fact if needed, is the work.

If a tool is in the room, it waits. Try unaided first. A tool may not invent what the spec lists, what the roster says, or what the spreadsheet totals. Point back to the work-hour AI box.

When to stop talking

When the two questions have been asked separately, one look has happened, and either knowledge has been named as missing or the sentence has been kept or changed. Ten to twenty-five minutes on a named artifact is the right ambition. If the table is becoming a lecture titled *Knowledge*, stop. Return to the page in hand. If it is becoming a bias workshop, stop. Return to the SKU.

Desk

Your vendor page. Your spreadsheet tab. What is this a number *of*? Average of what? Who is included? A spec you already wrote, versus what you would need to know to have sides on a vendor’s adjective. Dual perspective still needs knowledge of the thing. Prompting “what would someone who disagrees say?” on a product you cannot yet describe produces theatre.

This book will not assign you a domain you do not have. A paperback cannot supervise your job. Willingham’s precipice is aimed at school systems; the honest possibility here is repeated practice with ordinary artifacts *in the job you already hold*. Cue similarity next time the genre recurs. Last month’s vendor page does not automatically produce this month’s look. Near is not far. You will have to say, out loud, “this is like the tab,” or “this is like the roster.”

Floor-or-shift

Your roster, your SKU line, last quarter’s process document. An object you can hold: a part, a box, a label already in the bay. Who was in a position to know?

How do you know that happened? Then, separately, why might it have happened? Safety: ordinary noticing, not a sting. No public quiz. No undercover audit. A clinic try-it uses a protocol already in use — not a diagnosis of a named patient in this book.

The floor is not a different faculty. It is the same two questions in a different jacket. A fast theory at a handoff is still a theory. It is not better because it was first. If the bay is live, ask after, not during.

Practice that actually builds it

Desk, floor, clinic, shop *are* the practice. They do not replace the named move. Each try-it below names time, materials, safety, and the move. Knowledge first: the materials are yours, already in the job.



Figure 8: A generic box and a label with unreadable type beside a process document with greeking, no brand marks, no faces.

Try-it: Your vendor page (desk)

Time. 8–15 minutes.

Materials. One vendor one-pager you actually use. Not a famous civic site. Not a live controversy. A dated notebook.

Safety. Ordinary reading. You are not running an undercover audit of a shop. Skip pages that import a live fight.

The move. Name what you would need to know to have sides. Then either learn it or do not pretend the page settled it. “Industry-leading” is an adjective. The spec, the last-quarter incident log, the thing the team actually does — those are sides, if you know them.

This is blocked practice: the page is already yours; the move is named. Mixed practice is later, when a *new* vendor page arrives and you have to notice that this is like last month's page.

Try-it: Your spreadsheet tab (desk)

Time. 8–15 minutes.

Materials. A metric you already own. One tab. One number that already drives a sentence.

Safety. Ordinary. You are not publishing the tab. Confidential columns stay off public tools.

The move. What is this a number *of*? Average of what? Who is included? “Because the dashboard looked good” is not a because if “good” is undefined. Write the definition in a line a colleague who missed the meeting could follow. If you cannot, knowledge is the next move, not a maxim.

Illustration, not a reported workplace: a tile says 94. Someone says we are green. You ask what 94 is a number of. It is an average of last week's closed tickets, excluding two queues, including a retry that counts twice. The green was a color. The check is the definition. The sentence changes, or it stays and names the queues. Either is thinking if the definition was asked.

Try-it: Your roster, SKU, or process document (floor-or-shift)

Time. 5–10 minutes.

Materials. The posted roster, the SKU line, or last quarter's process document — whichever is already in use.

Safety. Ordinary looking. Not a labor-politics seminar. Not a sting. No public quiz in front of a supervisor.

The move. How do you know that happened? Then separately: why might it have happened? A shift rumor about who works when is a why-story until the roster is opened. “We always do X” is a why-story until last quarter's document is opened. The SKU line is a look at *what* the thing is. A story about why it sold is a different job.

Try-it: Object you can hold (floor-or-shift)

Time. 3–8 minutes.

Materials. A part, a box, a label already in the bay. Not a tasting of unknown things. Not a test you steal from a shelf.

The move. Who was in a position to know? Did you see it, or did someone tell you, or do you have a theory about why it would make sense? Point at the object. Then ask the two questions separately.

Talk / meeting box

Opening questions. “How do you know that happened?” Then, separately, “Why might it have happened?”

Not: “What is the bias here?” Not: “Did we think critically?”

Follow-ups (pick three to five, not all at once).

1. What do you mean by that verb?
2. Did you see it, or is this a story of why it would make sense?
3. What would we have to know about this vendor / this SKU / this roster to have sides?
4. What’s a second place we could look?
5. Can you write that so a colleague who missed the meeting could follow it?

How to wait. After you ask, a slow three. Stahl, labelled classroom-origin. After they stop, wait again. If they are mid-reason, let them finish. Look at the roster, the SKU, the process document, or the clock, not at the person, if the silence is hard. After “write the sentence,” that is task-completion time. Wait in private when the power difference is steep. Rowe is science class, not a stand-up RCT.

What a stuck silence usually means. They offered a why-story where a look was asked for. They think knowledge is a lecture. They think “I don’t know this domain” is a character flaw — it is the next move. Wait-time 1 was zero. They are guessing what you want. They think “critical” means finding fault. They think a finished-looking paragraph is already a check.

Your job is to pose, wait, and point back at the object — not to fill the silence with the right answer, and not to turn the meeting into a workshop on named biases.

One incorrect example to diagnose. A why-story offered as evidence. “How do you know the line stopped?” “Because the part is always wrong on Tuesdays.” What it usually means: the two questions collapsed, and a theory about Tuesdays did the work of a look. What to say: “That is a story about *why*. I am asking *whether*.” Then the count, the label, the protocol already in use. Then, separately, “Why might it have happened?” If the look goes the other way, correct the sentence.

If it isn’t clicking

Three diagnostics. No shame. The next move is small.

1. The two questions keep collapsing. Every “how do you know?” is answered with a why. Slow down. Name the mash once, in ordinary language. Ask only the first question this sitting. Point at one object you can hold. Do not add the second question until a look has happened. Going ahead with both at once writes the mash in ink. If they stare, model on posted hours versus the clock: “How do I know we started at seven? I looked at the card and the clock. Why we started late is a separate question.” Then return to their artifact.

2. Knowledge is being treated as a lecture, or as a defect. “I don’t know this domain” arrives as shame, or as a reason to skip the look, or as a reason for you to deliver a seminar. It is the next move. Name what you would need to know to have sides. Write that sentence. Then either learn the pocket — a colleague who holds it, last quarter’s document, the spec — or wait. Pretending the vendor page settled it is the miss. Asking a human colleague is allowed. A tool may explain a *named* page to you after you tried; it may not invent the spec.

3. A bias workshop has entered the hour. Named biases, System-1 as an insult, a list of ways people are fools. Slow down. Return to the artifact. The medical literature this chapter borrows, labelled, found bias-recognition training ineffective and knowledge reorganisation modestly useful. This book is not a white-collar copy of that trial. It is the same moral: the case first. What would you need to know about *this* vendor, *this* SKU, *this* roster? If the claim is false, ask a better question, then name the fact. Correct the reason without crushing the talk.

When to go ahead: you can ask the two questions separately, name what you would need to know to have sides, and look or wait — at the desk and on the floor. When to stay: every answer is still a why-story, or “good” is still undefined, or the hour

is still a list of biases.

Tools, including AI

Optional. Keep this short. Point back to the work-hour AI box.

After you have tried, you may: explain this vendor page *to yourself* from a named source (title, URL, page number, contract clause, spreadsheet tab); ask for extra questions, then vet them; use a labelled SCRIPT after you tried; ask for a hint after an attempt; ask a tool to diagnose a reason you already gave.

You may not: park an unsupervised chatbot as your only partner; ask the tool to write your reason and paste it; invent an expert identity (“you are a procurement professor”) and treat that voice as a person; ask the tool to invent what the spec lists, what the roster says, or what the spreadsheet totals; put confidential workplace data into a public consumer tool; hand yourself an agent as a thinking partner for the first attempt.

A fluent explanation of a domain you do not know is still not knowledge you have. Open the named source. If the tool named a paper, open the paper. Language models invent sources; in one study more than half of one model’s citations and nearly a fifth of a later model’s citations were fabricated.⁶⁵ If you asked a tool for a paper, open the paper. Then return to *your* vendor, *your* tab, *your* SKU.

What “done enough” looks like

Before you move on, you should be able to hear a why-story and ask “How do you know that happened?” without treating the theory as the look. You should be able to ask, separately, “Why might it have happened?” You should be able to name what you would need to know to have sides on a vendor page, a spreadsheet tab, a roster, a SKU, or a process document — and then look, or wait, or ask a human colleague.

Checklist

⁶⁵William H. Walters and Esther Isabelle Wilder, “Fabrication and Errors in the Bibliographic Citations Generated by ChatGPT,” *Scientific Reports* 13 (2023). About 55% of GPT-3.5 citations and 18% of GPT-4 citations fabricated in that study. In-band for adults writing at work. If you asked a tool for a paper, open the paper. Point back to the work-hour AI box.

- You can say, in ordinary language: a maxim without something to think about will not implement.
- You can hear “because the part was wrong” as a why-story where a look was asked for.
- You can write “I do not yet know enough about X to have sides” without treating it as a defect.
- There has been at least one desk try (vendor page or spreadsheet tab) and at least one floor-or-shift try (roster / SKU / process document, or an object you can hold).
- The hour did not become a bias workshop.
- A spoken sentence plus one written line that names the claim, the because, and the check — and “good” is defined.
- Placement is by skill, not job title. If you still mash the two questions, or still treat a fluent page as sides you do not have, you are still in this chapter’s early band of the move, whatever your title.

A “critical thinking workshop certificate” is a receipt, not the work. Knowledge first is not a diploma. It is the next look.

You understand this. You know what to say tomorrow. You know what a why-story sounds like when it is wearing a look’s clothes. You know what an undefined “good” looks like on a dashboard. You can do this at the desk and on the floor.

Chapter 4 — Two sources

Critical Thinking for Adults

Michael Gannotti

2 September 2026



Figure 9: A still-life from a high three-quarter view: two vendor pages with greeked type side by side, a recap and notes with greeking, an invoice and a purchase order with greeking. No people. No logos. No readable UI.

Why this matters

What's a second place we could look?

That is not a request for the word *corroboration*. It is not a news-literacy course for adults. It is one question, asked of a page already on the table, about a claim that has so far had only one speaker. The vendor page said *leading*. The page is a speaker with a job: a sale. The spec sheet is another speaker with a job: a list. Leave this page. Look at one other thing.

A fluent sentence is not a check. Leave this page and look at one other thing. Recap versus notes. Spec versus review. Invoice versus purchase order. Two handoff notes. Product page versus data sheet. Two vendor pages in the same

category. Those are this week's objects. They are ordinary. They already decide what to believe and what to do.

The struggle is worth it because finished-looking pages keep people. Professional fact checkers leave an unfamiliar page. Ph.D. historians, adult experts in close reading, often stay. A doctorate and a lifetime of print sourcing did not produce leaving.⁶⁶ Expertise in one domain did not produce the web move. That is Willingham in a browser. Vertical reading — logo, layout, abstract, references on the page — is the adult expert's trap. You already know how to read closely. Close reading of *this* page is still one speaker.

High-school students who stay on a finished-looking page often never learn who is behind it. That is why a working adult leaves a vendor page, a slide, and a shift-rumor email *now*.⁶⁷ That sentence is the civic why. The items from that trial are not assigned. Being an adult is not a licence to reprint them. This book is not a civic-media course.

What this idea unlocks is every later sitting that used to stop at a fluent sentence. Without a second look, the recap *is* the check. Without naming who wrote it, a finished layout does the work of trust. Without leaving this page, two tabs from the same site feel like two sources. Two rumors are not two sources. Leaving the page is not disloyalty to the vendor you like.

If you only remember one thing from this chapter: leave this page. Ask what's a second place we could look. Then look.

This is not a reprint of *Critical Thinking for Life's* leave-the-page chapter, and not

⁶⁶Sam Wineburg and Sarah McGrew, "Lateral Reading and the Nature of Expertise: Reading Less and Learning More When Evaluating Digital Information," *Teachers College Record* 121, no. 11 (2019). Ten professional fact checkers, ten Ph.D. historians (ages 39–69, mean 47), 25 Stanford undergraduates. Fact checkers left unfamiliar pages (on the order of half a minute); historians took longer and often stayed; seven of ten historians treated references on the page as conferring legitimacy. Civic items in the paper are research objects — not assigned here. Adult experts are in-band as *why leave this page*. Expertise in close reading of documents did not produce leaving.

⁶⁷Joel Breakstone, Mark Smith, Sam Wineburg, Amie Rapaport, Jill Carrell, Marshall Garland, and Anna Saavedra, "Students' Civic Online Reasoning: A National Portrait," *Educational Researcher* 50, no. 8 (2021). 3,446 high-school students; 96% never learned a climate site's fossil-fuel ties. High school, not this book's population. One civic-why sentence in the body: high-school students who stay on a finished-looking page often never learn who is behind it — that is why a working adult leaves a vendor page, a slide, and a shift-rumor email now. Items not assigned, including the fossil-fuel item, the voter-fraud item, and any live civic URL. Being an adult is not a licence to reprint the trial.

a reprint of the 11–15 or 15–18 two-source weeks. One pointer. Then teach on the job’s pages.

Understand it yourself

Here is the everyday picture.

Two vendor pages, same category of tool. The first says industry-leading. You want to believe it, or you already do. The move that teaches is slower than the move that ends the talk. Ask: “What’s a second place we could look?” Then wait. Then open the second page, or the spec sheet, or last quarter’s incident log. The first page is still allowed to be interesting. It is no longer the only speaker. If the spec list is thinner than the adjective, you may change the sentence. That is the work. It is not losing.

A second everyday picture, still ordinary. The recap is in the inbox. “We decided X.” Open the notes from the same hour. Two speakers. Where they disagree is the fact that needs a third look — the recording if you have one, the calendar, the original email. You do not have to pick a team tonight.

Here is the precise picture.

Two-source is corroboration scaled to work. The adult civic practice is: leave an unfamiliar page and see what someone else says. This book’s practice is: leave this vendor page, this recap, this spec, this invoice, this handoff note, this product page, and look at one other thing that is not the same speaker repeating themselves. Two tabs from the same site are one speaker in two windows. A roundup article that already compared the two vendors is a third speaker with a job: a summary. Struggle before rescue: try the two pages before the roundup.

Ennis already names the adult version you can steal: use credible sources, and notice agreement among sources. Expertise, conflict of interest, reputation, and position to observe sit on the same list. Hearsay is weaker than the observer’s report.⁶⁸ You do not have to install that list as a syllabus. You do have to leave a fluent page for a second named place.

⁶⁸Robert H. Ennis, “The Nature of Critical Thinking” (2013, last revised September 2015). Credibility criteria a working adult can steal: who is in a position to know; conflict of interest; reputation; agreement among sources; hearsay weaker than an observer’s report. Not a workplace syllabus.

Three questions, translated onto work objects as a design hypothesis, labelled, not a trial: who is behind this page; what would show the claim; what does a second look say.⁶⁹ Who wrote “industry-leading”? What would show it — the spec sheet, the notes, the purchase order, the roster? What does the other page say? Near transfer to another vendor page is the honest hope. Far transfer to civic web evaluation is an open question. Willingham predicts it will not happen without the later domain’s knowledge.⁷⁰ Practice on the known stake. Do not rehearse the later domain as content.

Wineburg and McGrew sat professional fact checkers, Stanford undergraduates, and Ph.D. historians in front of unfamiliar web pages. The fact checkers left. The historians often stayed. Historians in that study were adults, mean age in the mid-forties, trained for a lifetime to read documents closely. Seven of ten treated references *on the page* as conferring legitimacy.⁷¹ That is the expert trap, in-band. You can be good at your job’s documents and still stay. Close reading is a virtue. It is not leaving.

History-class studies that taught sourcing, corroboration, and a named writing move moved those moves *in history class*.⁷² Copy the shape: more than one document; a named writing move. Copy not the period, not the professional develop-

⁶⁹The three Civic Online Reasoning questions (who’s behind this; what’s the evidence; what do other sources say) translated onto work objects. Design hypothesis, labelled, not a trial. Sam Wineburg, Joel Breakstone, Sarah McGrew, Mark Smith, and Teresa Ortega, “Lateral Reading on the Open Internet” (2022 government-course study): six ~50-minute lessons in a high-school government course; modest incomplete gain. Journalism summaries that “doubled” scores are journalism. This book does not reprint those six lessons.

⁷⁰Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator* 31, no. 2 (Summer 2007), and “How to Teach Critical Thinking,” NSW Department of Education occasional paper (2019). Near is not far. Far transfer from vendor-page leaving to a cloaked live site is an open question; Willingham predicts it will not happen without knowledge of the later domain. Cue similarity; Gick and Holyoak, Detterman: noticing is the bottleneck.

⁷¹Wineburg and McGrew (2019), as in note 66. Historians stayed. That is the adult expert trap, said honestly in the body, not as a gotcha.

⁷²Jeffrey D. Nokes, Janice A. Dole, and Douglas J. Hacker, “Teaching High School Students to Use Heuristics While Reading Historical Texts,” *Journal of Educational Psychology* 99, no. 3 (2007): ages 16–17; multiple texts beat textbook on sourcing and corroboration; contextualization did not move. Avishag Reisman, “Reading Like a Historian,” *Cognition and Instruction* 30, no. 1 (2012): QED, 236 eleventh-graders; history class with professional development. Susan De La Paz and Mark Felton, “Reading and Writing from Multiple Source Documents in History,” *Contemporary Educational Psychology* 35, no. 3 (2010): grade 11. History-class evidence, labelled. Copy shape (more than one document; a named writing move), not the period, not the packet, not a civic site.

ment, not a civic site. A grade-eleven packet is not this week's gym.

Two sources can disagree without anyone having to pick a team tonight. Agreement among sources is useful when it is there. Its absence is also information. Two quotes for the same SKU can both be "offers." One of them can still be better supported by what is actually included. A product page can say *best*. The data sheet lists capacity. "Best" is often the page's word.

Wrong answers you should be able to hear

These are illustrations, not portraits.

1. **A fluent recap offered as a check.** "We decided X because the recap said we decided X." One speaker. What it usually means: testimony is doing the work of a look. What to say: "What's a second place we could look?" Then the notes.
2. **Two tabs from the same site, or two rumors, treated as two sources.** What it usually means: independent was skipped. What to say: "Is that a second speaker, or the same speaker twice?" Name both. If they are the same, look again.
3. **Staying for the logo, the layout, or the references on the page.** What it usually means: the adult expert trap. Vertical reading felt like work. What to say: "Who is behind this page? What would we have to leave this page to find out?" Then leave.
4. **Leaving treated as disloyalty.** "We already like this vendor." What it usually means: own-side is doing the work of a second look. What to say: leaving is not awarding the other vendor. It is asking the spec sheet the same question.
5. **A civic-web item assigned "because adults can handle it."** What it usually means: the later domain arrived as content. What to say: practice leaving a vendor page, a recap, an invoice. The later, heavier claim will need knowledge this book is not teaching. That is Willingham, not a dodge.

A sixth you will also hear: a roundup article offered as the first look. The roundup already compared them. Struggle before rescue: two pages first, then the roundup if you still need it.

Five-minute warm-up

Do this before the try, with two pages you already have — two vendor one-pagers, a recap and notes, an invoice and a purchase order — and a scrap of paper.

Minute 1. Write the claim. Write speaker one. Write speaker two. If you cannot name a second speaker that is actually independent, this is not yet a two-source sitting. Pick a different claim.

Minute 2. Translate three questions onto your object, in ordinary words: Who wrote this? What would show the claim? What does the other page say?

Minute 3. True or false, said aloud: “The recap is the check.” “Two tabs from the same site are two sources.” “If they disagree, we have to pick a team.” You are practicing the hearing, not racing.

Minute 4. Imagine a colleague saying “because the page said leading.” Decide what a rant would do. Decide what “what’s a second place we could look?” would do instead.

Minute 5. Write the sentence you will actually say: “What’s a second place we could look?” Put the pencil down. That sentence is the lesson.

If those five minutes are fluent for you, you can hear a fluent page with no second speaker, two tabs that are one speaker, and a disagreement treated as a fight. That is the job.

How to try it this week

The work hour already licensed session shape. This chapter fills it with a second speaker.

Session shape

Warm-up (2 minutes). Yesterday’s claim. “How did we know?” No second source yet if last time was only a how. No device required.

Short model (2–3 minutes). You leave a recap for the notes, or a vendor page for a spec, in two lines. You name both speakers. Then you stop. You are showing the leave, not delivering a lecture titled *Corroboration*.

Your attempt (8–15 minutes). You name the claim, name two places, look. Wait. Struggle before rescue: ask, wait, hint (“the spec is here”), then a short model on *your* scrap.

One good question. The exact wording you can say this week:

“What’s a second place we could look?”

If the object is not a vendor page, keep the shape: “What’s a second place for this recap?” “What’s a second place for this invoice?” “What’s a second place for this handoff?”

Mixed practice (5 minutes). A cousin of the same genre: if the first was two vendor pages, the second can be invoice versus purchase order, or two handoff notes.

A written line (3–5 minutes). The comparison, dated, sources named, so a colleague who missed the meeting could follow it. If “it” and “they” still need the meeting, rewrite.

Exact wording you can say in a meeting or at a handoff

- “What’s a second place we could look?”
- “What do you mean, *leading*?”
- “How do you know?”
- “Who’s behind this page?”
- “What would count as a second look that is not the same thread?”
- “Can you write the comparison so a colleague who missed the meeting could follow it?”

Ask one. Wait. A slow three. Stahl, labelled. After they stop, wait again. Look at the second page, the notes, the purchase order, or the data sheet, not at the person, if the silence is hard. At work the shrug often looks like refusal. Leave it. After “write the sentence,” that is task-completion time. Wait in private when the power difference is steep.

First try-it

Two vendor pages, desk, 10–15 minutes. Same category of tool or service; two vendors’ one-pagers. Opening: “What’s a second place we could look?” Leave the first page. Who is behind each. What is the claim. What does a second look say.

How to fade help

First time: you leave together, you name both speakers, they echo. Second: they leave, you wait, you hint (“the spec is here”). Third: they leave and say the sentence. Fourth: a cousin of the same claim — invoice versus purchase order, two

handoff notes — with the first page closed. The word *corroboration*, if it arrives at all, arrives after the look.

If they stay on the logo, name the trap once: finished layout is not a second speaker. If they treat two rumors as two sources, ask whether the speakers are independent. If they treat leaving as disloyalty, name it: the spec sheet is not a rival vendor. You may name a better reason if the claim is false. Hearing a wrong reason, then asking a better question, then naming a fact if needed, is the work.

If a tool is in the room, it waits. Try the two pages before the tool's summary. A labelled SCRIPT after you tried is a prompt for *you* to say, not a comparison to paste. Point back to the work-hour AI box.

When to stop talking

When two speakers have been named, a look has happened, and the sentence has been kept, changed, or marked as still needing a third look. Ten to twenty-five minutes on two pages is the right ambition. If two-source becomes a rant, return to two pages. If it becomes citizen-journalism, return to ordinary work stakes. If the table is becoming a fight, stop. Postpone. Pick a smaller claim.

Desk

Two vendor pages. Recap versus notes. Spec versus review — the spec the team wrote versus the reviewer's comments, or the spec versus what shipped. Invoice versus purchase order can live here too when the files are already on the desk. Safety: curious, not a courtroom. No public quiz of a junior in front of a vice-president. Cue similarity next time the genre recurs. Last month's two vendor pages do not automatically produce this month's leave. Near is not far. You will have to say, out loud, "this is like the two pages."

Knowledge first still holds. Leaving a vendor page for a spec sheet still requires knowing what the spec is a spec *of*. Two fluent adjectives are not two sides. If you cannot yet name what the product does in *your* workflow, the second page has nowhere to land. Chapter 3 sits underneath this one. Name what you would need to know. Then leave.

Floor-or-shift

Invoice versus purchase order when the paper is in the bay. Two handoff notes, after, not during, a live procedure. Product page versus data sheet for a thing already in the bay. Listed capacity, weight, warranty words on the sheet — not

stars. Safety: no public quiz, no undercover audit, no gotcha against a colleague, a vendor, or a shop, no lecture to a stranger's customer. A clinic try-it uses a protocol already in use — not a diagnosis of a named patient in this book.

The floor is not a different faculty. It is the same leave in a different jacket. A fast adjective at a vendor call is still one speaker. If the bay is live, compare after, not during. Two quotes for the same SKU are two speakers if they are independent. They are one speaker twice if they reprint each other's sentence.

Practice that actually builds it

Desk, floor, clinic, shop *are* the practice. They do not replace the named move. Each try-it below names time, materials, safety, and the move. At least one lives at the desk. At least one lives on the floor or shift.




Figure 10: Two vendor pages with greeked type side by side, no logos, no readable UI, no faces.

Try-it: Two vendor pages (desk)

Time. 10–15 minutes.

Materials. Same category of tool or service; two vendors' one-pagers you actually use or were already sent. A dated notebook.

Safety. Ordinary reading. You are not running a sting on a shop. Skip live-controversy pages.

The move. Leave the first page. Who is behind each. What is the claim. What does a second look say. Struggle before rescue: the two pages before the roundup article that already compared them.

Illustration, not a reported workplace: page A says industry-leading uptime. Page B, same category, lists a maintenance window in small type. You wanted A. You leave. The spec sheet for A names the same window. The sentence changes from *leading* to *the window is named, and it matches last quarter's tickets* — or it doesn't. Either is thinking if the leave happened.

Try-it: Recap versus notes (desk)

Time. 5–10 minutes.

Materials. The meeting recap and the notes from the same hour.

Safety. Ordinary. You are not putting a colleague on trial for their adjectives.

The move. Two sources on the same hour. The recap is a speaker with a job: a story of the hour. The notes are a speaker with a job: what was written down. Where they disagree is the fact that needs a third look.

This is blocked practice: both pages are already in your hand. Mixed practice is later, when you have to notice that *this* recap is like last Tuesday's recap.

Try-it: Spec versus review (desk)

Time. 8–15 minutes.

Materials. The spec the team wrote versus the reviewer's comments, or the spec versus what shipped.

Safety. Ordinary. Correction of the spec is allowed. “Never correct” is not this book’s rule.

The move. The spec is a theory of what the thing should be. The review is evidence of what someone found. Merging them is how a failed review becomes a story about intent. Leave the spec for the review. Name both. Write the disagreement so a colleague who missed the meeting could follow it.

Try-it: Invoice versus purchase order (floor-or-shift, or desk)

Time. 5–10 minutes.

Materials. The invoice and the purchase order for a thing you already ordered or received.

Safety. Ordinary looking. Not a gotcha against a vendor on the floor. Confidential numbers stay off public tools.

The move. Line items, quantities, the purchase-order number, the named requester — not the total in bold. The bold total is a fluent sentence in numerals. The lines are the look.

Try-it: Two handoff notes (floor-or-shift)

Time. 5–10 minutes after, not during, a live procedure.

Materials. Outgoing and incoming notes on the same bay, line, or patient. A protocol already in use if this is a clinic.

Safety. After, not during. No public quiz. No diagnosis of a named patient in this book. Noticing is not a sting.

The move. Where they disagree is the fact that needs a third look. Two speakers. Not one speaker twice. Write the disagreement so the next shift could follow it without having been there.

Try-it: Product page versus data sheet (floor-or-shift)

Time. 5–10 minutes.

Materials. The product page and the data sheet for a thing already in the bay.

Safety. Ordinary. No lecture to a stranger's customer. No tasting unknown things. No shoplifting a "test."

The move. Listed capacity, weight, warranty words on the sheet — not stars. The page is a speaker with a job: a sale. The sheet is a speaker with a job: a list. Leave the page.

Talk / meeting box

Opening question. "What's a second place we could look?"

Not: "Did we have a good meeting?" Not: "Is this vendor bad?"

Follow-ups (pick three to five, not all at once).

1. What do you mean, *leading*?
2. How do you know?
3. Who's behind this page?
4. What would count as a second look that is not the same thread?
5. Can you write the comparison so a colleague who missed the meeting could follow it?

How to wait. After you ask, a slow three. Stahl, labelled classroom-origin. After they stop, wait again. If they are mid-reason, let them finish. Look at the second page, the notes, the purchase order, or the data sheet, not at the person, if the silence is hard. After "write the sentence," that is task-completion time. Wait in private when the power difference is steep. Rowe is science class, not a stand-up RCT.

What a stuck silence usually means. They stayed on the logo. They treated references on the page as a check. They think two rumors are two sources. They think leaving the page is disloyalty to the vendor they like. Wait-time 1 was zero. They offered a why-story where a look was asked for. They think a finished-looking paragraph is already a check.

Your job is to pose, wait, and point back at the second object — not to fill the silence with the right vendor, and not to turn the meeting into a fight.

One incorrect example to diagnose. A fluent recap with no notes, offered as a check. "We decided X. I already wrote the recap. The recap is the record." What it usually means: one speaker finished the job in a finished-looking paragraph. What to say: "What's a second place we could look?" Then the notes. If the notes go the other way, correct the sentence. That is the work. It is not losing.

If it isn't clicking

Three diagnostics. No shame. The next move is small.

1. They stay. You stay. Logo, layout, references on the page, a copyright line, a finished deck. Vertical reading felt like work. Slow down. Name the trap once, in ordinary language: this page can look finished and still be one speaker. Historians stayed. Adult experts stayed. You are not deficient for staying. The next move is the leave. Set a small clock if it helps — not as a trial, as a nudge — and open the second named place. If they stare, model on recap versus notes, which you already hold. Then return to the vendor page.

2. Two rumors, two tabs, or a roundup treated as two sources. Independent was skipped. Slow down. Name both speakers. Ask whether they are the same speaker twice. Struggle before rescue: two pages first, then the roundup. If the second “source” is a forwarded email that quotes the first, it is still one speaker. The next move is a named place that does not live in the same thread: the spec, the notes, the purchase order, the roster, the data sheet.

3. Leaving feels like disloyalty, or two-source has become a rant. Own-side is doing the work, or disagreement was treated as a fight. Slow down on meaning. “What do you mean, *leading*?” Then the spec. Leaving is not awarding the other vendor. Two sources can disagree without a team tonight. If two-source becomes citizen-journalism, return to ordinary work stakes. If the claim is false, ask a better question, then name the fact. Correct the reason without crushing the talk. Ask a human colleague who holds the second file. A tool can diagnose a reason you already gave; you hear the diagnosis and decide. A tool does not get the first leave.

When to go ahead: you can leave a fluent page, name two speakers that are not the same speaker twice, look, and change the sentence if the look goes the other way — at the desk and on the floor. When to stay: the first page is still the check, or two rumors are still doing the work of two sources.

Tools, including AI

Optional. Keep this short. Point back to the work-hour AI box.

After you have tried the two pages, you may: explain this page *to yourself* from a named source (title, URL, page number, contract clause); ask for extra questions,

then vet them — throw out any that import a live controversy, a secret, a pretend expert, a colleague’s private performance, or a civic-trial item; use a labelled SCRIPT after you tried; ask for a hint after an attempt; ask a tool to diagnose a reason you already gave.

You may not: park an unsupervised chatbot as your only partner; ask the tool to write the comparison and paste it as yours; invent an expert identity; treat two model-generated summaries as two sources; ask the tool to invent what the spec lists; put confidential invoices or purchase orders into a public consumer tool; hand yourself an agent as a thinking partner for the first leave.

A fluent roundup from a machine is a third speaker with a job: a summary. It is not the two pages. Try the two pages first. Open *your* employer’s acceptable-use page. Then return to the spec.

What “done enough” looks like

Before you move on, you should be able to leave a fluent page and ask “What’s a second place we could look?” without treating the first speaker as the check. You should be able to take two vendor pages, a recap and notes, a spec and a review, an invoice and a purchase order, two handoff notes, or a product page and a data sheet, name two speakers, look, and write the comparison so a colleague who missed the meeting could follow it.

Checklist

- You can hear a fluent recap with no notes as one speaker.
- You can hear two tabs from the same site, or two rumors, as not yet two sources.
- You can name the adult expert trap in ordinary language: staying on a finished-looking page, including its references, is still one speaker.
- There has been at least one desk try (two vendor pages, recap versus notes, or spec versus review) and at least one floor-or-shift try (invoice versus purchase order, two handoff notes, or product page versus data sheet).
- Struggle before rescue: two pages before the roundup.
- A written comparison, dated, sources named, a colleague who missed the meeting could follow it.
- No civic-trial items were assigned. No live controversy entered as content.
- Placement is by skill, not job title. If you still stay, or still treat two rumors as

two sources, you are still in this chapter's early band of the move, whatever your title.

A "critical thinking workshop certificate" is a receipt, not the work. Leaving is the work. The title of the year's work, if a title is needed at all, is still the job you hold — not Critical Thinking I, and not a news-literacy badge.

You understand this. You know what to say tomorrow. You know what a finished-looking page without a second speaker looks like. You know what two rumors wearing two-source's clothes sound like. You can do this at the desk and on the floor.

Chapter 5

Graphs



Figure 11: A generic deck chart on a table, two bars with the axis starting above zero, a pencil, and the same numbers redrawn from zero on scrap paper; no product chrome, no logos, no people

Why this matters

What does this chart actually show?

That question unlocks a picture made of numbers. At work a graph is a claim made with ink. Someone chose the title, the units, where the axis starts, who is in the count, and what to leave off. Robert Ennis’s public list of thinking abilities includes elementary graphs and maths, which is why a deck chart belongs in this book without turning the book into a statistics course.⁷³ You do not need a campaign dashboard. You do not need a culture-war infographic. You need one picture already in the deck, one look at the axis, and one honest sentence about what the picture does and does not say.

A graph can be honest. A graph can also make a small change look like a leap, simply by where the picture starts. That is not a children’s worksheet. In one useful study of y-axis truncation — chopping a bar graph’s axis so it does not start at zero — about five in six participants across five studies saw the illustrated difference as larger than it was. Teaching about truncation helped a little and did not wipe the effect out. Graph literacy did not predict who was fooled. The sample was adult, including people with advanced degrees in quantitative fields. This is a useful study, not a promise that every workplace will see the same result.⁷⁴ The

⁷³Robert H. Ennis, “The Nature of Critical Thinking: Outlines of General Critical Thinking Dispositions and Abilities,” 2015 PDF: Ability #4, a recent addition, “Understand and use elementary graphs and maths.” This chapter is a thinking manual that uses elementary graphs; it is not a mathematics textbook. Susan N. Friel, Frances R. Curcio, and George W. Bright, “Making Sense of Graphs: Critical Factors Influencing Comprehension and Instructional Implications,” *Journal for Research in Mathematics Education* 32, no. 2 (2001): 124–158. Graph sense; Curcio’s three levels — read the data, read between the data, read beyond the data — as reported in this review. Frances R. Curcio, “Comprehension of Mathematical Relationships Expressed in Graphs,” *Journal for Research in Mathematics Education* 18, no. 5 (1987): 382–393, is the source of the three-level terminology; three-level language here is via Friel, Curcio, and Bright 2001. Access date for sources in these notes: 2 September 2026.

⁷⁴Brenda W. Yang, Camila Vargas Restrepo, Matthew L. Stanley, and Elizabeth J. Marsh, “Truncating Bar Graphs Persistently Misleads Viewers,” *Journal of Applied Research in Memory and Cognition* 10, no. 2 (2021): 298–311. Five studies; y-axis truncation led viewers to perceive illustrated differences as larger; 83.5 percent of participants across the five studies showed a truncation effect; Cohen’s d 0.48–1.26; graph literacy failed to predict the size of the effect; warnings reduced but did not eliminate it. Adult samples (MTurk, mean ages about 32–35; Study 5, Ph.D. students in quantitative fields and in the humanities). Even quantitative Ph.D. students with a warning immediately before the ratings still showed a truncation effect (Study 5, estimate 0.24). The authors’ recommendation is specific to bar graphs, where length encodes quantity; they do not generalise a “never truncate” rule to line graphs and dots. Useful study, not a promise that every workplace

move is still worth teaching: *where does this picture start?*

A journalist's illustrated book from 1954 named the same trick the Gee-Whiz Graph: chop the bottom, then change the scale so a modest rise climbs the page. The same book said a truncated bar chart deserves exactly the same reputation as a truncated line. That book is a labelled neighbour for lying-with-charts. It is not a trial.⁷⁵ The tricks did not expire. They live on the slide you will sit through on Tuesday.

That is worth the struggle because the job already meets graphs. A two-bar sales pair in a deck. A dashboard of tiles, a bar pair, and a sparkline on a shared screen. A board number on the floor. A count versus a rate. A relative percent with no base. Each of those can be a fluent caption — “jumped,” “soared,” “collapsed,” “dramatic” — offered as if the picture had already done the looking. The adult who can hear a caption offered as a check can practice this week. The adult who can lift a number, compare two bars, say what the picture does not tell them, and write one sentence a colleague who missed the meeting could follow has practiced deciding what to believe on a domain they actually know.

A fluent sentence is not a check. Leave this picture and look at the axis, the units, and the numbers. Then write what you actually saw.

This week you can learn three ordinary-language levels of graph reading, and how to ask “where does this picture start?” of a deck chart, then ask for one written sentence. Today, at the desk, you can cover the title, read the axes, and try the same numbers from zero. Today, on the floor, you can look at a board number or a dashboard tile and ask what this is a number *of*.

Understand it yourself

You do not need to be a statistician. You do need to hear “look at that leap” as a caption, not yet a look at the axis, and “the deck said we jumped” as a sentence

will see the same result.

⁷⁵Darrell Huff, *How to Lie with Statistics* (New York: W. W. Norton, 1954), pictures by Irving Geis. Journalist's illustrated popular book, not a randomised trial. Chapter 5, “The Gee-Whiz Graph”: chop the bottom of a line; change the ordinate so a modest rise climbs the page. Chapter 6, “The One-Dimensional Picture”: “A truncated bar chart has, and deserves, exactly the same reputation as the truncated line graph.” Chapter 2, “The Well-Chosen Average,” and Chapter 3, “The Little Figures That Are Not There,” are neighbours of the same family. Labelled neighbour for lying-with-charts, not a trial of this week's sitting.

that still has a spreadsheet.

Everyday picture. Two months of unit sales. April, 48,000. May, 56,000. Someone draws two bars. The axis starts at 40,000. April's bar is a short stub. May's bar looks twice as tall. "Jumped," someone says in the meeting. Now redraw the same numbers from zero. The leap shrinks. The numbers did not change. The claim "jumped" was the picture's.

That pair is an illustration, not a reported workplace. Your numbers will be different. The move will not.

Precise picture. A graph is a *claim made with ink*. Three jobs, in ordinary language, named by researchers who study how people read graphs:⁷⁶

- **Read the data.** Lift a number. What does this bar say? April, 48,000 units. Title, axes, units first. Who is in the count?
- **Read between the data.** Compare. Which month was higher, and by how much? People often manage the first job. Errors cluster here: subtracting the wrong way, interpolating a point that is not there, treating a chopped axis as if it were the size of the change.
- **Read beyond the data.** Predict, or notice what the graph does not say. Will June be higher? The picture does not tell you. Beyond needs prior knowledge related to the graph — what a unit sale *is*, who was in the count, that two months are not a year, that a relative percent still has a base. A maxim without that knowledge will not implement.⁷⁷

⁷⁶Robert H. Ennis, "The Nature of Critical Thinking: Outlines of General Critical Thinking Dispositions and Abilities," 2015 PDF: Ability #4, a recent addition, "Understand and use elementary graphs and maths." This chapter is a thinking manual that uses elementary graphs; it is not a mathematics textbook. Susan N. Friel, Frances R. Curcio, and George W. Bright, "Making Sense of Graphs: Critical Factors Influencing Comprehension and Instructional Implications," *Journal for Research in Mathematics Education* 32, no. 2 (2001): 124–158. Graph sense; Curcio's three levels — read the data, read between the data, read beyond the data — as reported in this review. Frances R. Curcio, "Comprehension of Mathematical Relationships Expressed in Graphs," *Journal for Research in Mathematics Education* 18, no. 5 (1987): 382–393, is the source of the three-level terminology; three-level language here is via Friel, Curcio, and Bright 2001. Access date for sources in these notes: 2 September 2026.

⁷⁷Daniel T. Willingham, "Critical Thinking: Why Is It So Hard to Teach?," *American Educator*, Summer 2007; and Willingham, "How to Teach Critical Thinking," NSW Department of Education, 2019. Thinking is not a bicycle skill you apply in any situation; processes are intertwined with content; maxims without knowledge will not implement. Friel, Curcio, and Bright 2001: "read beyond the data" requires prior knowledge related to the graph. Near practice on a deck chart does not, by itself, install unaided judgment on an unfamiliar later chart. Cue the similarity ("this is like

Those three jobs are school-age literature in origin. The three questions scale to a deck and a dashboard. They are not a personality of distrust. Some graphs start at zero and show what they claim to show. The move is the axis, the units, and the missing piece — not “all graphs lie.” You are teaching yourself how to *read* a picture, not how to manufacture one that misleads.

A table is a cousin of a graph. A spreadsheet tab is a table. A recap’s word *jumped* is a caption. The tab is a second look. You already practiced leaving a recap for the notes in the two-source chapter. This chapter asks what the *picture* actually shows, then what it does not.

A dashboard is a cousin with more furniture. Tiles across the top. A bar pair in the middle. A sparkline in the corner. No product chrome in this book’s pictures: the practice is the question, not the screenshot’s pixels. Each tile is still a claim. Ask what it is a number *of* before you trust the colour.

The truncation finding that is worth keeping is about **bar graphs**. Length encodes quantity. Starting the axis above zero stretches the difference. Line graphs and dots are a different encoding. A sales *line* cropped at 80,000 needs a different sentence than a sales *bar* cropped at 80,000: the line can still show the shape of change; the bar’s height is being asked to stand for the amount. How adults read those two encodings on decks is not a settled trial in the sources this book opened. Stay with the bars you have. When you meet a line, still ask where the axis starts, still lift the number, and do not pretend the bar finding already covered the line.⁷⁸

79

the deck chart”).

⁷⁸Brenda W. Yang, Camila Vargas Restrepo, Matthew L. Stanley, and Elizabeth J. Marsh, “Truncating Bar Graphs Persistently Misleads Viewers,” *Journal of Applied Research in Memory and Cognition* 10, no. 2 (2021): 298–311. Five studies; y-axis truncation led viewers to perceive illustrated differences as larger; 83.5 percent of participants across the five studies showed a truncation effect; Cohen’s *d* 0.48–1.26; graph literacy failed to predict the size of the effect; warnings reduced but did not eliminate it. Adult samples (MTurk, mean ages about 32–35; Study 5, Ph.D. students in quantitative fields and in the humanities). Even quantitative Ph.D. students with a warning immediately before the ratings still showed a truncation effect (Study 5, estimate 0.24). The authors’ recommendation is specific to bar graphs, where length encodes quantity; they do not generalise a “never truncate” rule to line graphs and dots. Useful study, not a promise that every workplace will see the same result.

⁷⁹Darrell Huff, *How to Lie with Statistics* (New York: W. W. Norton, 1954), pictures by Irving Geis. Journalist’s illustrated popular book, not a randomised trial. Chapter 5, “The Gee-Whiz Graph”: chop the bottom of a line; change the ordinate so a modest rise climbs the page. Chapter 6, “The One-Dimensional Picture”: “A truncated bar chart has, and deserves, exactly the same

A well-chosen average is a neighbour of the same family. “Average deal size” with no mean, no median, and no range is a caption wearing a number. Little figures that are not there — a steep line with no scale, “40% fewer” with no base — are the same kind of neighbour.⁸⁰ When a tile shows a relative percent, asking “of how many, from what, compared with what” is ordinary risk communication, not a theory of thinking.⁸¹ Then look.

Wrong answers you should be able to hear

Including from yourself.

1. The caption instead of the number. “Jumped.” “Soared.” “Collapsed.” “Dramatic.”

What it usually means. You read the picture’s drama, not the data. The axis was skipped. The adjective in the title did the looking.

What to say. “What does this chart actually show? Where does the axis start? Let’s lift the number first.”

2. A comparison with no amount. “May was way up.”

What it usually means. Between-the-data work started and stopped at a feeling. The subtraction was not done.

What to say. “Which month was higher, and by how much?” Wait. If you need it: “48,000 and 56,000. That’s 8,000. Is 8,000 a leap, or does the picture make it look

reputation as the truncated line graph.” Chapter 2, “The Well-Chosen Average,” and Chapter 3, “The Little Figures That Are Not There,” are neighbours of the same family. Labelled neighbour for lying-with-charts, not a trial of this week’s sitting.

⁸⁰Darrell Huff, *How to Lie with Statistics* (New York: W. W. Norton, 1954), pictures by Irving Geis. Journalist’s illustrated popular book, not a randomised trial. Chapter 5, “The Gee-Whiz Graph”: chop the bottom of a line; change the ordinate so a modest rise climbs the page. Chapter 6, “The One-Dimensional Picture”: “A truncated bar chart has, and deserves, exactly the same reputation as the truncated line graph.” Chapter 2, “The Well-Chosen Average,” and Chapter 3, “The Little Figures That Are Not There,” are neighbours of the same family. Labelled neighbour for lying-with-charts, not a trial of this week’s sitting.

⁸¹Gerd Gigerenzer and colleagues, natural-frequency work, including Gigerenzer and Ulrich Hoffrage, “How to Improve Bayesian Reasoning Without Instruction: Frequency Formats,” *Psychological Review* 102, no. 4 (1995): 684–704. Labelled risk communication when a dashboard shows a relative percent: ask of how many, from what, compared with what. Not a theory of critical thinking, and not this chapter’s course.

like one?”

3. Beyond the data with no knowledge, or a moral in place of a look. “This proves the year is a disaster.” Or: “Whoever drew this is trying to trick us.”

What it usually means. You jumped to a story the picture cannot carry, or you treated a chopped axis as a character flaw. The deck chart in front of you is two months. Two months are not a year, and a chopped axis is a design choice, not a personality.

What to say. “What does this picture *not* tell you? What would we still need to know?” Then, if needed: “We have two months. We do not have a year. The axis started at 40,000. That is the look we have.”

4. The recap as the chart, or the fluent paragraph as the axis. “The write-up said we jumped, so the numbers must show that.” “My paragraph already said dramatic.”

What it usually means. A fluent sentence from a speaker who already picked a word. The spreadsheet, or the axis, is still sitting there.

What to say. “What’s a second place we could look? Let’s lift one number. What does *jumped* mean if we use that number? Can you write that so a colleague who missed the meeting could follow it?”

5. Graph literacy as a vaccine, or a warning caption as the whole method. “I took a data course, so I won’t be fooled.” Or: the slide already says “axis truncated,” so the looking is done.

What it usually means. A useful study was treated as a personality. Graph literacy did not predict the size of the truncation effect. Warnings reduced the effect and did not eliminate it. Even people with years of quantitative training still showed it.⁸²

⁸²Brenda W. Yang, Camila Vargas Restrepo, Matthew L. Stanley, and Elizabeth J. Marsh, “Truncating Bar Graphs Persistently Misleads Viewers,” *Journal of Applied Research in Memory and Cognition* 10, no. 2 (2021): 298–311. Five studies; y-axis truncation led viewers to perceive illustrated differences as larger; 83.5 percent of participants across the five studies showed a truncation effect; Cohen’s *d* 0.48–1.26; graph literacy failed to predict the size of the effect; warnings reduced but did not eliminate it. Adult samples (MTurk, mean ages about 32–35; Study 5, Ph.D. students in quantitative fields and in the humanities). Even quantitative Ph.D. students with a warning immediately before the ratings still showed a truncation effect (Study 5, estimate 0.24). The authors’ recommendation is specific to bar graphs, where length encodes quantity; they do not generalise a “never truncate” rule to line graphs and dots. Useful study, not a promise that every workplace

What to say. “Let’s still lift the number. Where does this picture start? What would the same numbers look like from zero?”

Five-minute warm-up

Before the sitting, with no one else in the room:

1. Draw two bars: 48,000 and 56,000, axis starting at 40,000. Then redraw from zero. Say out loud, to the empty room: “What does this chart actually show?” Wait a slow three. Three seconds is a convention from classroom think-time, not a workplace study.⁸³
2. Write the three levels in ordinary language: lift a number; compare; say what it does not tell us. You will not lecture those words as vocabulary. You will have them.
3. Write one caption you expect — “jumped,” “soared,” “dramatic” — and the better question: “Where does this picture start?” Then, if needed, the redraw from zero. Write the sentence you will ask for at the end: *I claim X because Y. We checked by Z.*

If those three are fluent for you, you can hear a caption offered as a check. That is the job. You still lift the number.

How to try it this week

The sitting uses the work-hour shape already in the front of this book: warm-up, short model, your attempt, one good question, mixed practice, a written line. Fill it with one deck chart, one redraw from zero, or one dashboard already on a shared screen. Fifteen to twenty-five minutes on a named artifact. Then stop.

Warm-up (2–4 minutes). Yesterday’s two-source look, from memory. “What was the second place we looked?” Then: “Today the picture is the claim. We will ask what it actually shows.” No new chart yet.

will see the same result.

⁸³Robert J. Stahl, “Using Think-Time and Wait-Time Skillfully in the Classroom,” ERIC Digest ED370885 (1994). Think-time; about three seconds as a convention, not a sacrament; written for classroom teachers. A working adult can use the move. This is not a workplace wait-time trial; none was opened for this book.

Short model (5–8 minutes). Put one slide on the table. One comparison. One axis.




Figure 12: Two generic bars, April 48,000 and May 56,000, y-axis starting at 40,000; the same numbers redrawn from zero on a second scrap. No product chrome, no logos

You name the claim the picture seems to make: “jumped.” You lift April: 48,000. You lift May: 56,000. You wait. You ask where the axis starts. You redraw from zero. You say what you saw: the numbers did not change; the leap was the picture’s. Then you show the *wrong* cousin on purpose: treating “jumped” as the look, without the axis. “Can we tell for sure from the drama of the bars?” Wait. Then stop talking sooner than feels polite.

This pair is an illustration, not a reported workplace. Use *your* slide if you have one.

Exact wording you can steal.

On the picture: “What does this chart actually show?”

On the start: “Where does this picture start? Where does the axis begin?”

On the lift: “What number is this bar? What are the units?”

On the compare: “Which month was higher, and by how much?”

On the missing piece: “What does this picture *not* tell you? Who is included?”

On a caption: “I hear *jumped*. Let’s look at the numbers.”

On a relative percent: “Forty percent of what? From what? Compared with what?”

On a written sentence: “Can you write that so a colleague who missed the meeting could follow it? Claim, because, how we checked.”

On mind-change: “You thought it was a leap. We redrew from zero. You can say the new thing. That is not losing.”

Your attempt. After the model, you run the next pair — last month’s actual slide if you have it, or a simple table of two weeks — unaided. Product: one number lifted, one comparison, one thing the picture does not say, then one dated sentence a colleague who missed the meeting could follow. You wait, including with yourself. You do not “help” the subtraction by grabbing the answer. You do not finish the sentence.

One good question. After the attempt: “If we had started the axis at 40,000, could ‘jumped’ still be the numbers?” You generate the why.

Mixed practice. One deck redraw. One dashboard tile. One board number on the floor. Mix, so the move is not glued to sales bars.

Written line (2–5 minutes, unaided). “What does this chart actually show? Where does the axis start?” Spoken is enough for a first try. One dated line is the week’s win: title, one number, one comparison, one missing piece, how we checked. Tool closed.

How to fade help. You name all three jobs (lift, compare, missing) → you lift while a sticky note names the rest → you name all three → you catch a chopped axis without the prompt → you write the sentence without dictating it to yourself from a template. When you can say “this picture starts at 40,000” without a hint, put the sticky-note scaffold away. Bring it back when a new kind of picture arrives — a line, a sparkline, a floor board, a relative-percent tile.

When to stop talking. After one model and one question. You hold the key: is there a title, an axis, a number, a wait? You do not become a lecture about statistics. The word *truncation* can wait. The axis cannot. In a meeting, one good question beats five. After you ask, wait. If the power difference is steep, wait in private. You are not citing a science-class wait-time paper as a workplace trial.^{84 85}

Desk / floor-or-shift

Desk. A slide already in the deck. A spreadsheet tab already in the file. A relative percent already on a dashboard you own. Cover the title. Read the axes. Try from zero. Write the line.

Floor or shift. A board number already posted. A shared screen already looked at in the bay. A count versus the thing it is a number *of*. Ordinary noticing. No public quiz in front of a supervisor. No undercover audit. Noticing is not a sting.

Placement is by skill, not job title. A senior person who still reports “jumped” without the axis is still in the early band of this chapter. A newer person who catches “this starts at 40,000” and writes the check is already practicing the later catch. A two-day “data storytelling” workshop is a receipt, not the work.

Practice that actually builds it

Named try-its. Time, materials, safety, the move. At least one lives at the desk. At least one lives on the floor or shift. Desk, floor, clinic, and shop *are* the practice. They do not replace the named move.

Try-it 1 — A deck chart, desk

Time. 8–15 minutes. Then stop.

⁸⁴Robert J. Stahl, “Using Think-Time and Wait-Time Skillfully in the Classroom,” ERIC Digest ED370885 (1994). Think-time; about three seconds as a convention, not a sacrament; written for classroom teachers. A working adult can use the move. This is not a workplace wait-time trial; none was opened for this book.

⁸⁵Mary Budd Rowe, papers on wait-time in elementary science class (NARST 1972 / ERIC ED061103; *Journal of Research in Science Teaching*, 1974). Typical waits of about one second after a question; extending both waits to 3–5 seconds changed discourse. Setting: science class. Useful as a picture of adult hurry. Not a stand-up randomised trial and not evidence that wait-time teaches generic critical thinking. Do not mash Rowe 1972 and Stahl 1994 as if they were one paper.

Materials. One slide already built, or the 48,000 / 56,000 picture you just drew. One comparison. One axis. Pencil. Scrap paper for the redraw from zero. Dated notebook for the sentence.

Safety. Curious, not courtroom. No humiliation of whoever built the slide. No public quiz of a junior colleague in front of a vice-president.

The move. What is the claim the chart wants? Does the vertical axis start at zero? What would the same numbers look like from zero? Average of what? Who is included?

How to run it. Put the picture down. Opening question: “What does this chart actually show?” Wait. Lift one number. Compare two bars, with the amount. Ask where the axis starts. Redraw from zero. Last thoughts: “I used to think it was a jump. Now I think the numbers went up by 8,000, and the picture made that look larger.” That sentence is allowed. It is not a confession ritual. Then write: *I claim the jump looks large because the axis starts at 40,000. We checked by redrawing from zero.*

Desk. You draw the chopped axis if you do not have a real slide. You lift and compare. You prompt the redraw. One written line is enough.

Floor or shift, if the same chart is on a shared screen. Same three jobs, spoken, quieter. Write the line later, at the desk, dated.

Try-it 2 — From-zero versus truncated, desk

Time. 5–10 minutes.

Materials. A series that starts at, say, 80 instead of 0 — unit sales, a count, a rainfall-shaped pair if that is what you have. Scrap paper. Two drawings of the *same* numbers.

Safety. Ordinary paper. You are reading a picture, not designing a trick to deploy next Tuesday.

The move. Redraw from zero. The leap shrinks. The numbers did not change. Struggle before rescue: read the axis before the caption’s adjective (“dramatic,” “soaring,” “collapsed”).

How to run it. Cover the title. Read the axes. Draw the chopped version if you only have the numbers. Draw the from-zero version. Opening question: “What

does this chart actually show?” Then: “What would the same numbers look like from zero?” Write one sentence a colleague who missed the meeting could follow.

Desk. Two drawings side by side. Stay until you can say, unaided, whether “jumped” was the numbers or the picture.

Floor or shift. If a posted pair of bars is already on a board, you may not be able to redraw on the board. Sketch the from-zero version on a pad in a break. Same move. Public quiet.

This is a bar-graph sitting. If the picture in front of you is a line, still ask where the axis starts, still lift the number, and keep the sentence honest: the bar finding is about bars.⁸⁶

Try-it 3 — A dashboard described, floor or shared screen

Time. 5–10 minutes at a board or a shared screen you already look at.

Materials. Tiles, a bar pair, a sparkline — described. No product chrome in later art, and none needed in the sitting. The question, not the screenshot’s pixels.

Safety. Ordinary looking. You already have access. No undercover audit of a workplace you do not work. No gotcha against the person who owns the dashboard.

The move. Each tile is a claim. What is this a number *of*? Rate versus count. Relative versus absolute. Where does the bar pair’s axis start?

How to run it. Point at one tile. Opening question: “What does this chart actually show?” Wait. If the tile says “up 40%,” ask: of what, from what, compared with

⁸⁶Brenda W. Yang, Camila Vargas Restrepo, Matthew L. Stanley, and Elizabeth J. Marsh, “Truncating Bar Graphs Persistently Misleads Viewers,” *Journal of Applied Research in Memory and Cognition* 10, no. 2 (2021): 298–311. Five studies; y-axis truncation led viewers to perceive illustrated differences as larger; 83.5 percent of participants across the five studies showed a truncation effect; Cohen’s *d* 0.48–1.26; graph literacy failed to predict the size of the effect; warnings reduced but did not eliminate it. Adult samples (MTurk, mean ages about 32–35; Study 5, Ph.D. students in quantitative fields and in the humanities). Even quantitative Ph.D. students with a warning immediately before the ratings still showed a truncation effect (Study 5, estimate 0.24). The authors’ recommendation is specific to bar graphs, where length encodes quantity; they do not generalise a “never truncate” rule to line graphs and dots. Useful study, not a promise that every workplace will see the same result.

what. If the bar pair is chopped, name the start. If the sparkline has no scale, name that missing piece. Write the sentence later if the bay cannot hold a notebook.

Desk cousin. The same dashboard on your laptop, after, not during, a live huddle.

Floor or shift. One tile. One question. Then back to the work. Cue if needed: “this is like the deck chart.”

Try-it 4 — Posted numbers versus the clock or the count, floor

Time. 3–8 minutes of a shift you were already working.

Materials. A board number and the thing it is a number *of*: units on the board versus units you can count; posted hours versus the clock; a clinic board’s wait-time versus the clock you already watch. A protocol already in use — not a diagnosis of a named patient in this book.

Safety. Ordinary looking. No public quiz. No lecture to a stranger’s customer. No sideline interrogation during a live procedure. Noticing is not a sting.

The move. Rate versus count. Relative versus absolute. Two records of the same interval.

How to run it. Look at the board. Opening question: “What does this number actually show? What is it a number *of*?” Wait. Then look at the clock, or the count, or the thing itself. Write later: *I claim the board said 12 because I read the tile. We checked by counting. The count was 9.* A colleague on the next shift could follow that.

Desk cousin. A dashboard tile versus the spreadsheet tab it is supposed to summarise.

Floor or shift. Spoken is enough in the bay. Ink after.

Talk/meeting box

Exact opening question. “What does this chart actually show?”

Not: “How do people lie with statistics?” Not: “What is the vocabulary word for a chopped axis?” Not: “Did we have a good meeting?”

Follow-ups.

1. Where does the axis start? (meaning of the picture)
2. What would the same numbers look like from zero? (reason / check)
3. Average of what? Who is included? Always this high? (alternative)
4. What's a second place we could look — the spreadsheet tab, the notes, the count? (second source)
5. Can you write that so a colleague who missed the meeting could follow it? You may change your mind about “jumped.”

How to wait. After you ask, a slow three. Three seconds is a convention from classroom think-time, not a workplace study.⁸⁷ After they stop, wait again. If they are mid-reason, do not cut them off. Look at the graph, the notes, or the clock, not at the person, if the silence is hard. After “write the sentence,” that is task-completion time, not a three-second wait. At work the shrug often looks like refusal. Do not rescue it. Wait in private when the power difference is steep. You are not citing science-class wait-time as a stand-up trial.⁸⁸

What a stuck silence usually means. They read the title instead of the axes. They offered a why-story (“it looks like we soared”). They think graph literacy is a vaccine. Wait-time 1 was zero. The question was a recitation item in disguise. They think “critical” means finding fault with whoever built the slide. They think a finished-looking paragraph that cites a chart is already a check. They think noticing is snitching.

If they are stuck. Smaller question. “What number is this bar?” Wait. Point at the axis. Do not fill the silence with the right answer. Do not turn the meeting into a fight about who is trying to trick whom.

Blocked, then mixed

Blocked. The first sitting is one deck chart. Two bars. Chopped axis, then redraw from zero. Stay there until lift, compare, and missing piece can happen on that

⁸⁷Robert J. Stahl, “Using Think-Time and Wait-Time Skillfully in the Classroom,” ERIC Digest ED370885 (1994). Think-time; about three seconds as a convention, not a sacrament; written for classroom teachers. A working adult can use the move. This is not a workplace wait-time trial; none was opened for this book.

⁸⁸Mary Budd Rowe, papers on wait-time in elementary science class (NARST 1972 / ERIC ED061103; *Journal of Research in Science Teaching*, 1974). Typical waits of about one second after a question; extending both waits to 3–5 seconds changed discourse. Setting: science class. Useful as a picture of adult hurry. Not a stand-up randomised trial and not evidence that wait-time teaches generic critical thinking. Do not mash Rowe 1972 and Stahl 1994 as if they were one paper.

picture.

Mixed. Then a dashboard tile. Then a board number on the floor. Then a relative percent with no base. Mix so the move is not glued to sales bars. Near transfer to another deck chart, another tile, another board is the honest hope. Far transfer to an unfamiliar later chart, unaided, is not a promise.⁸⁹

One incorrect example to diagnose

A finished-looking paragraph in a recap. The paragraph is fluent. Topic sentence: “Unit sales jumped hugely in May.” The writer cites “the chart.” The axis is never named. The numbers are never lifted. The paragraph looks done.

This is an illustration, not a reported workplace.

What it usually means. A caption wearing a topic sentence. The mash of a why that makes sense — the bars *look* tall — with evidence that the claim is true.

What to say. “What does this chart actually show? Where does the axis start?” Wait. Redraw from zero. Then: “Can you write that so a colleague who missed the meeting could follow it?” If the sentence cannot survive the check, correct the sentence. That is the work. Silence in the face of a false caption is not kindness. “Never correct a colleague” is not this book’s rule.

If it isn’t clicking

The caption still arrives as the look. Slow down. Cover the title. Point at the axis. Lift one number out loud. Stay on two bars until the amount has a name. If a tool is finishing the caption, close the window. The unaided sentence is yours. If you cannot yet lift a number after several sittings, stay here. Do not add a statistics course.

⁸⁹Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, “How to Teach Critical Thinking,” NSW Department of Education, 2019. Thinking is not a bicycle skill you apply in any situation; processes are intertwined with content; maxims without knowledge will not implement. Friel, Curcio, and Bright 2001: “read beyond the data” requires prior knowledge related to the graph. Near practice on a deck chart does not, by itself, install unaided judgment on an unfamiliar later chart. Cue the similarity (“this is like the deck chart”).

The sitting becomes a sermon about liars, or a lecture about statistics. Hear it. Return to this picture. “Where does *this* axis start?” A chopped axis is a design choice. Some pictures start at zero and show what they claim. The move is the look, not a personality of distrust. If a real disagreement about the year is too hot for the table today, pick the two bars. Practice the move. A later, heavier claim will need knowledge this sitting does not have.⁹⁰

You treat literacy, a warning, or a degree as immunity. Graph literacy did not predict the size of the effect. Warnings reduced it and did not eliminate it. Advanced quantitative training did not make people immune.⁹¹ Hear that without shame. Then lift the number anyway. If you skip the axis because you “already know,” that is information. The next move is the redraw, not a new course title.

Go ahead when you can lift, compare with an amount, catch a chopped axis on a bar pair, and write one sentence a colleague who missed the meeting could follow. Stay longer when the caption still arrives as the look. A job title does not settle it. Ask a human colleague when the metric’s definition is the wall — what a “unit sale” is, who is in the count. That is knowledge first, not a character flaw. A human who already owns the spreadsheet is the right person. A chatbot is not.

Correct a wrong reason without crushing the talk. “I hear *jumped*. The strongest version of what you’re seeing is that May’s bar is taller. I’m not sure that reason works as a leap, because the axis starts at 40,000. What would the same numbers

⁹⁰Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, “How to Teach Critical Thinking,” NSW Department of Education, 2019. Thinking is not a bicycle skill you apply in any situation; processes are intertwined with content; maxims without knowledge will not implement. Friel, Curcio, and Bright 2001: “read beyond the data” requires prior knowledge related to the graph. Near practice on a deck chart does not, by itself, install unaided judgment on an unfamiliar later chart. Cue the similarity (“this is like the deck chart”).

⁹¹Brenda W. Yang, Camila Vargas Restrepo, Matthew L. Stanley, and Elizabeth J. Marsh, “Truncating Bar Graphs Persistently Misleads Viewers,” *Journal of Applied Research in Memory and Cognition* 10, no. 2 (2021): 298–311. Five studies; y-axis truncation led viewers to perceive illustrated differences as larger; 83.5 percent of participants across the five studies showed a truncation effect; Cohen’s *d* 0.48–1.26; graph literacy failed to predict the size of the effect; warnings reduced but did not eliminate it. Adult samples (MTurk, mean ages about 32–35; Study 5, Ph.D. students in quantitative fields and in the humanities). Even quantitative Ph.D. students with a warning immediately before the ratings still showed a truncation effect (Study 5, estimate 0.24). The authors’ recommendation is specific to bar graphs, where length encodes quantity; they do not generalise a “never truncate” rule to line graphs and dots. Useful study, not a promise that every workplace will see the same result.

look like from zero?” Then look. Then write.

Tools, including AI

Point back to the work-hour box. Keep this short.

After you have tried, you may ask a tool, on an account you hold, to explain *this page to yourself* from a named source — the slide number, the spreadsheet tab, the board’s posted title. Extra questions you then vet. Throw out any that import a live controversy, a secret, a pretend expert, a colleague’s private performance, or a civic-web item. A labelled SCRIPT after the try is a short sequence for *you* to say — not a paragraph to copy as your reason, your email, or your recap. A hint after an attempt is not the finished reason. A diagnosis of a reason you already gave is yours to hear and decide.

You may not park an unsupervised chatbot as the only partner. You may not ask a tool to invent last month’s numbers, whether the axis started at zero, or what the board said. You may not ask it to write your sentence and paste that text as yours. You may not invent an expert identity and treat that voice as a person. You may not grade a colleague’s slide from a detector score. You may not print a certificate of thinking. You may not hand yourself an agent as a thinking partner for the first attempt. You may not put confidential workplace data into a public consumer tool.

On some consulting tasks inside one model’s frontier, adults using the model were faster and higher-rated; on a designed outside-frontier task they were less likely to be right.⁹² A fluent memo is not a check that the chart’s leap was the numbers. A faster caption is performance. It is not yet your reason.

One crutch sentence, labelled, then teach: in a high-school math trial, an unguarded chatbot made practice look better and left students worse on the test with the window closed.⁹³ That is not a workplace trial. The closed window still applies.

⁹²Fabrizio Dell’Acqua and colleagues, “Navigating the Jagged Technological Frontier,” Harvard Business School working paper, 2023. About 758 consultants; inside-frontier tasks faster and higher-rated; outside-frontier, 19 percentage points less likely to be correct. In-band, labelled. Full working-paper PDF not line-checked this pass. A fluent memo is not a check that the task was inside the frontier.

⁹³Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman, “Generative AI Can Harm Learning,” *Proceedings of the National Academy of Sciences*, 2025. Field randomised trial, nearly 1,000 Turkish high-school math students. Unguarded GPT: practice looked better, then unaided isomorphic exam worse versus never-AI. One crutch sentence, labelled

After any tool, write the sentence from the picture itself, tool closed.

Open *your* employer's acceptable-use page. Label it. A fluent reason from a machine is performance, not your check. A thinking tool is not a secret friend.

What “done enough” looks like

Done enough is skill, not a job title, and not a course called Graph Thinking 1. The title a stranger can read is still the job you hold — analyst, lead, clerk, tech, nurse, manager. The dated note lives inside that file. It is not Critical Thinking I. A “critical thinking workshop certificate” is a receipt, not the work.

Checklist before moving on.

- Lifts a number from a graph or a table, with units.
- Compares two bars or two rows, with an amount, not only a feeling.
- Asks, or can answer, “where does this picture start?”
- Redraws a chopped-axis bar pair from zero, or can say what that redraw would do.
- Names one thing the picture does not tell them — who is included, what is missing, what two months are not.
- Hears “jumped” or “dramatic” as a caption, and, with a question, gets to a number.
- Writes one dated sentence a colleague who missed the meeting could follow: claim, because, how we checked.
- Changes the sentence when the redraw or the tab goes the other way, without treating that as losing.
- On a dashboard tile or a floor board, can use the cue “this is like the deck chart.”
- Does not treat graph literacy, a warning caption, or a degree as immunity.

Desk, done enough. Spoken sentence plus one written line that names the claim and the check, on a slide you already have.

Floor or shift, done enough. One tile or one board number named as a number of something, then a later dated line. Sometimes a closed-window paragraph after any tool.

not a workplace trial.

What can still be thin: unaided far transfer to an unfamiliar later chart with no cue; a multivariable data set; reading beyond without the domain knowledge; treating a cropped *line* as if the bar finding had already covered it. Those can wait. They are not a window that closed. Graph sense is not a bicycle. Knowledge of the metric, of who is in the count, of what the board is counting, still has to be in the room.⁹⁴ Adulthood is not a thinking birthday. A job is not arrival.

An adult who felt rusty should be able to say, after this chapter: I understand the three levels. I know what to say tomorrow at the desk and on the floor. I know what a caption offered as a check sounds like. I know what a finished-looking paragraph without a named axis looks like. I can lift the number. I can do this.

⁹⁴Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, “How to Teach Critical Thinking,” NSW Department of Education, 2019. Thinking is not a bicycle skill you apply in any situation; processes are intertwined with content; maxims without knowledge will not implement. Friel, Curcio, and Bright 2001: “read beyond the data” requires prior knowledge related to the graph. Near practice on a deck chart does not, by itself, install unaided judgment on an unfamiliar later chart. Cue the similarity (“this is like the deck chart”).

Chapter 6

A short written reason




Figure 13: An open dated notebook on a desk, a short reason in ink, two named sources, a pencil, no diploma seal, no logos, no people

Why this matters

Can you write that so a colleague who missed the meeting could follow it?

That question unlocks this book's added object. Talk is cheap in a team that already talks. A sentence on paper is a different demand. You have to pick a claim, pick a because, and pick a check that a colleague — someone who was not in the room, does not share the running joke, does not already know which "they" you mean — could follow.

The shape, once:

I claim X because Y. We checked by Z.

Optionally a second sentence: **Someone who disagrees would say...**

Dated. Stimulus named. Sources named.

That is closer to a decision log a stranger can read than it is to a shrug. It is not a five-paragraph course. It is not a programme for raising a national writing score. It is not a college thinking inventory. It is not a Watson–Glaser sitting. It is not generic critical thinking in a bottle.

Extra essay writing without an interlocutor is not the intervention. In a multi-year classroom study of dialogic practice, students who wrote more essays plus whole-class discussion did not match the dialogic group's dual-perspective on later individual essays. Extra writing did not help. The students were ages eleven to twelve at entry. Copy the *shape*: talk both sides, then write a product. Copy not the topics.⁹⁵ The missing interlocutor has to have been in the room, or on the floor. Talk first. Write. Then look. If the sentence cannot survive the check, correct the sentence.

A large research review of thinking instruction found, for graduate and adult samples, a small average effect whose confidence interval includes zero. Seventeen studies. That is the honest adult cell. It is not a promise that a paperback will move a standardised test, and it is not a workplace trial of this notebook sentence.⁹⁶

⁹⁵Deanna Kuhn and Amanda Crowell, "Dialogic Argumentation as a Bridge to Argumentative Thinking," *Psychological Science* 22, no. 4 (2011): 545–552. Ages 11–12 at entry; extra essay writing plus whole-class discussion did not match the dialogic group's dual-perspective on later individual essays. Copy shape (dialogic dual-perspective, then a written product), not topics. Access date for sources in these notes: 2 September 2026.

⁹⁶Philip C. Abrami, Robert M. Bernard, Eugene Borokhovski, David I. Waddington, C. Anne

Undergraduate numbers from the same review are the wrong band to steal. A combination of authentic problems, dialogue, and mentoring that looks larger is mixed-age school work, not a job result. Then teach the sentence.

That is worth the struggle because the job already produces fluent talk. A recap's *we decided*. A vendor's *leading*. A handoff's *always*. A finished-looking paragraph that cites "the call." A chart that "jumped." Each of those can be spoken and still vanish. Ink holds the claim still long enough for a colleague — and for next Tuesday — to see whether the because was a look or a why-story. The adult who can hear a finished-looking paragraph offered as a reason can practice this week. The adult who can write one dated sentence, then change it when the look goes the other way, has practiced deciding what to believe or do, with reasons, on a domain they actually know.

This week you can learn the four-part shape — claim, because, other side, check — and that extra essays without someone in the room are not the intervention. Today, at the desk, you can write one dated sentence after a meeting so a colleague who wasn't there could follow it. Today, on the floor, you can write one dated line after a handoff so the next shift could follow it.

Understand it yourself

You do not need to be a writing teacher. You do need to hear a fluent paragraph with no named source as a performance, not yet a check, and "because we wanted it more" as a why-story in ink.

Everyday picture. A meeting. Someone says the vendor is industry-leading. Someone else says last quarter's tickets were slow. Talk can loop. Then a notebook:

2 September 2026. Stimulus: Tuesday vendor call plus the spec sheet, clause

Wade, and Tonje Persson, "Strategies for Teaching Students to Think Critically: A Meta-Analysis," *Review of Educational Research* 85, no. 2 (2015): 275–314. Overall $g^+ = 0.30$, mixed ages, standardised thinking measures. Graduate and adult $k = 17$, $g = 0.21$; 95 percent confidence interval about -0.05 to 0.46 , not significantly greater than zero. Undergraduate $k = 126$, $g = 0.26$ is the wrong band to quote as workplace. Combination of authentic instruction, dialogue, and mentoring $g = 0.57$ is 19 mixed-age studies, not a job result. High school 16–18, $k = 71$, $g = 0.25$, contrast only. Labelled, not a promise that a paperback will move a standardised test. Abrami and colleagues, 2008 *RER*, remains the sourced sentence that improvement cannot be a matter of implicit expectation; do not pick 2008's larger adult cell and hide 2015.

4.2. *I claim we should not switch this month because the spec still lists a 48-hour response and last quarter's tickets sat at 72 hours. We checked by opening clause 4.2 and the ticket log for July. Someone who disagrees would say the new page's "leading" still captures how the demo felt.*

A colleague who missed the call can follow that. They could open the same clause and the same log.

That paragraph is an illustration, not a reported workplace. Your clause numbers will be different. The shape will not.

Precise picture. A short written reason is a product with parts. Claim: a sentence that can be true or not, or fair or not, about something we can look at. Because: evidence *that*, not only a story of why it would make sense. Check: a named look — a second page, a spreadsheet tab, a roster, a spec sheet, last quarter's process document, the axis on a chart. Other side: the strongest fair version of what someone who disagrees would say, written so they would nod. Date. Stimulus. Sources. If the colleague would have to have been there to know what "it" and "they" refer to, the reason is not yet a reason.

Paper first. Three boxes on scrap — Claim / Because / Unless — then a fifth line: the sentence a colleague could read plus the check.⁹⁷ A map on a screen can wait. The sentence cannot. Three boxes on paper is a type, not a university semester, and not a standardised-test gym.

You are allowed to correct the sentence. "Never correct a colleague" is not this book's rule. Hearing a wrong because, then asking a better question, then naming a better reason if needed, is the work. Silence in the face of a false written reason leaves the error in the file.

A school assignment that asks for a paragraph is a neighbour, not this chapter's course. National writing sittings at grade twelve, and counts of how often high-school assignments asked for a paragraph, are facts about *school writing*. They are not a thinking score, and they are not a workplace statistic this notebook will close.⁹⁸ Strategy instruction can improve writing quality in school grades. That is

⁹⁷Tim van Gelder, "Using Argument Mapping to Improve Critical Thinking Skills," in *The Palgrave Handbook of Critical Thinking in Higher Education*, ed. Martin Davies and Ronald Barnett (2015). University argument mapping; "Unfortunately, no such study has ever been conducted" — no large randomised trial, his words. For this book: three boxes on paper (claim, because, objection) before any app. Type, not the course, not a college inventory.

⁹⁸National Assessment of Educational Progress writing 2011, about 27 percent of twelfth-

writing quality, not a thinking inventory.⁹⁹ The portable piece: do not say “write it up.” Say the parts.

Struggle before rescue: write the sentence before the slide template and before the chatbot. A faster, prettier email can be a substitute for your reason.¹⁰⁰ If a tool named a paper, open the paper.¹⁰¹

Wrong answers you should be able to hear

Including from yourself.

1. A fluent paragraph with no named source. The page looks finished. Topic sentences. No check. “They” and “it” only make sense if you were in the room.

What it usually means. Length was treated as the job. A why-story wearing professional clothes.

What to say. “Can you write that so a colleague who missed the meeting could follow it? What is the claim, in one sentence? Because — is that evidence *that*, or a why-story? How did we check? Who is *they*?”

2. Extra writing as the whole method. “We’ll just write more.” Or a chatbot paragraph offered as help. Or a five-paragraph night.

graders at or above Proficient — school writing, not critical thinking, not a workplace statistic. Arthur N. Applebee and Judith A. Langer, high-school assignment counts (about one in five asking for a paragraph or more) — school writing. Steve Graham and Dolores Perin, 2007, writing-quality effects in grades 4–12 — writing quality, not a thinking inventory. Ennis–Weir essay test (1985) is a cousin of “a sentence a colleague could read,” not a workplace trial.

⁹⁹National Assessment of Educational Progress writing 2011, about 27 percent of twelfth-graders at or above Proficient — school writing, not critical thinking, not a workplace statistic. Arthur N. Applebee and Judith A. Langer, high-school assignment counts (about one in five asking for a paragraph or more) — school writing. Steve Graham and Dolores Perin, 2007, writing-quality effects in grades 4–12 — writing quality, not a thinking inventory. Ennis–Weir essay test (1985) is a cousin of “a sentence a colleague could read,” not a workplace trial.

¹⁰⁰Shakked Noy and Whitney Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science* 381 (2023): 187–192. 453 college-educated professionals; time down about 40 percent; quality up about 18 percent. MIT working paper (labelled draft): many treated writers submitted the first pass unedited. A faster email can substitute for effort. Not a leaving-the-page trial.

¹⁰¹William H. Walters and Esther Isabelle Wilder, 2023, fabricated citations in generated text (about 55 percent of GPT-3.5 citations and 18 percent of GPT-4 citations in that study). In-band for adults writing at work: if a tool named a paper, open the paper.

What it usually means. Writing practice was confused with thinking practice. The interlocutor left the room. A machine finished the sentence. Extra essays without dual-perspective talk did not produce the later dual-perspective page, in the study this book is copying the shape from, with eleven- and twelve-year-olds.¹⁰²

What to say. Talk first. Then one dated sentence. Then look. If a tool is in the room, it does not write this paragraph as yours. Point back to the work-hour box.

3. A steelman that is actually a straw, or steelman heard as giving up. “They just want to be difficult.” Or: “If I write their side, I lose.” Or a machine steelman pasted as yours.

What it usually means. Dual focus was skipped, or it was heard as surrender, or the first attempt was handed to a tool.

What to say. “Write it so they would nod. Then we still look. Changing the sentence if the check goes the other way is not losing.”

4. A test sitting treated as the gym. A college thinking inventory, a Watson–Glaser sitting, a five-paragraph drill, offered as this week’s hour.

What it usually means. An instrument was asked to do the work of the sentence. Those sittings explain what some schemes sit. They are not this notebook.¹⁰³

What to say. Date the notebook. Name the stimulus. Name the sources. Keep the title as the job you hold.

5. “Write it up” with no other side, offered as thinking. A culture that asks

¹⁰²Deanna Kuhn and Amanda Crowell, “Dialogic Argumentation as a Bridge to Argumentative Thinking,” *Psychological Science* 22, no. 4 (2011): 545–552. Ages 11–12 at entry; extra essay writing plus whole-class discussion did not match the dialogic group’s dual-perspective on later individual essays. Copy shape (dialogic dual-perspective, then a written product), not topics. Access date for sources in these notes: 2 September 2026.

¹⁰³Philip C. Abrami, Robert M. Bernard, Eugene Borokhovski, David I. Waddington, C. Anne Wade, and Tonje Persson, “Strategies for Teaching Students to Think Critically: A Meta-Analysis,” *Review of Educational Research* 85, no. 2 (2015): 275–314. Overall $g^+ = 0.30$, mixed ages, standardised thinking measures. Graduate and adult $k = 17$, $g = 0.21$; 95 percent confidence interval about -0.05 to 0.46 , not significantly greater than zero. Undergraduate $k = 126$, $g = 0.26$ is the wrong band to quote as workplace. Combination of authentic instruction, dialogue, and mentoring $g = 0.57$ is 19 mixed-age studies, not a job result. High school 16–18, $k = 71$, $g = 0.25$, contrast only. Labelled, not a promise that a paperback will move a standardised test. Abrami and colleagues, 2008 *RER*, remains the sourced sentence that improvement cannot be a matter of implicit expectation; do not pick 2008’s larger adult cell and hide 2015.

for a recap and never requires the other side in speech is running the comparison condition from that classroom study: more writing, less dual perspective.¹⁰⁴

What it usually means. Ink without an interlocutor. Own-side exposition wearing a date.

What to say. “Someone who disagrees would say... Write it so they would nod. Then the check.”

Five-minute warm-up

Before the sitting, with no one else in the room:

1. Write the shape on scrap: I claim X because Y. We checked by Z. Someone who disagrees would say...
2. Say out loud: “Can you write that so a colleague who missed the meeting could follow it?” Wait a slow three.¹⁰⁵ Then wait as if you were writing. That second wait is task-completion time.
3. Write one fluent fake you expect — a paragraph with no source, a why-story in ink, a machine sentence, an “it” that needs the meeting — and the better ask: claim, because, check, other side. Then look.

If those three are fluent, you can hear a finished-looking page offered as a reason. That is the job. You still write the sentence.

How to try it this week

Use the work-hour shape. Fill it with a claim you already have: a recap versus notes, a vendor call versus a spec, two handoff notes, a chart caption from Chapter 5, a “we always” versus last quarter’s process document. Fifteen to twenty-five minutes on a named artifact. Stop while it is still a sitting.

¹⁰⁴Deanna Kuhn and Amanda Crowell, “Dialogic Argumentation as a Bridge to Argumentative Thinking,” *Psychological Science* 22, no. 4 (2011): 545–552. Ages 11–12 at entry; extra essay writing plus whole-class discussion did not match the dialogic group’s dual-perspective on later individual essays. Copy shape (dialogic dual-perspective, then a written product), not topics. Access date for sources in these notes: 2 September 2026.

¹⁰⁵Robert J. Stahl, “Using Think-Time and Wait-Time Skillfully in the Classroom,” ERIC Digest ED370885 (1994). About three seconds as a convention, classroom origin. After “write the sentence,” that is task-completion time. Not a workplace trial.

Warm-up (2–4 minutes). Yesterday’s graph or yesterday’s two-source look, from memory. “What did we look at? What did we used to think?” Then: “Today we write it so a colleague who missed the meeting could follow it.”

Short model (5–8 minutes).




Figure 14: An open notebook, a dated line, two named sources, a check; desk; no person, no logo, no diploma seal

You talk first. A recap that uses *we decided*, and the notes. You say a sloppy spoken version on purpose. Then you write, slowly, where you can see it:

2 September 2026. Stimulus: recap of Tuesday’s stand-up plus the notes in the shared doc. I claim we did not decide to switch vendors because the notes list “gather two quotes” and no vote. We checked by opening the notes, 10:40 a.m. line. Someone who disagrees would say the recap’s word still captures how the room felt.

You wait. You ask whether a colleague who was not in the stand-up could follow that. Then you show the *wrong* cousin on purpose: a fluent paragraph that never names the notes, full of “it” and “they.” “Is that a check?” Wait. Then stop talking sooner than feels polite.

This pair is an illustration, not a reported workplace. Use *your* recap if you have one.

Exact wording you can steal.

On the page: “Can you write that so a colleague who missed the meeting could follow it?”

On the claim: “What is the claim, in one sentence?”

On the because: “Because what? Is that evidence *that*, or a why-story?”

On the check: “How did we check? What did we open?”

On the pronouns: “Who is *they*? What is *it*?”

On the other side: “Someone who disagrees would say... Write it so they would nod.”

On mind-change: “You may change the sentence if the check goes the other way. That is not losing.”

Your attempt. After the model, you write the next one — last meeting, last vendor call, last handoff — unaided. Product: one dated sentence. Claim, because, check. Sources named. You wait, including with yourself. You do not finish the sentence from a template. You do not paste a machine’s paragraph as yours.

One good question. After the attempt: “If a colleague opened only this sentence, could they find the file?” You generate the why.

Mixed practice. One meeting sentence. One vendor-call sentence. One handoff line. Mix, so the move is not glued to stand-ups.

Written line (the sitting is the line). Ten to twenty minutes is task-completion time, not a three-second wait. Tool closed.

How to fade help. You name all four parts (claim, because, other side, check) → you fill the check while a sticky note names the rest → you name all four → you catch an “it” that needs the meeting → you write without the boxes. When a colleague who wasn’t there could follow the sentence, put the sticky-note scaffold

away. Bring it back when a new genre arrives — a clinic note, a warehouse count, a review versus a spec.

When to stop talking. After one model and one question. You hold the key: is there a claim, a because, a named look, a wait? You do not become a lecture about essays. The word *argumentation* can wait. The sentence cannot. In a meeting, ask for the line after, not as a public quiz. If the power difference is steep, wait in private.

Desk / floor-or-shift

Desk. After a meeting, after a vendor call, after a thread already in the inbox. Dated notebook or a note in the project folder. The title is still the project. It is not Written Reasoning I.

Floor or shift. After a handoff, not during a live procedure. Two notes plus the check. The next shift could follow it without having been there. Ordinary noticing. No public quiz. No undercover audit.

Placement is by skill, not job title. A senior person whose recap still needs the meeting to make sense is still in the early band of this chapter. A newer person who names the file and the check is already practicing the later catch.

Practice that actually builds it

Named try-its. Time, materials, safety, the move. At least one lives at the desk. At least one lives on the floor or shift.

Try-it 1 — After a meeting, desk

Time. 10–20 minutes. Then stop.

Materials. A dated notebook or a note in the project folder. The recap already in the inbox, if there is one, and the notes if you have them.

Safety. Curious, not courtroom. No humiliation of whoever wrote the recap. No public quiz of a junior colleague in front of a vice-president.

The move. One dated sentence — claim, because, check — that a colleague who was not in the room could follow.

How to run it. Talk first, even if only with yourself: what do we think we decided? Then write. Opening question: “Can you write that so a colleague who missed the meeting could follow it?” Name the stimulus. Name the sources. If “it” needs the meeting, rewrite. Last thoughts: “I used to think we decided. Now I think the notes list a next step, not a vote.” Then look. If the sentence cannot survive the look, correct the sentence.

Desk. Ink at the desk. Title still the project.

Floor or shift cousin. If the “meeting” was a huddle in the bay, write after, in a break. Same shape.

Try-it 2 — After a vendor call, desk

Time. 10–20 minutes.

Materials. The call notes and one named source: the spec, the quote, the contract clause.

Safety. Ordinary looking. No gotcha against a vendor on a live call. After, not during. No confidential paste into a public tool.

The move. Same shape. Leave the glowing verb. Open the clause. If “it” needs the call, rewrite.

How to run it. Write the glowing verb you heard. Then write the check. Opening question: “Can you write that so a colleague who missed the call could follow it?” Someone who disagrees would say... Then look at the named source.

Desk. The spec sheet next to the notes.

Floor or shift cousin. After a walk-through, two quotes for the same SKU, written so the next person could follow the difference without having been on the floor.

Try-it 3 — After a handoff, floor

Time. 5–10 minutes after, not during.

Materials. The two notes plus the check: a board, a count, a protocol already in use. Clinic bay, shop floor, warehouse line, retail counter — the place you already work.

Safety. After, not during, a live procedure. No public quiz. No named patient in this book. No undercover audit. Noticing is not a sting.

The move. Write the disagreement so the next shift could follow it without having been there.

How to run it. Read outgoing. Read incoming. Where they disagree is the fact that needs a third look. Opening question: “Can you write that so the next shift could follow it?” Then look. Then write: *I claim the count is 9 because I counted the bay at 14:10. We checked against the outgoing note, which said 12. Someone who disagrees would say the board had not been updated.* A colleague who starts at 15:00 could follow that.

Desk cousin. Two meeting notes on the same hour, written so someone who missed both could still find the file.

Floor or shift. Spoken in the bay if ink is impossible; ink as soon as the bay allows.

Try-it 4 — Steelman on paper, desk

Time. 8–15 minutes.

Materials. A disagreement already in the thread. Scrap paper. Dated notebook.

Safety. Write it so they would nod. Inventing a better argument they did not make, then knocking it down, is not steelman. Asking a machine to write the steelman so you can paste it is write-my-reason.

The move. Someone who disagrees would say... Then the check.

How to run it. Write yours in one sentence. Write theirs so they would nod. Opening question: “Can you write that so a colleague who missed the meeting could follow it?” Then look at one named source. If the check goes the other way, change *your* sentence. That is not losing.

Desk. Thread plus one source.

Floor or shift cousin. “We always start with X” versus last quarter’s process document, written so the next shift could follow both sides.

Talk/meeting box

Exact opening question. “Can you write that so a colleague who missed the meeting could follow it?”

Not: “Write a five-paragraph essay.” Not: “What is the vocabulary word for a claim?” Not: “Did we have a good meeting?”

Follow-ups.

1. What is the claim, in one sentence? (meaning)
2. Because what? Is that evidence *that*, or a why-story? (reason)
3. Always? What else could be true? (alternative)
4. How did we check? Who is “they”? What’s a second place we could look? (check / second source)
5. What would someone who disagrees say? You may change the sentence if the check goes the other way. (steelman / revise)

How to wait. After you ask, a slow three.¹⁰⁶ After they stop, wait again. After “write the sentence,” that is task-completion time, not a three-second wait. Look at the notebook, the notes, the spec, or the roster, not at the person, if the silence is hard. At work the shrug often looks like refusal. Do not rescue it. Wait in private when the power difference is steep. You are not citing science-class wait-time as a workplace trial.¹⁰⁷

What a stuck silence usually means. They think a finished-looking paragraph is already a check. They think more writing is the treatment. They think a college inventory is the gym. They pasted a model’s paragraph as theirs. Wait-time 1 was zero. The question was a recitation item in disguise. They offered a why-story in ink. They think steelman on paper means giving up. They think “critical” means finding fault.

If they are stuck. Smaller question. “What is the claim, in one sentence?” Wait. Point at the object. Do not fill the silence with your paragraph. Do not turn the meeting into a fight.

¹⁰⁶Robert J. Stahl, “Using Think-Time and Wait-Time Skillfully in the Classroom,” ERIC Digest ED370885 (1994). About three seconds as a convention, classroom origin. After “write the sentence,” that is task-completion time. Not a workplace trial.

¹⁰⁷Mary Budd Rowe, wait-time papers, elementary science class, 1972/1974. Not a stand-up randomised trial. Do not mash Rowe and Stahl as if they were one paper.

Blocked, then mixed

Blocked. First sitting: one claim you already talked about, one look you can actually make, one dated line. Stay there until the parts have names.

Mixed. Then steelman on paper. Then a vendor call. Then a handoff. Mix so the move is not glued to the first object. Near transfer to another recap, another call, another handoff is the honest hope. Far transfer to an unseen later essay about a domain you do not know is not a promise.¹⁰⁸

One incorrect example to diagnose

A finished-looking paragraph in the project folder. Five sentences. Topic sentence. “The vendor is leading.” No clause number. No date. No log. “They said it on the call.” The page looks like work.

This is an illustration, not a reported workplace.

What it usually means. A fluent recap offered as a check. The mash of a why that makes sense — vendors describe themselves as leading — with evidence that this claim is true. A model’s paragraph, or a human one, can wear the same clothes.

What to say. “Can you write that so a colleague who missed the meeting could follow it? What’s a second place we could look?” Open a named page. If the sentence cannot survive the check, correct the sentence.

If it isn’t clicking

You write a lot and name nothing. Slow down. Three boxes on paper, then one line. You name the four parts. You fill the check. Fade the blanks. If a tool is finishing the paragraph, close the window. The unaided sentence is yours. If you cannot yet hold one claim and one look after several sittings, stay here. Do not add a five-paragraph night. Do not add a college inventory as the week’s hour.

You think steelman means giving up, or you write a straw. Hear it. “Write it so they would nod.” If needed, model two fair sentences, then write the next

¹⁰⁸Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, NSW occasional paper, 2019. Thinking is not a bicycle; maxims without knowledge will not implement. Near practice on a recap does not, by itself, install unaided judgment on an unseen later artifact.

two. Then look. Dual focus is promptable. It is not the default.¹⁰⁹ Correct a straw without crushing the talk: “Is that a version they would sign?” Inventing a better argument they did not make, then knocking it down, is not steelman.

You want more writing, or a test, or a machine, to do the thinking. Extra writing without an interlocutor is not the intervention.¹¹⁰ A college inventory is not this notebook.¹¹¹ A fluent machine paragraph is performance.¹¹² Return to talk first, then one dated sentence, then a look you can actually make. If writing itself is the wall — a freeze at the blank page, a workplace that only accepts a slide template — get a human colleague who already writes the file you have to write. You still own the thinking question: could a colleague who missed the meeting follow the because and the check?

Go ahead when you can write claim, because, check, and a fair other side on a known stake, and change the sentence when the look goes the other way. Stay longer when the page is still a performance. A job title does not settle it. Ask a human colleague when the domain is the wall — what the spec lists, what the

¹⁰⁹Deanna Kuhn and Amanda Crowell, “Dialogic Argumentation as a Bridge to Argumentative Thinking,” *Psychological Science* 22, no. 4 (2011): 545–552. Ages 11–12 at entry; extra essay writing plus whole-class discussion did not match the dialogic group’s dual-perspective on later individual essays. Copy shape (dialogic dual-perspective, then a written product), not topics. Access date for sources in these notes: 2 September 2026.

¹¹⁰Deanna Kuhn and Amanda Crowell, “Dialogic Argumentation as a Bridge to Argumentative Thinking,” *Psychological Science* 22, no. 4 (2011): 545–552. Ages 11–12 at entry; extra essay writing plus whole-class discussion did not match the dialogic group’s dual-perspective on later individual essays. Copy shape (dialogic dual-perspective, then a written product), not topics. Access date for sources in these notes: 2 September 2026.

¹¹¹Philip C. Abrami, Robert M. Bernard, Eugene Borokhovski, David I. Waddington, C. Anne Wade, and Tonje Persson, “Strategies for Teaching Students to Think Critically: A Meta-Analysis,” *Review of Educational Research* 85, no. 2 (2015): 275–314. Overall $g^+ = 0.30$, mixed ages, standardised thinking measures. Graduate and adult $k = 17$, $g = 0.21$; 95 percent confidence interval about -0.05 to 0.46 , not significantly greater than zero. Undergraduate $k = 126$, $g = 0.26$ is the wrong band to quote as workplace. Combination of authentic instruction, dialogue, and mentoring $g = 0.57$ is 19 mixed-age studies, not a job result. High school 16–18, $k = 71$, $g = 0.25$, contrast only. Labelled, not a promise that a paperback will move a standardised test. Abrami and colleagues, 2008 *RER*, remains the sourced sentence that improvement cannot be a matter of implicit expectation; do not pick 2008’s larger adult cell and hide 2015.

¹¹²Shakked Noy and Whitney Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science* 381 (2023): 187–192. 453 college-educated professionals; time down about 40 percent; quality up about 18 percent. MIT working paper (labelled draft): many treated writers submitted the first pass unedited. A faster email can substitute for effort. Not a leaving-the-page trial.

roster says. That is knowledge first.

Tools, including AI

Point back to the work-hour box. Keep this short.

After you have tried, you may ask a tool to explain this page to yourself from a named source — title, URL, page number, contract clause, spreadsheet tab — or to give extra questions you then vet. Throw out any that import a live controversy, a secret, a pretend expert, a colleague’s private performance, or a civic-web item. A labelled SCRIPT after the try is a prompt for *you* to say — not a paragraph to copy as your reason, your email, or your steelman. A hint after an attempt is not the finished reason. A diagnosis of a reason you already gave is yours to hear and decide.

You may not ask a tool to write your reason, email, or steelman, then paste that text as yours. You may not park an unsupervised chatbot as the only partner. You may not invent an expert identity. You may not grade a colleague’s page from a detector score. You may not print a certificate of thinking. You may not ask a tool to invent what the spec lists, what the roster says, or what the spreadsheet totals. You may not hand yourself an agent as a thinking partner for the first attempt. You may not put confidential workplace data into a public consumer tool. You may not paste undisclosed generated text as yours where a named employer forbids it.

On mid-level professional writing, one study found a chatbot cut time and raised grader scores and mostly substituted for effort — many treated writers submitted the first pass.¹¹³ A faster email is performance. It is not yet your reason. If a tool named a paper, open the paper.¹¹⁴ One college-writing sentence, then stop: a university writing study warned of “metacognitive laziness.” That is college writing, not a workplace trial.¹¹⁵ One crutch sentence, labelled: an unguarded

¹¹³Shakked Noy and Whitney Zhang, “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” *Science* 381 (2023): 187–192. 453 college-educated professionals; time down about 40 percent; quality up about 18 percent. MIT working paper (labelled draft): many treated writers submitted the first pass unedited. A faster email can substitute for effort. Not a leaving-the-page trial.

¹¹⁴William H. Walters and Esther Isabelle Wilder, 2023, fabricated citations in generated text (about 55 percent of GPT-3.5 citations and 18 percent of GPT-4 citations in that study). In-band for adults writing at work: if a tool named a paper, open the paper.

¹¹⁵Yizhou Fan and colleagues, *British Journal of Educational Technology*, 2025. 117 university

chatbot can make practice look better and leave you worse when the window is closed.¹¹⁶ After any tool, write the sentence from the file itself, tool closed.

Open *your* employer's acceptable-use page. Label it. A fluent reason from a machine is performance, not your check. A thinking tool is not a secret friend.

What “done enough” looks like

Done enough is skill, not a job title, and not a course called Written Reasoning I. The title a stranger can read is still the job you hold. The dated note lives inside that file. It is not Critical Thinking I. A workshop certificate is a receipt, not the work. A percentile on a college inventory is not the year's work.

The honest adult cell in a large review of thinking instruction is modest, and the interval includes zero.¹¹⁷ This notebook is still worth writing. The promise is a sentence a colleague could follow. The promise is not a moved test score.

Checklist before moving on.

- Writes one dated sentence: I claim X because Y. We checked by Z.
- Names the stimulus and the sources, so a colleague who missed the meeting could find the file.
- Hears a fluent paragraph with no named source as a performance, not yet a check.
- Writes a fair other side so the other person would nod, then still looks.
- Changes the sentence when the check goes the other way, without treating that as losing.

students; authors warn of “metacognitive laziness.” One sentence, labelled not a workplace trial.

¹¹⁶Hamsa Bastani and colleagues, *Proceedings of the National Academy of Sciences*, 2025. Turkish high-school math; unguarded chatbot made practice look better and left students worse unaided. One crutch sentence, labelled not a workplace trial.

¹¹⁷Philip C. Abrami, Robert M. Bernard, Eugene Borokhovski, David I. Waddington, C. Anne Wade, and Tonje Persson, “Strategies for Teaching Students to Think Critically: A Meta-Analysis,” *Review of Educational Research* 85, no. 2 (2015): 275–314. Overall $g^+ = 0.30$, mixed ages, standardised thinking measures. Graduate and adult $k = 17$, $g = 0.21$; 95 percent confidence interval about -0.05 to 0.46 , not significantly greater than zero. Undergraduate $k = 126$, $g = 0.26$ is the wrong band to quote as workplace. Combination of authentic instruction, dialogue, and mentoring $g = 0.57$ is 19 mixed-age studies, not a job result. High school 16–18, $k = 71$, $g = 0.25$, contrast only. Labelled, not a promise that a paperback will move a standardised test. Abrami and colleagues, 2008 *RER*, remains the sourced sentence that improvement cannot be a matter of implicit expectation; do not pick 2008's larger adult cell and hide 2015.

- Catches an “it” or a “they” that needs the meeting, and rewrites.
- Does not treat extra writing, a five-paragraph night, or a college inventory as the gym.
- After a handoff, can write a line the next shift could follow.
- After any tool, can write a closed-window sentence from the file itself.

Desk, done enough. One dated sentence after a meeting or a call, sources named, other side optional but attempted.

Floor or shift, done enough. One dated line after a handoff, so the next shift could follow it.

What can still be thin: unaided far transfer to a domain you do not know; a culture that only rewards volume; a freeze at the blank page that is a writing problem, not a thinking one. Those can wait, or they can go to a human who already teaches the writing part. They are not a window that closed. A short written reason is not a bicycle. Knowledge of the spec, the roster, the metric, still has to be in the room.¹¹⁸ Adulthood is not a thinking birthday. A job is not arrival.

An adult who felt rusty should be able to say, after this chapter: I understand the shape. I know what to say tomorrow at the desk and on the floor. I know what a finished-looking paragraph without a named source looks like. I know what a why-story in ink sounds like. I can write a reason a colleague who missed the meeting could follow. I can do this.

¹¹⁸Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, NSW occasional paper, 2019. Thinking is not a bicycle; maxims without knowledge will not implement. Near practice on a recap does not, by itself, install unaided judgment on an unseen later artifact.

Chapter 7

On the job



Figure 15: A generic bay with a handoff notebook, a dashboard described as tiles, a vendor-call pad; no faces, no logos, no readable UI

Why this matters

What is the one thing that changed?

What do you notice?

Those two questions unlock the furniture of ordinary work. Desk: a meeting, a vendor call, a dashboard, a recap, a spec. Floor or shift: a handoff, a shop floor, a clinic bay, a warehouse count, a retail counter. No dinner table as the spine. No kitchen as the lab. No Bronze-to-Gold award scheme. The work hour already licensed a desk try-it shape and a floor try-it shape. This chapter fills named types and shows they are the same moves in two jackets.

A job is many small domains. Receiving is not cycle count. Cycle count is not a vendor dispute. Thinking that is fluent in one pocket does not ride, unaided, into the next.¹¹⁹ Transfer *within* a job, across artifacts, is the practical question. The sources this book opened almost never study it. Cue the similarity. “This is like the recap and the notes.” “This is like the deck chart.” Near is the honest hope. Far is not a promise.

The short written reason from Chapter 6 can own a page. This chapter is where that page lives. Meeting: one good question, wait, steelman, then the sentence. Vendor call: leave the glowing verb; open the spec. Dashboard: what does this chart actually show? Handoff: two notes; where they disagree; a third look. Shop floor, clinic bay, warehouse count: one thing that changed, or one thing you notice, and how you know.

That is worth the struggle because the week already has these rooms. Each of them can be a lecture, a gotcha, or a named move. The adult who can hear a hope offered as a look — “it will rise because we wanted it to,” “we always start with X,” “the board said twelve” — can practice this week. The adult who can name the one thing that changed, or name one thing they notice and how they know, then write a sentence a colleague who missed the meeting could follow, has practiced deciding what to believe or do on a domain they actually know.

¹¹⁹Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, “How to Teach Critical Thinking,” NSW Department of Education, 2019. Thinking is not a bicycle; processes are intertwined with content; experts encapsulate. Nathan Kuncel, as reported in Huber and Kuncel, *Review of Educational Research* 86, no. 2 (2016): field-specific expertise does not readily transfer; a finite set of specific reasoning skills as measured by thinking tests is not a bicycle. Huber and Kuncel is college, not the job; the moral still disciplines a workplace book. Access date for sources in these notes: 2 September 2026.

There is no named adult or workplace critical-thinking intervention report in the clearinghouse this book re-opened on 2 September 2026.¹²⁰ Employer interviews and a conceptual framework exist. They are not a trial of this Tuesday’s huddle.¹²¹ Ten to twenty-five minutes of a named move, a few times a week, on actual job artefacts, is the right ambition — not under-implementation of a proven workplace programme, because there is none.

This week you can try one meeting plus one written sentence at the desk, and one handoff or one “what do you notice?” on a bay you already work. Today the same five moves, plus two-source, graph, steelman, and the written reason, in two jackets.

Understand it yourself

You do not need a new job, a lab, or a philosophy kit. You do need to hear two things changed when one was asked for, a hope offered as a result, and a conversation sliding toward a gotcha.

Everyday picture. Two handoff notes on the same bay. Outgoing says 12. Incoming counts 9. “What is the one thing that changed?” If the board was not updated *and* a unit left *and* someone counted twice, you do not have a look. You have a story about a shift. Name the one thing you will check. Then look. Then write so the next shift could follow it.

That pair is an illustration, not a reported workplace. Your bay will be different. The move will not.

Precise picture. Desk and floor are the same moves on different furniture. Wait. What do you mean. How do you know. What else could be true. How could we check. What’s a second place we could look. What does this chart actually show. What’s the strongest version of theirs. Can you write that so a colleague who missed the meeting could follow it. The place does not install the move. The

¹²⁰Institute of Education Sciences, What Works Clearinghouse product search, re-opened 2 September 2026: no named adult or workplace critical-thinking intervention report. Academic Concept Mapping (Nelson 2007; WWC Study 89713) is college-advising GPA, not a thinking inventory, not workplace. Project QUEST mentions critical thinking in weekly meetings; rated outcomes are education and labour-market.

¹²¹Yi Song, Katrina Crotts Roohr, and Dessi Kirova, 2024, *Thinking Skills and Creativity*: conceptual workplace-CT framework plus employer interviews, not a trial.

named question does.

A job title does not install the move either. Field-specific expertise does not readily transfer onto a generic test, and a generic test does not readily transfer onto the next pocket of the same job.¹²² The clinic nurse who can think with a protocol can still miss a truncated bar on a staffing slide. The analyst who can catch an axis can still mash a why-story into a handoff. Cue the similarity. Stay with the object in front of you.

Safety is ordinary. No public quiz in front of a supervisor. No undercover audit of a workplace. No gotcha against a colleague, a vendor, or a shop. No lecture to a stranger's customer. No sideline interrogation during a live procedure. A clinic try-it uses a protocol already in use — not a diagnosis of a named patient in this book. Noticing is not a sting.

Professional domains already name a skepticism or evaluation duty *inside* the work: audit, nursing, journalism, engineering. Those duties live in records and resources as domains. They are not a transferable faculty, and they are not this chapter's curriculum.

Dinner-table practice from a sibling book is a neighbour, not the spine. If you think at home, that is your life. This chapter's objects are the meeting, the handoff, the vendor call, the dashboard, the shop floor, the clinic bay.

Wrong answers you should be able to hear

Including from yourself.

1. Two things changed, or a hope offered as a look. “We used the other bin *and* a new count.” “It will be fine because we wanted it to.” “We always start with X.”

What it usually means. The look collapsed into a story. A why-story where a check was asked for.

¹²²Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, “How to Teach Critical Thinking,” NSW Department of Education, 2019. Thinking is not a bicycle; processes are intertwined with content; experts encapsulate. Nathan Kuncel, as reported in Huber and Kuncel, *Review of Educational Research* 86, no. 2 (2016): field-specific expertise does not readily transfer; a finite set of specific reasoning skills as measured by thinking tests is not a bicycle. Huber and Kuncel is college, not the job; the moral still disciplines a workplace book. Access date for sources in these notes: 2 September 2026.

What to say. “What is the one thing that changed? What did we keep the same? What did you see? What’s a second place we could look?” Then look. Then write.

2. Noticing heard as snitching. “If I say the board is wrong, I’m reporting someone.”

What it usually means. The move was heard as a personality, or as an audit.

What to say. “What do you notice? How do you know?” Point at the board and the count, not at a person. Ordinary looking. After, not during. Private if the power difference is steep.

3. The job conversation as a gotcha, or the huddle as a courtroom. A public quiz in front of a vice-president. A vendor call turned into a trap. A clinic question during a live procedure.

What it usually means. Ordinary looking was inflated into an audit or a fight.

What to say. After, not during. Private. Two notes, or the spec, or the clock. Then stop.

4. The dinner table as the week’s work, or a franchise in place of the bay. A kitchen hour offered as if it were the job. An overnight identity. A workshop pin. “Critical Thinking I” on a record.

What it usually means. The place was asked to do the work of the move, or a sibling book’s furniture was imported as the spine.

What to say. One question. One look. One dated sentence. Title still the job you hold.

5. A fluent recap of the shift offered as the check. “We had a good handoff.” “The dashboard looked good.” “The call went well.”

What it usually means. A caption wearing a meeting. The notes, the tile’s number, the spec, are still sitting there.

What to say. “What is the one thing that changed? What do you notice? How do you know? Can you write that so a colleague who missed it could follow it?”

Five-minute warm-up

Before the sitting, with no one else in the room:

1. Pick today's place: a meeting already on the calendar, a vendor call already booked, a dashboard you already open, a handoff you already do, a bay you already work. Write one opening: "What is the one thing that changed?" or "What do you notice?"
2. Wait a slow three out loud.¹²³ Write the check you will point at: the notes, the spec, the axis, the other note, the clock, the count.
3. Write the sentence you will ask for at the end: claim, because, how we checked. Then stop talking sooner than feels polite, in rehearsal.

If those three are fluent, you can hear two changes, a hope, and a gotcha. That is the job. You still look.

How to try it this week

Use the work-hour shape. Fill it with one place. Stop while it is still a sitting. You do not need every try-it this week. You need one at the desk and one on the floor or shift.

Warm-up (2–4 minutes). Yesterday's written reason, from memory. "What did we claim? How did we check?" Then: "Today the place is the job. Same moves."

Short model (5–8 minutes).

¹²³Robert J. Stahl, "Using Think-Time and Wait-Time Skillfully in the Classroom," ERIC Digest ED370885 (1994). About three seconds as a convention, classroom origin. Not a workplace trial.



Figure 16: Two generic handoff notes with greeking, a pad, a board described as a tile; no faces, no logos, no readable UI

You talk first. Two notes. You say a sloppy spoken version on purpose: “The handoff was fine.” Then you ask: “What is the one thing that changed?” Wait. You lift outgoing: 12. You lift incoming: 9. You wait. You ask how we know. You look at the board. You write:

2 September 2026. Stimulus: outgoing note, incoming count, bay board at 14:10. I claim the count is 9 because I counted. We checked against the outgoing note, which said 12, and the board, which still said 12. Someone who disagrees would say the board is the record.

You wait. You ask whether the next shift could follow that. Then you show the *wrong* cousin on purpose: a fluent “good handoff” with no number. “Is that a check?” Wait. Then stop talking sooner than feels polite.

This pair is an illustration, not a reported workplace. Use *your* notes if you have

them.

Exact wording you can steal.

On the place: “What is the one thing that changed?” “What do you notice?”

On meaning: “What do you mean, *fine*? What do you mean, *always*?”

On reason: “How do you know? Did you see it, or is this a story of why it would make sense?”

On alternative: “What else could be true? Always — can we think of a time it wasn’t?”

On check: “What’s a second place we could look?”

On the sentence: “Can you write that so a colleague who missed the meeting could follow it?”

On mind-change: “You thought it was 12. We counted 9. You can say the new thing. That is not losing.”

Your attempt. After the model, you run the next place unaided — today’s meeting, today’s call, today’s bay. Product: one named question, one look, one dated line. You wait. You do not grab the answer. You do not finish the reason.

One good question. After the attempt: “If we had changed two things, could we still know which one mattered?” You generate the why.

Mixed practice. One meeting. One handoff. One dashboard. Mix, so the move is not glued to the first bay.

Written line (2–5 minutes, unaided). Claim, because, check. Tool closed.

How to fade help. You name the place and the question → you look while a sticky note names the check → you name both → you catch a hope offered as a look → you write without the prompt. When you can say “this is like the recap and the notes” on a new artifact, put the sticky-note scaffold away. Bring it back when a new pocket of the job arrives.

When to stop talking. After one model and one question. You hold the key: is there a named object, a wait, a look, a sentence? You do not become a lecture about the culture of work. In a meeting, one good question beats five. After, not

during, a live procedure. Wait in private when the power difference is steep.^{124 125}

Desk / floor-or-shift

Desk. A meeting you were already in, plus five minutes after. A vendor call's notes plus one named source. A dashboard you already own. The sentence lives in the project folder.

Floor or shift. A handoff after, not during. A bay you already work. Two quotes for the same SKU after a walk-through. Posted hours versus the clock. A clinic board versus a protocol already in use. Ordinary noticing. No sting.

Placement is by skill, not job title. A lead who still offers "good handoff" as the check is still in the early band of this chapter. A newer person who names the one thing that changed is already practicing the later catch.

Practice that actually builds it

Named try-its. Time, materials, safety, the move. At least one lives at the desk. At least one lives on the floor or shift.

Try-it 1 — Meeting, desk

Time. 10–15 minutes of a meeting you were already in, plus 5 minutes after.

Materials. The notes, or the recap already in the inbox. Dated notebook.

Safety. Curious, not courtroom. No public quiz of a junior in front of a vice-president. Wait in private when the power difference is steep.

The move. One good question, wait, steelman, then the sentence.

How to run it. In the meeting, ask one: "What do you mean, *always*?" or "How do you know we decided that?" Wait. After: write claim, because, check, so a colleague who missed it could follow. Someone who disagrees would say...

¹²⁴Robert J. Stahl, "Using Think-Time and Wait-Time Skillfully in the Classroom," ERIC Digest ED370885 (1994). About three seconds as a convention, classroom origin. Not a workplace trial.

¹²⁵Mary Budd Rowe, wait-time papers, elementary science class, 1972/1974. Not a stand-up randomised trial. Do not mash Rowe and Stahl as if they were one paper.

Desk. Ink after. Title still the project.

Floor or shift cousin. A huddle in the bay. Same one question. Write after, in a break.

Try-it 2 — Vendor call, desk

Time. 10–15 minutes after.

Materials. The claim on the call and one named source: the spec, the quote, the contract clause.

Safety. After, not during. No gotcha on a live call. No confidential paste into a public tool.

The move. Leave the glowing verb; open the spec or the clause.

How to run it. Write the verb you heard — *leading, faster, reduced*. Opening question after: “What is the one thing that changed?” Then: “What’s a second place we could look?” Open the named source. Write the sentence.

Desk. Spec next to notes.

Floor or shift cousin. After a walk-through, two quotes for the same SKU, or a product page versus a data sheet. What each actually includes.

Try-it 3 — Dashboard, desk

Time. 5–10 minutes.

Materials. Tiles, a bar pair, a sparkline — described. Point back to Chapter 5.

Safety. You already have access. No undercover audit.

The move. What does this chart actually show? What is this a number *of*? Where does the axis start?

How to run it. Cover the caption. Lift one number. Compare if there is a pair. Write the line. Cue: “this is like the deck chart.”

Desk. Your laptop, after a huddle.

Floor or shift cousin. A shared screen in the bay. One tile. Spoken. Ink later.

Try-it 4 — Handoff, floor

Time. 5–10 minutes after.

Materials. Two notes on the same bay, line, or board.

Safety. After, not during, a live procedure. No public quiz. No named patient. Noticing is not a sting.

The move. Two notes; where they disagree; a third look.

How to run it. Read outgoing. Read incoming. Opening question: “What is the one thing that changed?” or “What do you notice?” Wait. Look at the board, the count, or the protocol already in use. Write so the next shift could follow it.

Desk cousin. Two meeting notes on the same hour.

Floor or shift. Clinic, shop, warehouse, retail — the place you already work.

Try-it 5 — Shop floor, clinic bay, or warehouse count

Time. 3–10 minutes.

Opening cousins. “What is the one thing that changed?” / “What do you notice?”

Materials. A part, a box, a label already in the bay; a clinic board and a protocol already in use; a count you were already doing.

Safety. Ordinary noticing, not a sting. No lecture to a stranger’s customer. No tasting unknown things. No diagnosis of a named patient in this book. No sideline interrogation during a live procedure.

The move. Name one thing you notice and how you know. Then a second place if one is cheap and safe.

How to run it. Stop at the object. Ask. Wait. Look. Write later. Cue if needed: “this is like the two notes.”

Desk cousin. An object you can hold, photographed or described in the file, plus the spec.

Floor or shift. Spoken in the bay. Ink after.

Try-it 6 — Vendor call on the floor

Time. 5–10 minutes after a walk-through.

Materials. Two quotes for the same SKU, or a product page versus a data sheet.

Safety. After the walk-through. No gotcha against a shop. No lecture in the aisle.

The move. What each quote actually includes. Listed capacity, weight, warranty words on the sheet — not stars.

How to run it. Leave the first page. Open the second. Write the difference so a colleague who missed the walk-through could follow it.

Desk cousin. Two vendor one-pagers in the same category, from Chapter 4.

Floor or shift. Pad in a break. Title still the job.

Talk/meeting box

Exact opening question. “What is the one thing that changed?” / “What do you notice?”

Not: “Did we have a good meeting?” Not: “What is the vocabulary word?” Not: “Who is at fault?”

Follow-ups.

1. What do you mean by that verb? (meaning)
2. How do you know? Did you see it, or is this a story of why it would make sense? (reason)
3. What else could be true? Always — can we think of a time it wasn’t? (alternative)
4. What’s a second place we could look? (check / second source)
5. Can you write that so a colleague who missed the meeting could follow it?
You may change your mind. (write / revise)

How to wait. After you ask, a slow three.¹²⁶ After they stop, wait again. If they are mid-reason, do not cut them off. Look at the notes, the graph, the roster, or the clock, not at the person, if the silence is hard. After “write the sentence,” that is task-completion time. At work the shrug often looks like refusal. Do not rescue it.

¹²⁶Robert J. Stahl, “Using Think-Time and Wait-Time Skillfully in the Classroom,” ERIC Digest ED370885 (1994). About three seconds as a convention, classroom origin. Not a workplace trial.

Wait in private when the power difference is steep. You are not citing science-class wait-time as a stand-up trial.¹²⁷

What a stuck silence usually means. They think noticing is snitching. They think the dinner-table chapter from a sibling book is the spine. They offered a why-story about the shift. Wait-time 1 was zero. They are guessing what you want. The question was a recitation item in disguise. They think “critical” means finding fault. They think a fluent recap of the hour is already a check.

If they are stuck. Smaller question. “What number is on the board?” Wait. Point at the object. Do not fill the silence with the right answer. Do not turn the huddle into a fight.

Blocked, then mixed

Blocked. First sitting: one place you already work, one opening, one look. Stay there until the question has a name.

Mixed. Then a meeting. Then a call. Then a dashboard. Then a handoff. Mix so the move is not glued to one bay. Cue the similarity when the pocket changes: receiving is not cycle count; the recap is not the notes; the tile is not the spreadsheet. Near transfer across artifacts in the same job is the honest hope. Far transfer to an unseen later domain is not a promise.¹²⁸

One incorrect example to diagnose

A fluent recap of the shift, posted in the thread. “Handoff was clean. Board looks good. We always start with the left bay.” No count. No second note. No time. The paragraph looks like work.

This is an illustration, not a reported workplace.

¹²⁷Mary Budd Rowe, wait-time papers, elementary science class, 1972/1974. Not a stand-up randomised trial. Do not mash Rowe and Stahl as if they were one paper.

¹²⁸Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, “How to Teach Critical Thinking,” NSW Department of Education, 2019. Thinking is not a bicycle; processes are intertwined with content; experts encapsulate. Nathan Kuncel, as reported in Huber and Kuncel, *Review of Educational Research* 86, no. 2 (2016): field-specific expertise does not readily transfer; a finite set of specific reasoning skills as measured by thinking tests is not a bicycle. Huber and Kuncel is college, not the job; the moral still disciplines a workplace book. Access date for sources in these notes: 2 September 2026.

What it usually means. A caption wearing a shift. The mash of a why that makes sense — we wanted a clean handoff — with evidence that it was so.

What to say. “What is the one thing that changed? What do you notice? How do you know?” Wait. Open the two notes. Then: “Can you write that so the next shift could follow it?” If the sentence cannot survive the check, correct the sentence. That is the work. Silence in the face of a false recap is not kindness.

If it isn’t clicking

The place keeps sliding into a gotcha, or into dinner. Slow down. After, not during. Private. Pick a cooler object: the recap, the spec, the two notes, the tile. Dinner-table practice, if you do it, is extra. It is not this chapter. If a real fight is too hot, postpone. Practice the move on the vendor page, the deck number, or the roster.

Noticing is heard as snitching, or “critical” is heard as finding fault. Hear it. “What do you notice?” is about the board, the count, the notes — not about a person. Point at the object. Correct the *reason* without crushing the talk. “I hear that the board said 12. The strongest version of that is the board is a record. I’m not sure that reason works as the count, because we counted 9. What’s a second place we could look?” Then look. Then write. “Never correct a colleague” is not this book’s rule.

The sitting becomes a tour of the job instead of a named move. Hear it. Return to one question. One look. One dated line. A job is many pockets. You will not finish them this week. Cue the similarity the next time the genre recurs. If the wall is knowledge — you do not yet know what the spec lists, what the protocol requires — that is the next move, not a character flaw. Ask a human colleague who already owns that pocket.

Go ahead when you can run one desk sitting and one floor sitting, write the sentence, and change it when the look goes the other way. Stay longer when “good meeting” is still the check. A job title does not settle it.

Tools, including AI

Point back to the work-hour box. Keep this short.

After you have tried, you may ask a tool to explain this page to yourself from a named source, or to give extra questions you then vet. Throw out any that import a live controversy, a secret, a pretend expert, a colleague's private performance, or a civic-web item. A labelled SCRIPT after the try is for you to say — not a paragraph to copy as your reason, your email, or your steelman. A hint after an attempt is not the finished reason. A diagnosis of a reason you already gave is yours to hear and decide.

You may not park an unsupervised chatbot as the only partner. You may not ask a tool to invent what the roster says, what the bay counted, or what the spec lists. You may not ask it to write your reason and paste that text as yours. You may not invent an expert identity. You may not grade a colleague from a detector score. You may not print a certificate of thinking. You may not hand yourself an agent as a thinking partner for the first attempt. You may not put confidential workplace data into a public consumer tool.

On some consulting tasks, adults using a model were faster and higher-rated inside the frontier and less likely to be right outside it.¹²⁹ A fluent recap of the shift is not a check that the count was 12. After any tool, write the sentence from the notes, the board, or the spec, tool closed. One crutch sentence, labelled: an unguarded chatbot can make practice look better and leave you worse when the window is closed.¹³⁰

Open *your* employer's acceptable-use page. Label it. A thinking tool is not a secret friend.

What “done enough” looks like

Done enough is skill, not a job title, and not a course called Workplace Thinking 1. The title a stranger can read is still the job you hold. The dated note lives inside that file. It is not Critical Thinking I. A workshop certificate is a receipt, not the work.

Checklist before moving on.

¹²⁹Fabrizio Dell'Acqua and colleagues, Harvard Business School working paper, 2023. About 758 consultants; inside-frontier faster and higher-rated; outside-frontier 19 percentage points less likely to be correct. In-band, labelled. Full PDF not line-checked this pass.

¹³⁰Hamsa Bastani and colleagues, *Proceedings of the National Academy of Sciences*, 2025. Turkish high-school math. One crutch sentence, labelled not a workplace trial.

- Can run one named question in a meeting, wait, and write the sentence after.
- Can leave a glowing verb on a vendor call and open a named source.
- Can ask what a dashboard tile actually shows, pointing back to Chapter 5.
- Can read two handoff notes, name where they disagree, and look a third time.
- Can ask “what is the one thing that changed?” or “what do you notice?” on a bay they already work, without turning it into a sting.
- Writes a dated line a colleague who missed it could follow.
- Changes the sentence when the look goes the other way, without treating that as losing.
- Cues similarity when the pocket changes: “this is like the recap and the notes.”
- Does not treat dinner, a pin, or a job title as arrival.

Desk, done enough. One meeting plus one written sentence this week.

Floor or shift, done enough. One handoff or one “what do you notice?” on a bay you already work, then a later dated line.

What can still be thin: unaided transfer from receiving to a vendor dispute with no cue; a clinic pocket you have not yet learned; a dashboard metric whose definition you do not own. Those can wait. They are not a window that closed. A job is not a bicycle. Knowledge of *this* pocket still has to be in the room.¹³¹ Adulthood is not a thinking birthday.

An adult who felt rusty should be able to say, after this chapter: I understand that desk and floor are the same moves. I know what to say tomorrow in a meeting and after a handoff. I know what a hope offered as a look sounds like. I know what a fluent recap of the shift looks like when it is not yet a check. I can do this at the desk and on the floor.

¹³¹Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, “How to Teach Critical Thinking,” NSW Department of Education, 2019. Thinking is not a bicycle; processes are intertwined with content; experts encapsulate. Nathan Kuncel, as reported in Huber and Kuncel, *Review of Educational Research* 86, no. 2 (2016): field-specific expertise does not readily transfer; a finite set of specific reasoning skills as measured by thinking tests is not a bicycle. Huber and Kuncel is college, not the job; the moral still disciplines a workplace book. Access date for sources in these notes: 2 September 2026.

Chapter 8

Approaching disagreement at work



Figure 17: Two empty pads on a table, a spec sheet next to a vendor page with greeking, a roster next to a forwarded email with greeking; no faces, no logos

Why this matters

Say yours, then say theirs.

That sentence unlocks disagreement at work. Not a debate club. Not a current-events hour. Not a lesson on who is allowed to be upset. A sequence, practiced on a vendor page, a recap versus the notes, two quotes for the same SKU, a review versus a spec, a teammate’s “we always,” a truncated chart, a shift rumor about who works when. You listen. You ask what they mean. You ask for a reason. You open a second source. You wait. You may change your mind. You say the strongest fair version of theirs. A fluent sentence is not a check. Then you write the check so a colleague who missed the meeting could follow it.

You do not have to win the thread. You have to listen, then ask, then check.

Dual focus — your side and theirs — is promptable. It is not the default. Own-side talk is the default, including among adults who are already more strategic than teenagers.¹³² Copy the *shape* of classroom argumentation: a counter, evidence, a fair version of the other side. Copy not the topics those classrooms sometimes used. This chapter stays with ordinary workplace stakes.

If a vendor page says “industry-leading” and the thread repeats “industry-leading,” repeating “industry-leading” as the thing to knock down practices the rumor. Repeating the check — open the spec sheet, open last quarter’s incident log — practices the look. Deck translation: repeating “jumped” is the rumor; repeating the source spreadsheet and the axis start is the check. Shift-rumor translation: repeating “nights” is the rumor; repeating the posted roster is the check. That is a translation of a handbook about correction, using a vendor claim, a deck claim, or a roster claim. It is not a workplace trial of the handbook.¹³³

¹³²Matthew Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, no. 2–3 (2001): 135–153. Community-college young adults versus seventh- and eighth-graders; adults more strategic, still own-side. Lab topic of that study is not a book exercise; copy shape. Keith E. Stanovich, Richard F. West, and Maggie E. Toplak, “Myside Bias, Rational Thinking, and Intelligence,” *Current Directions in Psychological Science* 22, no. 4 (2013): 259–264. Myside bias shows little relation to intelligence, including in university samples, often literally zero in naturalistic between-subjects designs. Access date for sources in these notes: 2 September 2026.

¹³³Stephan Lewandowsky, John Cook, Ullrich Ecker, and 19 co-authors, *The Debunking Handbook 2020*. Backfire effects occur only occasionally; do not refrain from correcting out of fear of backfire. Continued-influence effect is the sticky problem. Detailed refutation outperforms a bare “that’s false.” Illusory truth: repetition increases judged truth. Work translation in this chapter —

A later, heavier claim will need knowledge you do not have yet. That is a fact about knowledge, not a dodge.¹³⁴ If a real workplace disagreement is too hot for the table today, pick the vendor page, the deck number, the shift rumor, or the invoice. Practice the moves.

High-school students who stay on a finished-looking page often never learn who is behind it — that is why a working adult leaves a vendor page, a slide, and a shift-rumor email *now*. That is one civic-why sentence. The items from that sitting are not assigned here.¹³⁵ Adult experts in another study often stayed on the page too. Leaving is the habit. It is not automatic because you were hired.¹³⁶

That is worth the struggle because the job already disagrees. Switch or stay. Recap versus notes. Two quotes. Spec versus review. “Always” versus last quarter’s process document. A chart that looks huge. A forwarded email about who works when. Each of those can be a fight, or a sequence. The adult who can hear “because we wanted it more” as a starting point, not a check, can practice this week. The adult who can let the other person finish, say their claim so they would nod, look at one other thing, and write the check has practiced deciding what to believe or do, with reasons, on a domain they actually know.

This week you can learn the sequence on a vendor or a deck at the desk, and on a shift rumor or an “always” on the floor, and that a fluent “because we always do” is still a starting point, not a check. Today you can say yours, then say theirs, look at one named source, and write the check.

Understand it yourself

You do not need to be a debate coach. You do need to hear “because we wanted it more” as a wish, “because we always” as a fluent start, “the vendor said leading”

vendor page, deck chart, shift rumor about who works when — is a translation, not a workplace randomised trial of the handbook.

¹³⁴Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, NSW occasional paper, 2019. Sequence and revisit; maxims without knowledge will not implement; a later heavier claim needs the later domain’s knowledge. Near is not far.

¹³⁵Joel Breakstone and colleagues, 2021, high-school civic online reasoning sitting (3,446 students). One civic-why sentence in the body; items not assigned; not this book’s population.

¹³⁶Sam Wineburg and Sarah McGrew, 2019. Professional fact checkers left an unfamiliar page; Ph.D. historians often stayed. Adult experts in-band as why leave this page. Pointed back to Chapter 4. Civic items in that paper are research objects, not exercises here.

as a sentence that still has a spec, and “the email said nights” as a sentence that still has a roster.

Everyday picture. Two people, one SKU, two quotes. Each has a sentence. If you pick a winner in the first second, you have a referee hour. If you say “say yours, then say theirs,” wait, ask what *includes* means, ask how they know, and look at what each quote actually lists, you have a work hour. Five to ten minutes when it actually happens, or as a calm replay. If the stakes are too hot, postpone. Practice on a cooler claim: the recap, the deck chart, the posted roster.




Figure 18: Two empty chairs facing each other across a table, two pads, no people

Precise picture. A *claim* is a sentence that can be true or not, or fair or not, about something we can look at. You do not have to win. You have to **listen**, then **ask what they mean**, then **ask for a reason**, then **leave this page for a second source**, then **wait**, then **allow a change of mind**. A fluent sentence is not a check. A *steelman* is the strongest fair version of what they mean, said so they would nod.

It is not agreement. After you say it, you still look. Inventing a better argument they did not make, then knocking it down, is not steelman. Asking a machine to write the steelman so you can paste it is write-my-reason. A graph can mislead by scale; “looks huge” is a caption, not a number. Then **write the check**: I claim X because Y. We checked by Z. Someone who disagrees would say... Changing your mind when the check goes the other way is the work. It is not losing. “I’m not sure yet” is allowed. Fake certainty is a fluent sentence pretending to be a check.

A short sequence you can keep, in ordinary language: restate theirs so they would nod; name any points of agreement; name what you still need to look at; then look.¹³⁷ You will not use all of it as a script in the first second. You will have it.

Reasons can be better or worse. Sensitivity to context — last quarter’s tickets are not the same object as a feeling about the demo; a spec sheet is not a reviewer’s word *leading*; a roster is not a forwarded email — is not the same as “anything goes.” You can train the talk moves and still leave someone believing that reasons do not rank. Take that seriously. Hearing a wrong reason, then asking a better question, then naming a better reason if needed, is the work. Silence in the face of a false reason is not kindness. “Never correct a colleague” is not this book’s rule.

Intelligence does not cancel own-side preference. University samples still show it, often with no relation at all.¹³⁸ A job is not a thinking birthday. Continued influence is the sticky problem: you opened the spec sheet, the vendor overclaimed, and ten minutes later the thread still says “industry-leading.” Repeat the check, not the rumor.¹³⁹ Backfire — making the rumor stronger by correcting it — is oc-

¹³⁷Steelman shape in ordinary language, after secondary reports of a four-step restatement practice associated with Rapoport and with Dennett. Page-level wording from those books was not independently opened this pass; the ordinary-language sequence in the body is the portable piece.

¹³⁸Matthew Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, no. 2–3 (2001): 135–153. Community-college young adults versus seventh- and eighth-graders; adults more strategic, still own-side. Lab topic of that study is not a book exercise; copy shape. Keith E. Stanovich, Richard F. West, and Maggie E. Toplak, “Myside Bias, Rational Thinking, and Intelligence,” *Current Directions in Psychological Science* 22, no. 4 (2013): 259–264. Myside bias shows little relation to intelligence, including in university samples, often literally zero in naturalistic between-subjects designs. Access date for sources in these notes: 2 September 2026.

¹³⁹Stephan Lewandowsky, John Cook, Ullrich Ecker, and 19 co-authors, *The Debunking Handbook 2020*. Backfire effects occur only occasionally; do not refrain from correcting out of fear of backfire. Continued-influence effect is the sticky problem. Detailed refutation outperforms a bare “that’s false.” Illusory truth: repetition increases judged truth. Work translation in this chapter — vendor page, deck chart, shift rumor about who works when — is a translation, not a workplace

casional, not a reason to skip the look. A detailed look outperforms a bare “that’s false.”

Four lines on a pad are a cousin: we listen; we ask what you mean; we give a reason; we can change our minds; we write a reason a colleague who missed the meeting could follow. Posting is not practice. Practice is the next real disagreement.

Wrong answers you should be able to hear

Including from yourself.

1. Listening means awarding the vendor, or giving up the starting spot. “If I let them talk, we switch.” Or: “If I steelman the email, nights is the schedule.”

What it usually means. They think the sequence is a trick to make them lose.

What to say. “Say yours, then say theirs. Then we look at the spec, or the notes, or the roster. Changing your mind if the check goes the other way is not losing. We still decide what to do.”

2. A fluent sentence with no check. “Because we always do.” “Because the vendor said leading.” “Because the graph looks huge.” “Because the email said nights.” “Because it sounds right.” “Because the deck is done.” “Because the model wrote it.”

What it usually means. A starting point offered as the look. The mash of a why that makes sense with evidence that it is so. Own-side “because we wanted it more” wearing work clothes.

What to say. “What do you mean, *leading*? What do you mean, *always*? What do you mean, *jumped*? What do you mean, *nights*? How do you know? What’s a second place we could look?” Then look. Then: “Can you write the check so a colleague who missed the meeting could follow it?”

3. You supplied the meaning and argued with that. They said “faster.” You heard “always faster,” and now you are in a fight with a sentence they did not say.

What it usually means. Meaning was skipped. Argument ate the sitting.

randomised trial of the handbook.

What to say. “What do you mean, *faster*? What do you mean, *we decided*?” Wait. Then the reason. Then the check.

4. Repeat the rumor to knock it down. “Leading, leading, leading — see, they’re not.” Or a singsong of “they said nights” after you already looked at the roster. Or repeating “jumped” as the thing to kill, instead of opening the spreadsheet and the axis.

What it usually means. The sticky sentence got practiced. The check did not.¹⁴⁰

What to say. Repeat the check: the spec sheet, the notes, the axis from zero, the other quote, the posted roster, last quarter’s process document. Lead with what you looked at.

5. A performance of a fight, or a performance of being “the nice one,” or a live controversy imported as the object. Volume as proof. Instant surrender with no look, to please the room. A steelman that is actually a straw. A finished-looking paragraph offered as the verdict. A thread that leaves the vendor page for a fight this book does not rehearse.

What it usually means. The sitting became theatre. The object left the table. Steelman was heard as a personality. Length was heard as a look. The week’s practice was asked to carry a claim it does not have knowledge for.

What to say. Stop. Point at the vendor page and the spec, the recap and the notes, the two quotes, the chart and the axis, the roster and the email. Smaller claim. “Say yours, then say theirs. What can we look at?” If it is still theatre, postpone. If the claim is too hot, pick the cooler object. Practice the moves.

Five-minute warm-up

Before the next real disagreement, or with a cooler claim on purpose:

1. Write the sequence on a scrap: listen; what do you mean; how do you know; two sources / leave this page; wait; you may change your mind; fluent ≠

¹⁴⁰Stephan Lewandowsky, John Cook, Ullrich Ecker, and 19 co-authors, *The Debunking Handbook 2020*. Backfire effects occur only occasionally; do not refrain from correcting out of fear of backfire. Continued-influence effect is the sticky problem. Detailed refutation outperforms a bare “that’s false.” Illusory truth: repetition increases judged truth. Work translation in this chapter — vendor page, deck chart, shift rumor about who works when — is a translation, not a workplace randomised trial of the handbook.

check; steelman; a graph can mislead by scale; then write the check. You will not use all of them as a script in the first second. You will have them.

2. Write today's opening: "Say yours, then say theirs." Wait a slow three out loud. It will feel long.¹⁴¹
3. Write one fluent fake you expect — "because we always," "the vendor said leading," "looks huge," "the email said nights," "the recap said we decided" — and the check you will point at: last quarter's process document, spec sheet, axis from zero, posted roster, notes, the other quote.

If those three are fluent, you can hear a wish, a mash, and a rumor being practiced. That is the job. You still talk. The object still sits there.

How to try it this week

Use the work-hour shape. Fill it with one ordinary stake: a vendor page, a recap, two quotes, a review versus a spec, an "always," a truncated chart, a shift rumor about who works when. Fifteen to twenty-five minutes on a named artifact. Stop while it is still a sitting. If the real fight is too hot, pick the cooler object.

Warm-up (2–4 minutes). Yesterday's on-the-job look, from memory. "What did we notice? How did we know?" Then: "Today we say ours, then theirs."

Short model (5–8 minutes). Two pads. You talk first. A vendor page that says *leading*, and a spec sheet. You say a sloppy spoken version on purpose: "They're leading, so we switch." Then: "Say yours, then say theirs." You wait. You steelman: "You mean we should switch because last month's tickets were slow." You wait. You open clause 4.2. You write:

2 September 2026. Stimulus: vendor one-pager plus spec clause 4.2 and July ticket log. I claim we should not switch this month because the spec still lists 48-hour response and July tickets sat at 72 hours. We checked by opening clause 4.2 and the log. Someone who disagrees would say the demo still felt faster, and last month's tickets were slow.

You wait. You ask whether a colleague who missed the call could follow that. Then you show the *wrong* cousin on purpose: repeating "leading" as the thing to

¹⁴¹Robert J. Stahl, "Using Think-Time and Wait-Time Skillfully in the Classroom," ERIC Digest ED370885 (1994). About three seconds as a convention, classroom origin. Not a workplace trial.

kill, without opening the spec. “Is that a check?” Wait. Then stop talking sooner than feels polite.

This pair is an illustration, not a reported workplace. Use *your* page if you have one.

Exact wording you can steal.

On the sequence: “Say yours, then say theirs.”

On meaning: “What do you mean, *leading*? What do you mean, *always*? What do you mean, *jumped*? What do you mean, *nights*?”

On reason: “How do you know?”

On alternative: “Always — can we think of a time it wasn’t?”

On check: “What’s a second place we could look?”

On steelman: “What’s the strongest fair version of what they mean? Say it so they would nod.”

On the rumor: “We looked at the spec.” Not a singsong of the glowing verb.

On the sentence: “Can you write that so a colleague who missed the meeting could follow it?”

On mind-change: “You may change your mind. That is not losing.”

Your attempt. After the model, you run the next stake unaided. Product: yours, theirs so they would nod, one named look, one dated line. You wait. You do not grab the answer. You do not finish their steelman. You do not paste a machine’s steelman as yours.

One good question. After the attempt: “If we had only repeated the verb, would we have practiced the rumor or the check?” You generate the why.

Mixed practice. One vendor. One deck. One roster rumor. Mix, so the move is not glued to one object.

Written line (2–5 minutes, unaided). Claim, because, check, other side. Tool closed.

How to fade help. You name the whole sequence → you steelman while a sticky note names the check → you name both → you catch a rumor being practiced →

you write without the prompt. When you can say theirs so they would nod and still look, put the sticky-note scaffold away. Bring it back when a new genre arrives.

When to stop talking. After one model and one question. You hold the key: did we listen, ask meaning, ask a reason, look, wait? You do not become a lecture about disagreement as a personality. In a meeting, one good question beats five. Wait in private when the power difference is steep.¹⁴² ¹⁴³ If the sitting becomes a fight, stop. Postpone. Pick a smaller claim.

Desk / floor-or-shift

Desk. Vendor disagreement. Deck disagreement. Recap versus notes. Review versus spec. Steelman in the thread, then the sentence in the project folder.

Floor or shift. Shift rumor about who works when. Two quotes for the same SKU. “We always.” Posted roster versus a forwarded email. After, not during. Ordinary looking. No sting.

Placement is by skill, not job title. A senior person who still fortifies their side and skips theirs is still in the early band of this chapter. A newer person who can hold two sentences and a look is already practicing the later catch.

Practice that actually builds it

Named try-its. Time, materials, safety, the move. At least one lives at the desk. At least one lives on the floor or shift.

Try-it 1 — Vendor disagreement, desk

Time. 8–15 minutes.

Materials. The vendor page and one named source: the spec, the contract clause, last quarter’s ticket log.

¹⁴²Robert J. Stahl, “Using Think-Time and Wait-Time Skillfully in the Classroom,” ERIC Digest ED370885 (1994). About three seconds as a convention, classroom origin. Not a workplace trial.

¹⁴³Mary Budd Rowe, wait-time papers, elementary science class, 1972/1974. Not a stand-up randomised trial. Do not mash Rowe and Stahl as if they were one paper.

Safety. After a call, not as a trap on the call. No gotcha against a vendor. No confidential paste into a public tool.

The move. Opening cousin: “You mean we should switch because last month’s tickets were slow.” Then the contract clause.

How to run it. Say yours. Say theirs so they would nod. Opening question: “Say yours, then say theirs.” Wait. Open the named source. Repeat the check, not the glowing verb. Write the sentence.

Desk. Page plus clause.

Floor or shift cousin. After a walk-through, product page versus data sheet.

Try-it 2 — Deck disagreement, desk

Time. 5–10 minutes.

Materials. One slide, the source spreadsheet, scrap paper for a from-zero redraw. Point back to Chapter 5.

Safety. Curious, not courtroom. No humiliation of whoever built the slide.

The move. “You mean unit sales were up last month, not that they ‘jumped’ for the year.” Then the axis.

How to run it. Say yours. Say theirs so they would nod. Opening question: “Say yours, then say theirs.” Cover the title. Lift the number. Ask where the axis starts. Redraw from zero. Repeat the spreadsheet and the axis start, not a singsong of “jumped.” Write the check.

Desk. Slide plus tab.

Floor or shift cousin. A shared-screen bar pair in the bay. Spoken. Ink later.

Try-it 3 — Shift rumor, floor

Time. 5–10 minutes.

Materials. The forwarded email and the posted roster, plus a named manager’s written note if you already have one.

Safety. After, not during a live rush. Private if the power difference is steep. Ordinary looking. No undercover audit. No public quiz. A rumor about *who works when* — not a rumor about a person's character.

The move. “You mean nights because the forwarded email said nights.” Then the posted roster.

How to run it. Say yours. Say theirs so they would nod. Opening question: “Say yours, then say theirs.” Open the roster. Repeat the roster, not a singsong of “nights.” Write so the next shift could follow it.

Desk cousin. A recap that says “we decided” versus the notes.

Floor or shift. Posted hours versus the clock is a neighbour, if the rumor is about when the desk is open.

Try-it 4 — Two quotes, floor

Time. 5–10 minutes.

Materials. Same SKU; two quotes. What each actually includes.

Safety. After the walk-through. No lecture in the aisle. No gotcha against a shop.

The move. Say yours, then say theirs. Then the line items.

How to run it. “You mean quote B is cheaper.” “You mean quote A includes delivery.” Open both. Write the difference so a colleague who missed the walk-through could follow it.

Desk cousin. Invoice versus purchase order, from Chapter 4.

Floor or shift. Pad in a break.

Try-it 5 — “We always,” floor

Time. 3–8 minutes.

Materials. Last quarter's process document; the last three actual tickets.

Safety. No humiliation. A counterexample is a time it wasn't, not a personality.

The move. What do you mean, *always*? Then the document and the tickets.

How to run it. Say yours. Say theirs so they would nod. “You mean we start with X because that is the process.” Open last quarter’s document. Open the last three tickets. Repeat the document, not a singsong of “always.” Write the line.

Desk cousin. “We always decide in stand-up” versus the notes.

Floor or shift. Shop, clinic protocol already in use, warehouse, retail — the place you already work.

Talk/meeting box

Exact opening question. “Say yours, then say theirs.”

Not: “Did we have a good meeting?” Not: “Who is right?” Not: “What is the vocabulary word for steelman?”

Follow-ups.

1. What do you mean by that verb — *leading, always, jumped, nights?* (meaning)
2. How do you know? (reason)
3. Always — can we think of a time it wasn’t? What else could be true? (alternative)
4. What’s a second place we could look? (check / second source)
5. What’s the strongest fair version of what they mean? You may change your mind. Can you write that so a colleague who missed the meeting could follow it? (steelman / revise / write)

How to wait. After you ask, a slow three.¹⁴⁴ After they stop, wait again. If they are mid-reason, do not cut them off. Look at the spec, the graph, the roster, or the notes, not at the person, if the silence is hard. After “write the sentence,” that is task-completion time. At work the shrug often looks like refusal. Do not rescue it. Wait in private when the power difference is steep. You are not citing science-class wait-time as a stand-up trial.¹⁴⁵

What a stuck silence usually means. They think steelman means giving up. They think listening means awarding the vendor. They repeated the rumor instead of the

¹⁴⁴Robert J. Stahl, “Using Think-Time and Wait-Time Skillfully in the Classroom,” ERIC Digest ED370885 (1994). About three seconds as a convention, classroom origin. Not a workplace trial.

¹⁴⁵Mary Budd Rowe, wait-time papers, elementary science class, 1972/1974. Not a stand-up randomised trial. Do not mash Rowe and Stahl as if they were one paper.

check. They imported a live controversy. Wait-time 1 was zero. They are guessing what you want. The question was a recitation item in disguise. They offered a why-story where a look was asked for. They think “critical” means finding fault. They think a finished-looking paragraph is already a check.

If they are stuck. Smaller question. “What do you mean, *leading?*” Wait. Point at the object. Do not fill the silence with the right steelman. Do not turn the meeting into a fight. If a live controversy has entered the room, stop. Return to the vendor page, the deck, the roster, or the invoice.

Blocked, then mixed

Blocked. First sitting: one cooler claim you already have, one steelman so they would nod, one named look. Stay there until the sequence has a name.

Mixed. Then a vendor. Then a deck. Then a roster rumor. Then an “always.” Mix so the move is not glued to one object. Near transfer to another vendor page, another recap, another rumor about who works when is the honest hope. Far transfer to a heavier claim in a domain you do not know is not a promise.¹⁴⁶

One incorrect example to diagnose

A thread. The first message forwards “nights.” Five replies repeat “nights.” Someone writes a fluent paragraph: “Nights is the plan; everyone knows; we always do nights.” No roster. No named note. The paragraph looks like work.

This is an illustration, not a reported workplace. The rumor is about *who works when*.

What it usually means. A sticky sentence got practiced. The mash of a why that makes sense — emails report schedules — with evidence that this claim is true. Continued influence: the look, if it happened, did not travel.¹⁴⁷

¹⁴⁶Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, NSW occasional paper, 2019. Sequence and revisit; maxims without knowledge will not implement; a later heavier claim needs the later domain’s knowledge. Near is not far.

¹⁴⁷Stephan Lewandowsky, John Cook, Ullrich Ecker, and 19 co-authors, *The Debunking Handbook 2020*. Backfire effects occur only occasionally; do not refrain from correcting out of fear of backfire. Continued-influence effect is the sticky problem. Detailed refutation outperforms a bare “that’s false.” Illusory truth: repetition increases judged truth. Work translation in this chapter —

What to say. “Say yours, then say theirs. You mean nights because the forwarded email said nights. What’s a second place we could look?” Open the posted roster. Repeat the roster, not a singsong of “nights.” Then: “Can you write that so a colleague who missed the thread could follow it?” If the sentence cannot survive the check, correct the sentence.

If it isn’t clicking

Steelman is heard as giving up, or listening is heard as awarding the vendor.

Slow down. “Say yours first. We still look. Changing your mind if the check goes the other way is not losing. We still decide what to do.” Model two fair sentences, then they write the next two. Dual focus is promptable. It is not the default.¹⁴⁸ If needed, write both sides on two pads so the object stays visible.

The rumor keeps getting practiced, or a live controversy keeps arriving as the object. Hear it. Repeat the check: the spec, the axis, the roster, the notes, the other quote. Lead with what you looked at. A detailed look outperforms a bare “that’s false.”¹⁴⁹ If the thread has left ordinary work stakes, stop. Postpone. Pick the vendor page, the deck number, the shift rumor about who works when, or the invoice. Practice the moves. The later, heavier claim will need knowledge this sitting does not have.¹⁵⁰ Backfire is occasional. It is not a reason to skip the look,

vendor page, deck chart, shift rumor about who works when — is a translation, not a workplace randomised trial of the handbook.

¹⁴⁸Matthew Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, no. 2–3 (2001): 135–153. Community-college young adults versus seventh- and eighth-graders; adults more strategic, still own-side. Lab topic of that study is not a book exercise; copy shape. Keith E. Stanovich, Richard F. West, and Maggie E. Toplak, “Myside Bias, Rational Thinking, and Intelligence,” *Current Directions in Psychological Science* 22, no. 4 (2013): 259–264. Myside bias shows little relation to intelligence, including in university samples, often literally zero in naturalistic between-subjects designs. Access date for sources in these notes: 2 September 2026.

¹⁴⁹Stephan Lewandowsky, John Cook, Ullrich Ecker, and 19 co-authors, *The Debunking Handbook 2020*. Backfire effects occur only occasionally; do not refrain from correcting out of fear of backfire. Continued-influence effect is the sticky problem. Detailed refutation outperforms a bare “that’s false.” Illusory truth: repetition increases judged truth. Work translation in this chapter — vendor page, deck chart, shift rumor about who works when — is a translation, not a workplace randomised trial of the handbook.

¹⁵⁰Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, NSW occasional paper, 2019. Sequence and revisit; maxims without knowledge will not implement; a later heavier claim needs the later domain’s knowledge.

and it is not a reason to skip correction.

The sitting becomes theatre — volume, instant surrender, a straw, a finished-looking paragraph, a machine steelman pasted as yours. Stop. Point at two named sources. Smaller claim. “Say yours, then say theirs. What can we look at?” Close the tool. The unaided steelman is yours. Correct a straw without crushing the talk: “Is that a version they would sign?” If it is still theatre, postpone. Ask a human colleague when the domain is the wall — what the spec lists, what the roster says. That is knowledge first, not a character flaw.

Go ahead when you can say yours, say theirs so they would nod, look at one named source, write the check, and change the sentence when the look goes the other way. Stay longer when “because we always” is still the check. A job title does not settle it.

Tools, including AI

Point back to the work-hour box. Keep this short.

After you have tried, you may ask a tool to explain this page to yourself from a named source, or to give extra questions you then vet. Throw out any that import a live controversy, a secret, a pretend expert, a colleague’s private performance, or a civic-web item. A labelled SCRIPT after the try is for you to say — not a paragraph to copy as your reason, your email, or your steelman. A hint after an attempt is not the finished reason. A diagnosis of a reason you already gave is yours to hear and decide.

You may not ask a tool to write your steelman and paste that text as yours. You may not park an unsupervised chatbot as the only partner. You may not invent an expert identity. You may not grade a colleague from a detector score. You may not print a certificate of thinking. You may not ask a tool to invent what the spec lists, what the roster says, or what the spreadsheet totals. You may not hand yourself an agent as a thinking partner for the first attempt. You may not put confidential workplace data into a public consumer tool.

A fluent machine steelman is performance, not your check. After any tool, write theirs so they would nod, from the page itself, tool closed. One crutch sentence, labelled: an unguarded chatbot can make practice look better and leave you worse

Near is not far.

when the window is closed.¹⁵¹ Open *your* employer’s acceptable-use page. Label it. A thinking tool is not a secret friend.

What “done enough” looks like

Done enough is skill, not a job title, and not a course called Disagreement 1. The title a stranger can read is still the job you hold. The dated note lives inside that file. It is not Critical Thinking I. A workshop certificate is a receipt, not the work.

This heading also carries the life of the habit. Growing fluency is more miles on the same road, not a new faculty and not a thinking birthday. Revisit the same question next Tuesday.¹⁵² “I used to think... now I think...” after a look is allowed, and audible.¹⁵³ Ground rules — we listen; we ask what you mean; we give a reason; we can change our minds; we write a reason a colleague who missed the meeting could follow — then the next real disagreement. Posting the rules is not practice. Placement is by skill, not job title. Not a diploma. Not a certificate of thinking. Not a closing window. How often is enough? The honest pitch is *often enough that the question is familiar*. Frequency at the desk grain is unknown in the sources this book opened. A paperback cannot supervise.

Checklist before moving on.

- Lets the other person finish.
- Asks what they mean before arguing.
- Asks how they know, and hears a wish as a wish.
- Leaves this page for a second named source — spec, notes, other quote, axis, roster, process document.
- Says theirs so they would nod, then still looks.
- Repeats the check, not the rumor.
- Writes a dated sentence a colleague who missed the meeting could follow.

¹⁵¹Hamsa Bastani and colleagues, *Proceedings of the National Academy of Sciences*, 2025. Turkish high-school math. One crutch sentence, labelled not a workplace trial.

¹⁵²Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, NSW occasional paper, 2019. Sequence and revisit; maxims without knowledge will not implement; a later heavier claim needs the later domain’s knowledge. Near is not far.

¹⁵³Robert H. Ennis, public list of dispositions: “look back and change one’s mind when the evidence warrants” as ordinary-language cousin of “I used to think... now I think...” Audible after a look. Not a diploma.

- Changes the sentence when the check goes the other way, without treating that as losing.
- Stays with ordinary work stakes when a hotter claim arrives.
- Does not treat steelman as giving up, or listening as awarding the vendor.

Desk, done enough. Say yours, then say theirs, on a vendor or a deck, then write the check.

Floor or shift, done enough. A shift rumor about who works when, or an “always”; repeat the roster or the process document, not the rumor.

What can still be thin: unaided far transfer to a heavier claim in a domain you do not know; a culture that rewards volume; own-side talk that returns the moment no one rings the bell.¹⁵⁴ Those can wait, or they can be next Tuesday’s revisit. They are not a window that closed. Approaching disagreement is not a bicycle. Knowledge of *this* spec, *this* roster, *this* axis, still has to be in the room.¹⁵⁵ Adulthood is not a thinking birthday. A job is not arrival.

An adult who felt rusty should be able to say, after this chapter: I understand the sequence. I know what to say tomorrow at the desk and on the floor. I know what an own-side “because” sounds like when it is not yet a check. I know what a rumor being practiced sounds like. I know what a finished-looking paragraph without a named source looks like. I can say mine, then theirs. I can look. I can write the check. I can do this.

Next Tuesday, open the same kind of object. Ask the same opening. Write the new line. That is the habit. It is not a diploma ceremony.

¹⁵⁴Matthew Felton and Deanna Kuhn, “The Development of Argumentative Discourse Skill,” *Discourse Processes* 32, no. 2–3 (2001): 135–153. Community-college young adults versus seventh- and eighth-graders; adults more strategic, still own-side. Lab topic of that study is not a book exercise; copy shape. Keith E. Stanovich, Richard F. West, and Maggie E. Toplak, “Myside Bias, Rational Thinking, and Intelligence,” *Current Directions in Psychological Science* 22, no. 4 (2013): 259–264. Myside bias shows little relation to intelligence, including in university samples, often literally zero in naturalistic between-subjects designs. Access date for sources in these notes: 2 September 2026.

¹⁵⁵Daniel T. Willingham, “Critical Thinking: Why Is It So Hard to Teach?,” *American Educator*, Summer 2007; and Willingham, NSW occasional paper, 2019. Sequence and revisit; maxims without knowledge will not implement; a later heavier claim needs the later domain’s knowledge. Near is not far.

Chapter 9

Records

There is often **no** “Critical Thinking I” as a job title. A stranger — a manager, a reviewer, a regulator, a later hiring panel, an editor, a partner — already reads the job’s own record. The honest adult work product is a **dated note**, a **short written reason plus the check**, and an **optional two-source file, inside the job’s real title**. Never Critical Thinking I. Never a LinkedIn badge as the proof. Never “I sat Watson–Glaser” as the year’s work.

This week you can date one short written reason plus the check so a colleague who wasn’t there could almost read it, and keep the title as the job you hold. Today you try the question. You write the line.

Illustration. An open notebook, a dated reason, two named sources, a check. No diploma seal. No logo. No “Critical Thinking I.” No people.

What you are making

You are making a **dated notebook**, a **short written reason plus the check**, and, stacked over the year, an **optional two-source file**. Stimulus named. Sources named. You are not making a national diploma. You are not making Critical Thinking I. You are not making a thinking licence.

Titles a stranger already understands: the job you hold — analyst, nurse, clerk, lead, tech, manager. Put the questions and the short written reason inside that file. Combining objects is allowed. Combining does not mint a new title. Two vendor pages plus a deck chart plus a dated sentence is still the project, still the workpaper, still the chart, still the story file.

Here is an illustration, labelled as such, not a reported workplace:

2 September 2026. Question: Does the vendor's "40% reduction in defects" slide match the methods note in the appendix we were sent on 28 August? Looked: slide 12; appendix p. 4 (sample n = 18, no control line). Check: the 40% is before/after, not versus the parallel line. What would change my mind: a methods paragraph with the comparator and a longer window. Files: vendor-q3.pdf; appendix-28aug.pdf.

That is readable by a manager who was not in the room. It is not a badge. Place by the job, not by the brand.

Auditor: the workpaper. **Nurse:** the chart. **Reporter:** the story file with sources. **Engineer:** the calculation package or the hazard memo. **Analyst:** the dated note in the repo. **Manager:** the decision log with the because and the unless.

A regulated job already names documentation as part of the work. In audit, for example, the workpaper is how a stranger later sees why the evidence was enough.¹⁵⁶ That is a domain duty. It is not proof of a transferable faculty, and it is not a reason to retile the year.

Homeschool transcripts, university entrance lists, and athletic-eligibility cores are out of band here except one sentence: those pages name English, science, history, sometimes philosophy — not Critical Thinking I — and this book does not import them as workplace law.¹⁵⁷

Instruments explain, they are not the gym

Some graduate schemes, some colleges, and some employers sit a thinking inventory. A sitting is a sitting. A score is a score. This book will not become a test-prep franchise.

¹⁵⁶Public Company Accounting Oversight Board, AS 1000, including documentation of the auditor's work; effective for audits of fiscal years beginning on or after 15 December 2024 (later for smaller firms). International Standard on Auditing 200 names a questioning mind and a critical assessment of evidence as a professional attitude. Audit domain, not a transferable faculty. Access date for sources in these notes: 2 September 2026.

¹⁵⁷Homeschool records and transcripts, University of California A–G, NCAA core: one sentence, out of band. Those pages do not name Critical Thinking I as a required course. This book does not import them as workplace law.

Watson–Glaser III is a timed inventory some organisations use in screening. It is a product with a model of recognising assumptions, evaluating arguments, and drawing conclusions. Independent commentary has raised construct-validity questions, including instructions that can punish people who know more logic, not less.¹⁵⁸ Treat it as what some schemes sit. Do not treat it as this year’s work. This book does not claim that most companies require it, and it does not print a 2026 sitting volume or a 2026 fee.

The California Critical Thinking Skills Test is college-calibrated, timed, scenario-based. It is designed for undergraduate and graduate students. A score is a score. It is not the meaning of critical thinking, and it is not this notebook.¹⁵⁹

INSIGHT Business Professional and related business forms are employment cousins of that family: mindset plus workplace scenarios, screening and development. Quote-on-request. Still a sitting.

Cornell Level Z is a publisher form listed for advanced high school, college, graduate students, and other adults, with employment among named uses. Still a vendor sitting, still not a diploma.

The Halpern Critical Thinking Assessment was a constructed-response-then-forced-choice inventory for adults. It has since been commercially retired from the system that sold it.¹⁶⁰ Do not send a 2026 adult to buy it as if it were in print.

LSAT Analytical Reasoning — the “logic games” section — was discontinued after June 2024.¹⁶¹ Do not send a career-changer to a discontinued section. SAT and

¹⁵⁸Pearson TalentLens, Watson–Glaser III product pages and Pearson Efficacy Report, April 2020 — publisher materials, not an independent assurance. Forty items from an item bank, typically about 30 minutes. Kevin Possin, “Excuse Me, but What Constitutes a Demonstration of Critical Thinking?,” *Informal Logic* 34 (2014): construct-validity issues, including instructions that can punish logic knowledge. Robert M. Bernard and colleagues, 2008: Watson–Glaser subscales highly intercorrelated; better treated as one general scale. “Thousands of organisations” is publisher copy. This book does not invent “most companies require Watson–Glaser,” a 2026 sitting volume, or a 2026 list price.

¹⁵⁹Insight Assessment, California Critical Thinking Skills Test product page: designed for undergraduate and graduate students; about 55 minutes, 40 scenario-based items. College-calibrated. Explain, not the meaning, not the gym. INSIGHT Business Professional: employment cousin, quote-on-request.

¹⁶⁰Heather Butler, *Journal of Intelligence*, 2024: the Halpern Critical Thinking Assessment “was available for a fee through the Vienna Test System ... but has since been retired.” Residual licensing not independently confirmed this pass.

¹⁶¹Law School Admission Council: Analytical Reasoning (“logic games”) discontinued after

ACT remain reading objects at these depths, not thinking diplomas. SAT Subject Tests were discontinued in June 2021. There is no AP Critical Thinking.

If your employer sits one of these, sit it if you choose to, or if the scheme requires it. Do not confuse the sitting with the work. Do not refuse a sitting your employer sits on this book's advice. Then return to the check. The year's honest page is still the dated note inside the job.

What employers rate is not a gym either

A 2026 spring employer survey put Communication and Teamwork at 4.5 on a five-point importance scale and Critical Thinking at 4.4. Not first. About 93.5 percent rated critical thinking very or extremely important; 50 percent rated recent graduates very or extremely proficient on it. The survey had 185 total respondents, including 142 members, 19.9 percent of eligible member respondents. It is a survey of members plus some nonmembers, not a trial, and not a probability sample of all employers.¹⁶² Older copy that put critical thinking first is not current.

A different international employer survey, a different phrase: analytical thinking remains a top core skill, with seven in ten companies considering it essential.¹⁶³ Do not collapse that phrase into a college inventory or into Watson–Glaser.

A U.S. occupational database lists critical thinking as a skill across hundreds of occupations — 894 shown on the page this book opened — and no occupation titled Critical Thinking I.¹⁶⁴ Flexibility in hiring for people without a degree is

June 2024. SAT Subject Tests discontinued June 2021. College Board and ACT reading tests may include graphics as literacy items; they are not thinking diplomas. There is no AP Critical Thinking.

¹⁶²National Association of Colleges and Employers, *Job Outlook 2026 Spring Update*, survey 12 February–17 March 2026. 185 total respondents; 142 NACE employer members = 19.9 percent of eligible member respondents; 43 additional nonmember responses. Importance (5-point; n = 153): Communication and Teamwork 4.5, Critical Thinking 4.4, Professionalism 4.3. Share rating the competency very or extremely important / share rating recent graduates very or extremely proficient: Critical Thinking 93.5 percent / 50.0 percent. Survey, not a trial. Not first. Competency definition (March 2025 revision) is a framework, not a trial.

¹⁶³World Economic Forum, *Future of Jobs Report 2025*, published 7 January 2025. “Analytical thinking remains the top core skill for employers, with seven out of 10 companies considering it as essential.” Phrase is analytical thinking. Different survey, different phrase. Journalism / employer survey, not a randomised trial.

¹⁶⁴O*NET OnLine, skill 2.A.2.a Critical Thinking, site updated 19 May 2026 as opened: 894 oc-

hiring flexibility. It is not a licence to skip the check.

None of those pages is the gym. The gym is Tuesday's artifact and the sentence a colleague could follow.

Place by skill, not job title

Can you lift a number from a chart and say where the picture starts? If not, you are still in the graphs chapter. Can you write claim, because, check, so a colleague who missed the meeting could follow it? If not, you are still in the written-reason chapter. Can you say yours so the other person would nod, then look at the spec? If not, you are still in the opening of disagreement. High-sounding *titles* cannot skip those moves. A director may still be in the early band of a chapter. A newer hire may already write the check unaided. Skill first.

What to keep in the file

Keep what a stranger can use.

A dated note. Stimulus named. One question. What you looked at. The date is 2 September 2026 or whatever day it actually was — not “Q3 thinking.”

A short written reason plus the check. I claim X because Y. We checked by Z. Someone who disagrees would say... Sources named. If “it” needs the meeting, rewrite.

An optional two-source file. Two vendor pages, recap versus notes, spec versus review, invoice versus PO, two handoff notes, product page versus data sheet. Named. Dated. What would change your mind.

Keep them *inside* the job's real title. The auditor's workpaper. The nurse's chart. The reporter's story file. The engineer's calc package. The analyst's repo note. The manager's decision log.

Life of the habit, here: the note is how the habit leaves a trace. Same question next Tuesday, dated. “I used to think... now I think...” after a look. Not a diploma

cupations shown. No occupation titled Critical Thinking I. SHRM skills-first copy about accepting nondegree pathways is hiring flexibility, not a skip-the-check licence.

ceremony. Growing fluency is more miles on the same road, not a new faculty and not a thinking birthday.

What not to keep as the proof: a workshop certificate; a LinkedIn badge; a percentile; a pin; a franken-title. Those can be receipts. They are not the work.

How much is enough? A line after a meeting is enough for a first week. A two-source file stacked over a quarter is enough for a later reader. A binder labelled “Critical Thinking, Level 2” is not more lawful than the line. Frequency the evidence actually supports at the desk grain is unknown. The honest pitch is often enough that the question is familiar, and dated enough that next Tuesday can find last Tuesday.

If writing itself is the wall — a freeze at the blank page, a workplace that only accepts a slide — keep the parts anyway: claim, because, check, date, named sources. A human colleague who already writes the file you have to write can help with the writing part. You still own the thinking question: could a colleague who missed the meeting follow it?

If a later reader asks

A manager who sees “Critical Thinking I, A” has been given a slogan. A manager who sees “Q3 vendor file — two sources; dated note; 40% slide versus appendix p. 4” has been given a year they can use.

A later hiring panel will read the job you held. Put the two-source file inside that job. Adding “Critical Thinker” to a résumé does not make the year more lawful. It may confuse the next stranger.

If a scheme sits Watson–Glaser, or a college inventory, or Cornell Z, the sitting can be named as a sitting. “I sat an inventory in March” is a fact about a morning. It is not the year’s work. The year’s work is the file.

You try this week

Date one short written reason plus the check so a colleague who wasn’t there could almost read it. Keep the title as the job you hold. If you already have two named sources, clip them to the same note. Next Tuesday, same kind of object, new date. That is the record. It is not a diploma.

Chapter 10

Resources

This chapter names programs a working adult actually meets. It does not rank them. It does not sell them. Fit is load, style, cost shape, and whether a desk or a shift can use it this week. Vendor or charity means the organisation sells or trains. Reach numbers on an organisation's page are that organisation's copy, not a census and not an effect size.

Whenever Philosophy for Children is named below, copy this sentence with it: a large English effectiveness trial in 2021, 198 schools, highest security rating, found **0 months** extra progress on reading and maths, including for pupils eligible for free school meals. That trial was Key Stage 2 / Year 6, about age 10–11, **not this book's trial**.¹⁶⁵ Dialogue may still be worth the hour as dialogue. It is not a proven engine of generic critical thinking or of adult workplace attainment at scale. An earlier, smaller efficacy trial had found about **+0.12** in reading and **+0.10** in maths; that result did not replicate at effectiveness scale.¹⁶⁶ Teacher glow is not the outcome.

What Works Clearinghouse, searched again on 2 September 2026, still has **no**

¹⁶⁵Education Endowment Foundation / National Foundation for Educational Research, Philosophy for Children effectiveness trial, 2021. 198 English primary schools; 5/5 padlock security rating; free-school-meals reading 0.02 (0 months); whole-cohort reading and maths **0 months**. Labelled Key Stage 2 / Year 6 $\approx$ 10–11, **not this book's trial**. Access date for sources in these notes: 2 September 2026.

¹⁶⁶Stephen Gorard, Nadia Siddiqui, and Beng Huat See, EEF efficacy trial, 2015. 48 English primary schools; Year 5; reading +0.12, maths +0.10. Primary outcome was attainment, not a thinking test. That efficacy result did not replicate at effectiveness scale in 2021. Still KS2, still not this book's trial.

named adult or workplace critical-thinking intervention report.¹⁶⁷ Neighbouring reports exist for college advising and for labour-market programmes that mention the phrase in weekly meetings. They are not a thinking franchise for this job. This book does not put a clearinghouse logo on a philosophy kit.

Paper map before any app.

This week you can name *why* tonight's object fits — usually two named pages plus one question plus a one-page reason, or a paper map. Today you try that question. You are not on a shopping task.

Illustration. Two pages, a paper map of three boxes, a desk, a dated notebook. Not a Gold pin. Not a wheel. Not a readable app. Not a logo. No people.

How to read a program

Four questions do the work.

Load. Are you asking one question on two pages you already opened? Sitting with a four-hundred-page manual? Booking a two-day seminar? Building a brown-bag series?

Style. Two named sources and a dated sentence? A paper map of claim / because / unless / check? A short published guide with a set of terms in ordinary language? A university catalogue? An ethics PD hour?

Cost shape. Nothing tonight. A pad you already own. A seminar priced for organisations. Fees this book did not verify. Confirm a current price on the page before you buy. This chapter does not invent one.

Desk or shift without a school. Can one adult use it this week on a recap, a vendor page, a deck chart, a handoff note, or does it assume a class, a trainer, a pin, a login, or a franchise hour?

¹⁶⁷Institute of Education Sciences, What Works Clearinghouse product search, re-opened 2 September 2026: no named adult or workplace critical-thinking intervention report; no Philosophy for Children report (the English effectiveness trial of that school programme found 0 months on KS2 reading and maths, labelled Year 6 $\approx$ 10–11, not this book's trial); no argument-mapping report; no named Paul–Elder intervention report. Academic Concept Mapping (Nelson 2007; WWC Study 89713) is college-advising GPA. Project QUEST mentions critical thinking in weekly meetings; rated outcomes are education and labour-market.

What a stranger can still read is the job you hold. Combining two honest objects does not mint a new title. Never Critical Thinking I.

Kitchen, desk, shift: nothing spent

Two named pages or two named people. A graph already in the slide deck. A short written reason — claim, because, unless, check — dated. Record: a dated note in the project folder, not a course name. Default.

This is not “failing the specialist.” It is the object this book already taught. Wait. One question. One look. One line.

Fit. Tonight, one adult, low load, low cost.

Title. The job you hold. Never Critical Thinking I.

Paper argument map

Claim / because / unless / check. Three boxes on a pad. Dated. In the file. Method, not an app. Paper first.

University argument-mapping courses have promising pre–post numbers on college inventories. There is no large randomised trial — the author’s words.¹⁶⁸ Three boxes on paper in a meeting is a type, not that semester, and not a standardised-test gym. Software with a login is optional if the organisation already pays. A thread tool is optional if the team already lives there. None outranks the pad. This chapter does not print a 2026 subscription price.

Fit. Visual map, no login, low cost. A site with a licence the organisation already has is a moderate extra, still after paper.

Title. The project file. Never Argument Mapping I.

¹⁶⁸Tim van Gelder, “Using Argument Mapping to Improve Critical Thinking Skills,” in *The Palgrave Handbook of Critical Thinking in Higher Education*, ed. Martin Davies and Ronald Barnett (2015). University; “Unfortunately, no such study has ever been conducted” — no large randomised trial, his words. Pre–post college-inventory numbers are university. For this book: three boxes on paper before any app. Harrell 2008 is a university philosophy-course neighbour: even pencil-and-paper mapping moved inventory scores in that course — still university.

A brand with elements and standards

The Foundation for Critical Thinking offers two-day seminars for leaders and managers, a set of elements and standards in ordinary language, custom webinars, and longer programs. That is a **brand**, not a meta-analysis, and the clearinghouse has no intervention report under that name.¹⁶⁹ ¹⁷⁰ Organisation reach numbers are Foundation copy, not a census. Do not franchise the wheel. Do not put a level number on a résumé as if it were a licence. Costs are quote-on-request; this chapter does not invent a 2026 fee.

Ordinary language is allowed: purpose, question, information, inference, assumption, point of view — said as those words, not as a badge. The move is still the look. The file is still the proof.

Fit. Named vocabulary for managers who want a shared list of terms, high load if you book the seminar. Not Abrami. Not a wheel of evidence.

Title. Never Paul–Elder as the job. Never a certified-thinker line as a licence.

University continuing education

Catalogues exist. Extension certificates in philosophy or ethics are neighbours — philosophy and ethics, not a critical-thinking gym. Massive open courses titled “critical thinking for leaders” are catalogue objects. This stream did not open a 2026 enrollment total or rank them. This chapter does not print a 2026 tuition. The certificate is a receipt. The file is the proof.

Fit. Transcribed university object, high load, high cost shape. Ethics is not this book’s gym. The dated note still lives inside the job you hold.

¹⁶⁹Institute of Education Sciences, What Works Clearinghouse product search, re-opened 2 September 2026: no named adult or workplace critical-thinking intervention report; no Philosophy for Children report (the English effectiveness trial of that school programme found 0 months on KS2 reading and maths, labelled Year 6 $\approx$ 10–11, not this book’s trial); no argument-mapping report; no named Paul–Elder intervention report. Academic Concept Mapping (Nelson 2007; WWC Study 89713) is college-advising GPA. Project QUEST mentions critical thinking in weekly meetings; rated outcomes are education and labour-market.

¹⁷⁰Foundation for Critical Thinking, business and professional groups pages, opened 2 September 2026. Two-day seminars, elements and standards, custom webinars, six-week and twelve-week programs. Brand language. “100,000 educators and professionals” is Foundation copy. 2026 fee quote-on-request, not printed here.

Ethics professional development

The Association for Practical and Professional Ethics lists **Deliber8** as ethics PD for business, nonprofit, and government leaders, typically with the APPE conference, hosted with a university practical-ethics centre.¹⁷¹ Adult ethics PD. Not a thinking sitting. Not a high-school ethics bowl. Not an undergraduate bowl. Steal the **format** — collaborative reason, not debate-win — for a staff brown-bag. Skip this year's live political case set. Fees not verified this pass.

Fit. Ethics PD, moderate to high load if you travel. Format yes. This year's case set no.

Title. Never Ethics Bowl I. Never Deliber8 as the job.

Instruments, pointed back

Watson–Glaser, the California Critical Thinking Skills Test, INSIGHT, Cornell Level Z: a sitting, employer- or school-gated, a score. Explain, not the gym. The Halpern inventory is commercially retired. Chapter 9 already named them. If a scheme sits one, sit it if you choose to. Then return to the check.

School objects, one sentence

Philosophy for Children, a UK training charity now under a new name, a UK philosophy-in-schools charity, novels written for a circle of enquiry, and a high-school ethics bowl are **out of band** except this sentence. Whenever Philosophy for Children is named: a large English effectiveness trial in 2021, 198 schools, highest security rating, found **0 months** extra progress on reading and maths, labelled Key Stage 2 / Year 6 $\approx$ 10–11, **not this book's trial**.¹⁷² Not adult workplace training.

¹⁷¹Association for Practical and Professional Ethics, Deliber8 pages, opened 2 September 2026, with the Center for Practical Ethics / Prindle. Adult ethics PD for business, nonprofit, and government leaders. Not a thinking sitting. Not a high-school ethics bowl. Not an undergraduate bowl. 2026 fee not verified this pass.

¹⁷²Education Endowment Foundation / National Foundation for Educational Research, Philosophy for Children effectiveness trial, 2021. 198 English primary schools; 5/5 padlock security rating; free-school-meals reading 0.02 (0 months); whole-cohort reading and maths **0 months**. Labelled Key Stage 2 / Year 6 $\approx$ 10–11, **not this book's trial**. Access date for sources in these notes: 2 September 2026.

An earlier efficacy trial's +0.12 / +0.10 did not replicate at scale.¹⁷³

Professional domains as domains

These are jobs. They are not a transferable faculty. They are not Watson–Glaser. They are not a licence this book can print.

Audit — professional skepticism. A questioning mind; a critical assessment of evidence; evaluate supporting *and* contradicting evidence; not assuming management is honest or dishonest.¹⁷⁴ Audit. The analogue this book may steal: do not assume the comfortable source is honest or dishonest; look at the contradicting page; write down why the evidence is enough. The workpaper is a dated two-source file.

Nursing — clinical judgment. A measurement model for an entry-level nursing exam names recognising cues, analysing cues, prioritising hypotheses, generating solutions, taking actions, evaluating outcomes. It neither defines nor redefines clinical judgment. The next-generation exam launched 1 April 2023.¹⁷⁵ Nursing. Not a generic thinking gym. Not a college inventory. A nurse's stranger-readable product is the chart. This book will not run a nursing curriculum.

Journalism — verification. A fact-checking network's code asks for more than one named source on key elements, among other commitments, as organisational signatory status — not a personal thinking badge.¹⁷⁶ A reporter's two-source file

¹⁷³Stephen Gorard, Nadia Siddiqui, and Beng Huat See, EEF efficacy trial, 2015. 48 English primary schools; Year 5; reading +0.12, maths +0.10. Primary outcome was attainment, not a thinking test. That efficacy result did not replicate at effectiveness scale in 2021. Still KS2, still not this book's trial.

¹⁷⁴Public Company Accounting Oversight Board, AS 1000, opened 2 September 2026: questioning mind; critical assessment of audit evidence and other information; evaluate supporting and contradicting evidence; not assuming management is honest or dishonest. Effective for audits of fiscal years beginning on or after 15 December 2024 (later for smaller firms). IAASB ISA 200: same attitude-words as a requirement. Audit. Not Watson–Glaser. Not a transferable faculty.

¹⁷⁵National Council of State Boards of Nursing, Clinical Judgment Measurement Model / Next Generation NCLEX pages. Measurement model for NCLEX; “neither defines nor redefines clinical judgment”; entry-level nurses; Next Generation NCLEX launched 1 April 2023; Layer 3: recognise cues, analyse cues, prioritise hypotheses, generate solutions, take actions, evaluate outcomes. Nursing. Not a generic thinking gym.

¹⁷⁶International Fact-Checking Network, Code of Principles, opened 2 September 2026: five commitments; assessor criteria; more than one named source on key elements; organisational signatory status, not a personal thinking badge.

is the adult analogue. Leaving the page is a neighbouring civic habit. It is not a workplace trial, and it is not reprinted here as a civic course.

Engineering — judgment and safety. Accreditation criteria for engineering programs name design with public health, safety, and welfare; informed judgments; and engineering judgment on experimental data.¹⁷⁷ Accreditation criteria, not a thinking sitting. The stranger-readable product is the calculation package, the hazard analysis, the design memo with assumptions listed.

Transfer from those trainings onto a generic inventory was not opened. A job is a domain. It is not all domains.

Fit table (use; do not rank)

Tonight's need	Object	Load	Note
Nothing spent	Two pages + one question + a dated reason	Low	Not "Critical Thinking I."
Visual map, no login	Paper claim / because / unless / check	Low	Paper first.
Visual map with a site	A mapping or thread tool the organisation already pays	Moderate	Still after paper.
Named vocabulary for managers	Foundation two-day	High	Brand, not a meta-analysis.
Transcribed university object	Continuing education / Extension	High	Ethics ≠ thinking gym.
Ethics PD	Deliber8	Moderate–high	Format yes; this year's case set no.

¹⁷⁷ABET Engineering Accreditation Commission Criteria, 2026–2027, opened 2 September 2026: student outcomes (2) design with public health, safety, and welfare; (4) informed judgments; (6) engineering judgment on experimental data. Accreditation criteria, not a thinking sitting.

Tonight's need	Object	Load	Note
Sitting	Watson–Glaser / college inventory / INSIGHT / Cornell Z	Gated	Instrument, not the gym.
School objects	Philosophy for Children / ethics bowl / enquiry novels	Out of band	P4C: 2021 0 months , KS2 / Year 6 $\approx$ 10–11, not this book's trial. ¹⁷⁸

You try this week

Which fit object matches tonight — usually two pages plus one question plus a one-page reason, or a paper map. Open the two pages. Ask one question. Write the dated line. Keep the title as the job you hold. If a branded hour is named in the building, remember the school trial's **0 months**, labelled Key Stage 2, and then teach the move on *your* vendor, *your* spreadsheet, *your* roster anyway. The file is the proof.

¹⁷⁸Education Endowment Foundation / National Foundation for Educational Research, Philosophy for Children effectiveness trial, 2021. 198 English primary schools; 5/5 padlock security rating; free-school-meals reading 0.02 (0 months); whole-cohort reading and maths **0 months**. Labelled Key Stage 2 / Year 6 $\approx$ 10–11, **not this book's trial**. Access date for sources in these notes: 2 September 2026.

A Note on Sources

Studies named in the chapters are listed in Notes, in one series at the back. That is where the full citations live, so the teaching pages can stay a teaching voice.

Some items the research behind this book did not open stay out of the teaching voice. I did not invent a coefficient, a statute, or a product feature to fill a hole. If a study is in the chapter, it is in the notes. If we could not open it, it is not used as a finding here.

Program hours, credit claims, and product descriptions in the resources chapter are taken from the companies' own pages. Access date 2 September 2026.