

ALMANAC ACADEMY - MASTERY PHASE

Cognitive Fallacies

Mental Debugging

Debug your own mind. This course teaches learners to identify, isolate and overcome the cognitive biases and logical fallacies that distort reality and human decision-making - from the machinery of perception to deepfakes, scams and filter bubbles.

22 Units

66 Lessons

220 Practice Questions

4 Chapters

22,000 XP Available

Advanced level - Track 14 of 14 - 22,000 XP Available

PARENT & TEACHER STUDY GUIDE

Course data current as of August 14, 2026 - Track 14 of 14 - Cognitive Fallacies

almanac.academy

About this guide

This study guide condenses the complete Almanac Academy course "Cognitive Fallacies" (Mental Debugging) into a printable reference for parents and teachers. Cognitive Fallacies is an advanced curriculum in mental debugging: understanding how the human mind systematically misreads reality, and building the habits that catch those errors before they cause harm. It moves from the machinery of perception and memory, through the psychology of judgment, to the practical defense toolkit for a world of manipulation and misinformation.

Where this fits: the final track in the Mastery Phase (Tracks 12-14), where learners debug their own decision-making. Cognitive Fallacies follows the official certification tracks (CCSE and EFP) and closes the entire 14-track Almanac journey - a capstone in clear thinking.

How to use this guide

- Skim a chapter's outcomes before your child starts it - they are the destination for that week of learning.
- After each unit, ask your child to explain one lesson back to you in their own words. If they can teach it, it has stuck.
- Use the sample questions in the Assessment section to practise the three formats your child will meet.
- Treat the vocabulary list at the back as a quick checkpoint - every term is taught in the lesson it belongs to.
- Pace is flexible: each unit works in short sessions, and XP accumulates as mastery grows.

Chapters at a glance

Ch 1	The Machinery of Mind 6 units - 18 lessons - 60 practice questions	3
Ch 2	Judgment: Framing, Losses & Social Gravity 5 units - 15 lessons - 50 practice questions	7
Ch 3	The Defense Toolkit: Debiasing & Rhetoric 5 units - 15 lessons - 50 practice questions	10
Ch 4	The Digital Battlefield & The Capstone 6 units - 18 lessons - 60 practice questions	13

CHAPTER 1

The Machinery of Mind

The course opens where all biases begin: in the hardware. Learners discover that the brain is a prediction engine, attention is a narrow spotlight, perception sketches rather than photographs, and memory rebuilds rather than replays. Then they meet the shortcuts - heuristics - that usually work but reliably fail at probability.

WHAT YOUR CHILD WILL BE ABLE TO DO

- Explain why the brain relies on predictions and shortcuts - and when 'useful wrongness' becomes a liability
- Describe what attention actually does, including change blindness and the cost of switching
- Understand that perception is a fast sketch built on assumptions, not raw reality
- Explain how memory reconstructs scenes from scraps - and why vivid memories aren't necessarily true
- Identify the three great heuristics (representativeness, availability, affect) and apply base rates before updating
- Recognize the law of small numbers and the conjunction trap

Your Brain, a Best-Guess Machine

Before the signals finish arriving, it takes a smart guess about what's most likely out there and only then checks whether the guess fits.

- Shortcuts for Survival

Think of the brain as a thrifty scientist with a tight deadline: it makes a provisional guess, compares that guess with what the senses report, and updates on the fly. Predict compare update act-over and over-so you can move through the world without stalling. You see this most clearly when the guess peeks through. In the phoneme restoration effect, a cough replaces part of a word in a sentence, yet listeners still "hear" the missing sound because context fills it in (Warren, 1970). Your eyes do something similar every second: there's a literal blind spot where the optic nerve exits the retina, but you don't notice a hole in your vision because the brain paints over it using the best available story. These shortcuts are not sloppy; they're efficient. They buy speed and save energy, and most of the time that's exactly what you want.

- How Expectations Shape Experience

Expectations don't sit quietly in the background; they lean on the steering wheel. Give people a random number before a judgment task and that number tugs their estimates-a classic anchoring demo where a rigged wheel landing on 10 or 65 pulls answers toward it (Tversky & Kahneman, 1974). Attention follows the plan too. When you're busy counting basketball passes in a lab video, many viewers never notice a person in a gorilla suit stroll through the scene. It isn't stupidity; it's a spotlight doing what it was told to do (Simons & Chabris, 1999). Heuristics and attention save time in familiar situations, but the same moves can mislead in noisy, novel, or high-pressure moments.

- Useful Wrongness

Being "wrong" can be useful when it keeps you alive long enough to be right. Veterans in risky jobs often notice a felt mismatch-something about the pattern is off-and act before they can explain why. In Gary Klein's field work, firefighters evacuated a kitchen moments before the floor collapsed; the cues didn't fit the expected pattern, so they updated fast (Klein, 1998). Stress changes the balance. Under heavy load, attention narrows, habits take the wheel, and the brain bets harder on the first good-enough story-speedy, but more error-prone in the wrong context (Sapolsky, 2017). You can counter this with tiny prompts: What else could this be? What base rate am I ignoring? What would change my mind? A one-sentence reflection after decisions strengthens the predict compare update act loop for next time.

Attention Isn't a Spotlight

It behaves more like a narrow, jumpy spotlight that brightens one small patch and leaves the rest in twilight.

- What Attention Actually Does

Picture a friend at a crowded party. They seem to follow your story effortlessly while a dozen other conversations buzz around you. That's selective attention doing its job. In the lab, it looks more precise: if a cue suggests "look left," people are quicker to detect a target that appears on the left-evidence that attention shifts and boosts processing at expected spots (Posner cueing). The boost comes with a price. What you don't spotlight gets less processing. That's why the cocktail party effect feels magical until your own name pops up from the noise; your brain was down-weighting almost everything else until a highly relevant cue tugged the spotlight.

- When You Miss the Obvious

When you lock the spotlight tightly, the rest of the scene goes soft. In inattentional blindness demos, focusing on a counting task leads many viewers to miss a person in a gorilla suit walking through the play (Simons & Chabris, 1999). In change blindness, even big alterations slip by if they happen during a brief disruption. In the street-corner door study, a stranger asking for directions is secretly swapped for a different person while a door passes between them; surprisingly many people don't notice the switch (Simons & Levin, 1998). The lesson isn't "people are oblivious." It's that attention is selective by design. You trade width for depth, often wisely-until context changes and the trade bites you.

- Blink, Switch, and Spillover

Attention also has a rhythm. When two targets appear in a rapid stream, the second one is often missed if it appears about a quarter to half a second after the first-the attentional blink (Raymond, Shapiro & Arnell, 1992). Outside the lab, that rhythm shows up as the feeling of being a beat behind when lots of important things happen at once. We're tempted to brag about multitasking, but most of the time it's rapid task switching with switch costs-extra time and errors as the spotlight detaches and re-locks. On the road, even hands-free phone conversations impair driving because of cognitive load (Strayer & Johnston, 2001). And when you leave one task mid-stream, a bit of your attention clings to it in the next task-attentional residue-which is why unfinished work keeps tugging your thoughts (Leroy, 2009). A few tiny habits help: announce what you're not monitoring, set simple cues to widen your scan ("name one thing outside the focus"), and, when switching, write a ten-second "last line" note so the old task lets go.

Perception Can't See "Raw Reality"

Most days the sketch is so good you never notice the pencil marks.

- Seeing with Assumptions

Stare at two equal lines with different arrowheads and you'll swear one is longer. Your brain is reading a flat drawing as a bit of 3-D-corners, corridors, perspective-and it quietly "corrects" the lengths to fit that story (Müller-Lyer). The same hunch plays out with the Ponzo rails: a bar "farther away" looks larger. My favorite party trick is the hollow-mask: a concave face that looks convincingly like a normal face because faces are supposed to bulge outward. Expectations aren't decoration; they're part of what you see.

- When Senses Disagree

Your senses don't file separate reports; they meet and settle on one version. In the McGurk effect (McGurk & MacDonald, 1976), the mouth shapes "ga," the audio says "ba," and you often hear "da." In the sound-induced flash demo, one flash paired with two beeps often looks like two flashes (Shams, Kamitani & Shimojo, 2000). The point isn't that a sense "wins," it's that the system prefers a single, workable story-even if that story bends a detail.

- Owning a Body the Brain Invents

Body ownership feels obvious until you nudge the timing. Watch a fake hand being stroked while your hidden hand is stroked in perfect sync, and the fake begins to feel like yours (Botvinick & Cohen, 1998). Break the rhythm and the feeling fades. Line up the cues and your inner body map updates-because the goal is coherent control, not perfect reporting.

Memory Reconstructs, Not Records

Memory behaves less like a video file and more like a story you retell.

- Reconstruction on the Fly

Ask witnesses, "How fast were the cars going when they hit?" and you'll get lower speed estimates than if you ask when they smashed. A week later, the "smashed" group is more likely to "remember" broken glass that never existed (Loftus & Palmer, 1974). What you recall is the story rebuilt from traces plus suggestion; wording seeps into the rebuild. A simple habit helps in real life: prefer open questions ("What did you notice?"), and write down details before hearing other accounts.

- Vivid Isn't Veridical

Flashbulb memories feel frozen, but follow-ups tell a humbler tale. When people describe where they were during a major event and are re-tested months later, the details often shift-even as confidence stays sky-high (Talarico & Rubin, 2003). Emotion flags an experience as important; it doesn't lock every pixel. Treat confidence as a feeling about the story, not a guarantee about the facts.

- Smart Guesses and Source Confusion

Your brain is a meaning hunter. After studying words like bed, pillow, dream, nap, many people later "remember" sleep even though it never appeared-that's the DRM effect (Deese; Roediger & McDermott). The guess usually serves you, but it blurs source monitoring: what was in the list vs. what your mind supplied, who told you vs. where you inferred it. You can reduce mix-ups by noting the source ("saw in report," "heard from Ana") alongside the fact.

Heuristics 101 (The Good Shortcuts)

Your brain carries a pocketful of shortcuts.

- Representativeness in the Wild

Pattern-matching is powerful. In classic studies, a description that "sounds like" a particular profession nudges people to choose it and forget the base rates. In the "Tom W." problem (Kahneman & Tversky, 1973), a personality sketch leads many to pick "computer science" even when told only a small fraction of students are in that field. The mind grabs the stereotype and treats it like a probability. A simple counter-move: ask, Out of 100 people here, how many are actually X?

- Availability's Vivid Pull

If examples pop up easily, they feel common. After heavy news coverage of a plane crash, people estimate flying risks higher (Lichtenstein, Slovic, Fischhoff et al., 1978). In another demo, folks say more words start with the letter K than have K as the third letter-because first letters are easier to search in memory (Tversky & Kahneman, 1973). Availability tracks ease, not truth; ease is influenced by recency, vividness, and media.

- Affect: When Feelings Steer

Feelings act like quick labels. When people like a technology, they judge it as low risk and high benefit; when they dislike it, they flip both judgments-an inverse link called the affect heuristic (Finucane, Alhakami, Slovic & Johnson, 2000). The label can be a useful first pass, but it can also hide trade-offs. A two-step habit helps: name one risk in a "liked" option and one benefit in a "disliked" one before deciding.

Probabilities Our Intuition Misreads

Our gut is great with stories and shaky with chance.

- Base Rates First

Imagine a disease that affects 1% of people and a test that's "95% accurate." A positive test sounds decisive-until you do the frequency version. Out of 1,000 people, about 10 have the disease; the test catches most of them (say 9). But among the 990 healthy people, 5% will still test positive-about 50 false alarms. So roughly 59 positives include only 9 real cases. The base rate explains why most positives aren't sick in this setting. Thinking in counts ("out of 1,000") reliably improves reasoning (Gigerenzer & Hoffrage, 1995).

- Small Samples, Loud Stories

Flip a coin six times and get HHHTTT. It feels "balanced." Get HHHHHT and it feels rigged. Both are equally likely for a fair coin; small strings are just noisy. We routinely over-read tiny streaks-a habit Tversky & Kahneman dubbed belief in the law of small numbers (1971). In sports, the "hot hand" was once thought to be pure illusion; later work finds there can be small real effects, but our eyes still exaggerate them (Gilovich, Vallone & Tversky, 1985; Miller & Sanjurjo, 2018). The safe move is to treat small samples as suggestive, not conclusive.

- The Conjunction Trap

Meet Linda: bright, outspoken, and concerned with social justice. Which is more likely? 1) Linda is a bank teller. 2) Linda is a bank teller and active in the feminist movement. Many people pick option 2 because the description fits the stereotype. But a specific story cannot be more likely than a broader one that includes it. That's the conjunction fallacy (Tversky & Kahneman, 1983). The fix is to zoom out: compare the simple category to the detailed one and ask which contains the other.

CHAPTER 2

Judgment: Framing, Losses & Social Gravity

Numbers come with handles, losses loom larger than gains, and we reason in tribes. This chapter maps the social and emotional forces that tilt judgment - from anchoring to confirmation bias, stress to sunk costs.

WHAT YOUR CHILD WILL BE ABLE TO DO

- Explain anchoring and framing - and why 'lives saved' versus 'lives lost' flips preferences
- Describe loss aversion, the endowment effect and sunk-cost reasoning
- Distinguish real patterns from regression, post hoc traps and illusory correlation
- Understand how stress dims the prefrontal cortex and how tribalism makes beliefs sticky

Numbers, Risk, and Framing

Numbers are not neutral-they come with handles.

- Anchors Set the Stage

High listing prices don't just tempt sellers; they pull buyers and even experts. In a classic field-style study, real-estate professionals toured the same house but were given different listing prices. Their appraisals shifted toward the anchor-even when they insisted they weren't influenced (Northcraft & Neale, 1987). Judges aren't immune either: a random "dice" suggestion for years in prison nudged sentencing upward (Englich, Mussweiler & Strack, 2006). A small guardrail helps-write your own range using base rates before you see anyone else's number.

- Frames Change Choices

Say a program will "save 200 lives," and most people prefer the sure thing. Say it will "result in 400 deaths," and many switch to the risky option-even though both describe the same expected outcome (Tversky & Kahneman, 1981). Medicine shows the same move: a surgery framed as "90% survival" feels safer than "10% mortality." When you feel your gut flip on wording, restate it as absolute frequencies ("out of 1,000 patients...").

- Ratios and Scope

If a jar has 9 winning red tickets out of 100 and another has 1 red out of 10, many people pick the first jar because 9 feels bigger-even though 1/10 is the better chance (Denes-Raj & Epstein, 1994). That's ratio bias. And when problems scale up, our feelings can go flat: people report similar willingness to pay to save 2,000 birds or 200,000 birds from oil ponds (Desvousges et al., 1992). To stay oriented, compare to a baseline denominator ("per 100k") and check whether the total actually matters to your decision.

Losses, Time, and Sunk Costs

Some tilts in judgment are so consistent you can feel them in your bones.

- When Losses Loom

Picture two gambles: one offers a sure \$50 and the other a 50/50 shot at \$100. Many people pocket the sure thing. Flip the same numbers into losses—a sure loss of \$50 versus a 50/50 chance to lose \$100—and preferences often reverse. That signature flip is loss aversion in action (Kahneman & Tversky, 1979). We lean toward the certain gain and toward the risky way to dodge a sure loss. You can feel it in small moments too: deleting a draft hurts more than the same draft ever "felt good" to write.

- Ownership and Staying Put

Hand people a coffee mug to "own" and, minutes later, they demand more money to sell it than they'd pay to buy the same mug. That's the endowment effect (Kahneman, Knetsch & Thaler, 1990). The label "mine" changes value. We also stick with the status quo even when alternatives are better—status quo bias (Samuelson & Zeckhauser, 1988). Real-world defaults show how strong this is: countries with opt-out organ donation have dramatically higher consent rates than opt-in countries (Johnson & Goldstein, 2003). Defaults aren't neutral; they're quiet recommendations.

- Time, Plans, and Sunk Costs

Ask people today: \$10 now or \$12 tomorrow? Many grab the ten. Ask \$10 in 52 weeks or \$12 in 53 weeks and most pick the twelve. That preference reversal reveals hyperbolic discounting—we discount the near future steeply (Ainslie, 1975; Laibson, 1997). We also tell rosy stories about timelines—the planning fallacy—then run late (Buehler, Griffin & Ross, 1994). And once we've paid, we feel pressure to keep going: the sunk cost effect (Arkes & Blumer, 1985). People sit through a bad movie they bought tickets for, or pour hours into a doomed project because "we've come this far." A kinder move is to set a "kill switch" rule in advance and to forecast using an outside view: what happened to projects like this?

Patterns, Causes, and Coincidence

Brains love stories about cause and effect.

- Regression Isn't Retribution

Extreme performances drift back toward average on their own. Praise often follows great results; the next attempt is usually less extreme, so it looks like praise "made it worse." Criticism often follows terrible results; the next attempt is usually better, so punishment "worked." In reality the statistics are doing most of the work. Kahneman tells a story from flight-training where instructors were convinced of punishment's power until they saw the regression pattern laid out (Kahneman, 2011). The practical nudge: judge methods over many cycles, not one swing up or down.

- Sticky Pairings and Post Hoc Traps

Striking pairs glue together in memory even when they're unrelated. Early studies showed clinicians "saw" links between certain test signs and diagnoses that weren't there (illusory correlation; Chapman & Chapman). In everyday life, people often believe joint pain tracks rainy days. When researchers had patients and weather stations keep careful logs, the association shrank or vanished (Redelmeier & Tversky, 1996). And beware the post hoc habit: if B follows A, your mind tags A as the cause. Without a comparison group, many "cures" will seem to work simply because symptoms were about to improve anyway.

- When Totals Lie

In Simpson's paradox, subgroup trends can reverse in the aggregate because the sizes of the subgroups differ. UC Berkeley's 1973 admissions looked biased against women overall, but department-level data suggested women tended to apply to more competitive departments; within many departments, admit rates for women were as high or higher (Bickel, Hammel & O'Connell, 1975). The practical habit is simple: when the stakes are real, check the slices before trusting the soup.

Emotion, Stress, and Choice (Sapolsky focus)

Stress isn't just a feeling; it's a body-wide mode switch.

- Threat Mode

When something feels urgent, your body mobilizes: heart rate up, breathing quick, muscles ready. The HPA axis releases cortisol while other systems release adrenaline and noradrenaline. Attention tunnels onto the possible threat; the brain reaches for practiced scripts. That's the right trade when seconds matter. It's the wrong trade for budgeting, hiring, or ethics. A simple guardrail: label the mode-"I'm in threat mode"-and, if stakes are high and timing allows, create a pause before committing.

- Prefrontal on a Dimmer

Under stress, the prefrontal cortex-the part that juggles rules, resists impulses, and switches strategies-can lose fine control (Arnsten, 2009). You feel more certain and faster, but also less flexible. Lab stressors like a cold-pressor task or time pressure nudge people toward habit-friendly, high-salience choices and away from slow comparisons (see Sapolsky, 2017). That's why "decide when calm" works: you're letting the prefrontal system do its job before the dimmer turns down.

- Chronic Load, Chronic Tilt

Long stress isn't just a lot of short stress. Baselines move: sleep erodes, patience shortens, short-term relief starts winning. McEwen calls the wear-and-tear allostatic load. The fix isn't heroics; it's context design. Schedule hard choices for calm hours, write if-then rules in advance (premade defaults and "kill switches"), and keep the basics-sleep, daylight, movement, and supportive people-because biology sets the stage for judgment.

Social Minds, Tribal Minds

We don't reason in a vacuum; we reason in tribes.

- Why Beliefs Get Sticky

Once you've told yourself a tidy story, new facts get sorted by that plot. In the classic death-penalty study, participants read the same mixed-quality evidence; both pro- and anti- groups judged supporting studies as higher quality and left more entrenched (Lord, Ross & Lepper, 1979). That's confirmation bias with a dash of motivated reasoning-reasoning recruited to defend identity (Kunda, 1990). A tiny counter-ritual: write one result that would change your mind before you read the evidence.

- Groups as Gravity

Social pressure is subtle until it isn't. Asch showed that many people will report the wrong line length if everyone else does first. Milgram showed that authority cues can push ordinary people much farther than they expect. And when like-minded people talk, discussions often drift toward more extreme positions-polarization (Moscovici & Zavalloni, 1969). These aren't just lab curiosities; you see them in meetings where a confident early comment anchors the room.

- Dissonance and Identity

When actions and beliefs clash, it feels uncomfortable-cognitive dissonance (Festinger, 1957). A small nudge can make people later report they enjoyed a dull task if they were paid \$1 to tell someone it was fun (Festinger & Carlsmith, 1959): the cheap lie conflicts with "I'm honest," so the mind edits the belief to reduce the tension. Identity matters too: we bend stories to protect who we think we are. Helpful habit: steelman the smartest opposing view, then state what would change your mind in one sentence.

CHAPTER 3

The Defense Toolkit: Debiasing & Rhetoric

Awareness is not enough - learners build the toolkit. From premortems and calibration to recognizing ad hominem, propaganda patterns and manipulated charts, this chapter turns insight into skill.

WHAT YOUR CHILD WILL BE ABLE TO DO

- Apply the debiasing habits: make dissent easy, run premortems, use reference classes, track and calibrate
- Name the rhetorical fallacies: ad hominem, tu quoque, red herring, straw man, slippery slopes and circular reasoning
- Recognize propaganda patterns - fear, identity, repetition and manufactured consensus - and practice inoculation
- Spot manipulated data: truncated axes, missing denominators and bad encodings
- Evaluate sources by incentives and verify with people, paper and pixels

The Debiasing Toolbox

There isn't one magic fix for bias-there's a handful of small, sturdy habits.

- Make Dissent Easy

If you only hunt for matching facts, you'll find them. A quick counter is to consider the opposite: before you evaluate, write one pattern of evidence that would point the other way (Lord, Lepper & Preston, 1984). You can formalize this with a "red team" or a role swap-steelman the best opposing case, then switch back. The point isn't to be contrarian; it's to give the truth two doors to enter.

- Premortems and Reference Classes

A premortem starts by assuming the plan already failed. Name specific reasons, then fix what you can now (Klein, 2007). For timelines and budgets, don't trust the glow from inside the plan; step outside to a reference class-what happened to projects like this? (Lovallo & Kahneman, 2003; Flyvbjerg et al.). The outside view feels colder and is usually kinder to future you.

- Track and Calibrate

A decision journal is a one-minute note: date, context, options, your prediction, and your confidence. Later you check: did 70%-confidence bets come true roughly 70% of the time? That's calibration (Lichtenstein & Fischhoff, 1977). A simple Brier score (squared-error measure; Brier, 1950) keeps you honest. Over time, you'll feel which 60% is soft and which 60% is solid.

Rhetorical Fallacies: Persuasion Without Proof

Some arguments win by tugging your social instincts rather than presenting good reasons.

- Attacks and Distractions

You'll hear, "Don't trust her argument-she once failed a class." That's ad hominem: attacking the person to avoid the claim (Walton). A cousin is tu quoque-"you're a hypocrite, so you're wrong"-which shifts attention from the point to the speaker (Aikin & Talisse). And a red herring/whataboutism drags in a side issue: "What about corruption in City B?" when we're judging a policy in City A. A sturdy counter is to name the move and restate the original claim in its strongest form before evaluating it.

- Misframing the Choice

A straw man replaces someone's actual position with a weaker version that's easy to knock down. A false dilemma pretends there are only two options ("either security or freedom") when more exist. The motte-and-bailey toggles between a bold, hard-to-defend claim (the bailey) and a safe, vague one (the motte) whenever pressure arrives. Your antidotes: quote the exact claim, list real alternatives, and ask the speaker to stick with one clear thesis.

- Slopes, Authorities, and Circles

A slippery slope says one step guarantees disaster without showing the mechanism. Appeal to authority is only fallacious when the cited person lacks relevant expertise or when their status substitutes for evidence; quoting a genuine expert with data is fine (Cialdini; Walton). Circular reasoning (begging the question) assumes what it must prove ("This policy is illegal because it's unlawful"). The fix is to ask for links in the chain, relevant credentials plus evidence, and an independent reason for the conclusion.

Propaganda Patterns & Manipulation Techniques

Some persuasion doesn't try to win a fair argument-it tries to tilt the field.

- Fear, Identity, and Repetition

Repetition is a volume knob for believability. Hear a claim enough times and it starts to feel true-the illusory truth effect-even when you knew it was shaky at first (Hasher, Goldstein & Toppino, 1977; Fazio et al., 2015). Add identity cues-flags, slogans, insider language-and the message rides your "us vs. them" circuitry (Cialdini). A simple habit helps: write one fact check source before you share; ask, "Where did this number first appear?" Familiar isn't the same as verified.

- Firehose & Manufactured Consensus

The firehose of falsehood floods feeds with lots of claims, across many channels, very fast, without sticking to one story (Paul & Matthews, RAND, 2016). The aim is not to persuade with one strong argument; it's to exhaust attention and make detailed rebuttals look slow. A related move is astroturfing-phony grassroots accounts or "letters from the public"-to fake consensus. Counters: trace to the origin (first source), check independence (are outlets copying one post?), and compare against baselines (what would typical data look like?).

- Inoculation and Prebunking

Minds can get "vaccinated" against tricks. Inoculation gives a tiny dose of the misleading tactic plus a quick refutation, so later you recognize the move and resist it (McGuire, 1964). Modern prebunking does this with short examples and games (Roozenbeek & van der Linden, 2019). You can prebunk yourself: name the technique ("loaded language," "fake expert," "false balance"), state why it misleads, and write the fact pattern you'd expect if the claim were true.

Spotting Manipulated Data & Charts

Tiny design choices-where a y-axis starts, which dates you include, whether you show per-person or totals-can swing the story without changing a single number.

- Scales and Baselines

If a bar chart's y-axis starts at 40 instead of 0, small differences look huge. Line charts can mislead by picking a convenient start date that hides earlier context. Cumulative curves always rise; sometimes a daily or per-week view tells the truer story. Ask: What happens if the axis starts at zero? What if I switch to per-capita instead of totals? (Tufte; Cairo).

- Denominators and Slices

"Up 300%!" sounds dramatic-until you see it was 1 to 4 cases. Percentages need counts, and rates need the denominator ("per 100k"). Aggregates can hide reversals: by subgroup, a pattern may flip-Simpson's paradox-because the mix of cases changed. When the stakes are real, check multiple slices and report the denominator alongside the claim.

- Design Tricks and Bad Encodings

Our eyes read position and length well, but we're worse with area and angle (Cleveland & McGill, 1984). That's why 3-D pies and exploded slices confuse more than they clarify. Dual y-axes can make unrelated lines seem to move together. Heavy color gradients suggest precision that isn't there. Prefer simple bars/lines, one scale per chart, and annotations that say exactly what changed and when (Cairo; Tufte).

Media Literacy - Sources, Incentives, and Fact-Checking

Media literacy - sources, incentives, and fact-checking not all "news" is built the same.

- Sources and Incentives

Ask what the outlet is for. Is it subscriber-funded, ad-driven, or advocacy-backed? Does it list a masthead and publish corrections? "Native ads" and "sponsored content" often look like articles but are paid placements. Press releases through newswires can show up as "news" with minimal editing-classic churnalism. A quick tell: lots of brand quotes, no independent voices, no linked data. Your first move is to read sideways to learn who you're dealing with.

- Verification: People, Paper, Pixels

Work from the outside in. For people, check independent bios and past work. For paper, follow the link trail to primary docs: datasets, court filings, methods sections. For pixels, watch for old images recaptioned as new, and do a quick reverse-image or context check (date, place, original caption). Open-source sleuths routinely spot reused footage by matching skylines and weather. If a chart is viral, ask for the denominator, the time window, and who made it.

- Corrections, Uncertainty, and Preprints

Science moves in drafts. Preprints share results before peer review; they can be useful and also wrong. Even peer-reviewed work gets retracted or revised. Strong outlets show hedging language ("may," "estimate," confidence intervals) and keep a visible corrections log. Press releases often sand the edges-remember the "chocolate weight-loss study" hoax that some outlets amplified (Bohannon, 2015)? Treat bold claims like bright lights: slow down and look for the circuit.

CHAPTER 4

The Digital Battlefield & The Capstone

The final chapter confronts the modern attention economy: algorithmic feeds, deepfakes, social engineering and conspiracy thinking. It closes with the capstone - decision hygiene and a personal playbook for life.

WHAT YOUR CHILD WILL BE ABLE TO DO

- Explain what feeds optimize and how filter bubbles shape exposure
- Understand deepfakes and cheapfakes - and the quick checks that catch them
- Recognize social engineering hooks and dark patterns, and dodge them
- Understand why conspiracy thinking attracts - and how to have bridge-building conversations
- Read risk honestly: absolute vs. relative numbers, correlation vs. causation, confidence intervals
- Build a personal decision playbook: checklists, premortems, noise reduction and cool-head guardrails

Algorithms, Feeds, and Filter Bubbles

Your feed isn't a timeline; it's a forecast.

- What Feeds Optimize

Feeds predict what keeps you there: clicks, dwell time, comments, reshares. Those are proxies for value, not value itself. Emotional, novel, and identity-flavored posts tend to score well, so they float. A/B tests and rapid rollouts let platforms nudge behavior at scale. Rule of thumb: if it spikes arousal or flattery, your feed may be rewarding the feeling, not the fact.

- Bubbles, Echoes, and Exposure

People cluster with their own-homophily-and algorithms can amplify that, creating echoes. But the "filter bubble" story is not one-size-fits-all: studies find both narrowing and meaningful cross-cutting exposure depending on platform, design, and user habits. Either way, the safe bet is to assume drift toward comfort unless you steer on purpose.

- Reclaim Your Feed

Treat your feed like a garden. Plant diverse follows (credible sources with different priors). Add lists/RSS for must-see reporting so algorithms can't bury it. Flip the switches: turn off autoplay, pause or clear watch/history, set time limits, and use "see less" on outrage bait. When a claim matters, go upstream to the first source before you share.

Deepfakes & Synthetic Media - Detection and Context

Deepfakes & synthetic media - detection and context a convincing fake doesn't need perfect pixels; it needs a believable story and a rushed viewer.

- What Counts as a Deepfake?

"Deepfake" gets used loosely. At one end are AI-generated images, video, and audio-face swaps, voice clones, fully synthesized scenes from diffusion or other models. At the other end are cheapfakes-spliced-up or slowed video, deceptive crops, swapped captions, misleading thumbnails. Both ride the same psychology: vivid + familiar + repeated feels true. Your first defense is to label what you're seeing: Is this likely AI-generated, or just edited? Different tools catch different tricks.

- Quick Checks That Catch a Lot

Slow it down. Watch mouth-audio sync, eyelines, and hand/object interactions frame by frame. Scan for impossible details (warped text, jewelry that jumps sides, inconsistent shadows/reflections). Do a reverse-image on key frames and search a quote to see if it exists elsewhere first. For audio, check for robotic prosody or breath/noise patterns that repeat. Then ask context questions: who filmed this, where, and when? Was it first posted by a known source, or by an anonymous account with recent creation and recycled content?

- Provenance, Labels, and When to Escalate

Some platforms and publishers now attach content credentials (provenance metadata that can show capture, edits, and tools). Watermarking systems exist but aren't yet universal, so treat labels as signals, not guarantees. For high-stakes clips-finance, safety, elections-escalate: look for newsroom debunks, fact-checkers, or dedicated OSINT analysts. Your personal rule of thumb: don't amplify until you've verified upstream and can name the original source, date, and place.

Scams, Social Engineering, and Dark Patterns

Scams don't beat your iq; they ride your instincts.

- How People Get Hooked

Social engineers use the same levers good marketers do: authority ("I'm from IT"), urgency ("respond in 5 minutes"), scarcity ("last slot"), liking/reciprocity ("I did you a favor"), and consistency ("you always help quickly"). Real plays: a "CEO" calls finance for an urgent wire (sometimes with an AI-cloned voice), a courier text links to a "missed delivery," or your authenticator floods with push requests (MFA fatigue). Name the lever out loud; it breaks the spell.

- The Social Engineer's Toolkit

Common moves: phishing/spearphishing (broad vs targeted), pretexting (invented roles), baiting (USB drops), quishing (QR-phishing), SIM-swaps (number hijack), and account-recovery abuse. Defenses that travel well: verify out-of-band (call a known number), never share one-time codes, prefer passkeys/hardware keys to SMS codes, turn on number-matching for push approvals, and keep least privilege on finance accounts.

- Dark Patterns and How to Dodge Them

A roach motel is easy to enter, hard to exit-subscriptions that hide "cancel." Confirmshaming scolds you into clicking ("No, I hate saving money"). Hidden costs appear at checkout; trick questions flip opt-outs; forced continuity charges after trials. Your counters: hunt for the neutral choice ("No thanks"), read labels around buttons, uncheck pre-selected boxes, and use virtual cards with spend caps for trials. If a flow fights you, that's information about the company.

Conspiracy Thinking - Why It Hooks & How to Engage

Conspiracy thinking - why it hooks & how to engage when life feels uncertain, the brain goes hunting for patterns and intentions.

- Why Conspiracies Attract

After a shocking event, tidy stories with villains feel more proportional than "a small error snowballed." That's proportionality bias. Add the brain's hair-trigger agency detection-seeing intention in noise-and a dry patch of uncertainty, and conspiracies can feel like relief (Whitson & Galinsky, 2008; Douglas et al., 2017). None of this means people are foolish; it means the mind prefers meaning to randomness, especially when control feels low (van Prooijen & Acker, 2015).

- How They Spread Online

False news often spreads faster and farther than true because it's novel and emotional (Vosoughi, Roy & Aral, 2018). Engagement-tuned feeds reward hot takes; JAQing ("just asking questions") keeps moving the goalposts; and like-minded clusters supply instant validation. A quick counter is structural: follow upstream to primary sources, compare with base rates, and look for falsifiable claims-ones that stake out what would count as being wrong.

- Bridge-Building Conversations

Lead with curiosity, not combat. Try a mini motivational-interviewing arc: ask open questions, reflect back ("So the timelines don't add up for you"), affirm shared values ("we both care about being accurate"), then offer a clear, sourced alternative. Use a fact-sandwich (fact myth fact), ask a confidence question ("what makes it a 7/10?"), and agree on one disconfirming test you'll both accept. Avoid shaming; you're protecting the relationship so truth has room to land.

Risk, Numbers, and What They Really Mean

Big headlines love big percentages.

- Absolute vs. Relative (and Base Rates)

A vitamin "halves" your risk. If the baseline is 2 in 10,000, halving to 1 in 10,000 is both true and tiny. That difference is the absolute risk (one fewer in 10,000), while the headline used relative risk (50%). When you see a scary percent, ask for the base rate and write it as "X in 1000." Doctors and patients make better calls with natural frequencies (Gigerenzer).

- Correlation, Causation, and Study Design

If ice cream sales and drownings rise together, summer-not sundaes-is the cause. Correlation is a clue; causation needs a mechanism and a design that rules out confounders. Randomized trials help; observational studies are useful but vulnerable to hidden differences. Whenever possible, ask: what else moved with the treatment? Is there a plausible pathway? Do multiple designs point the same way?

- p-Values, Intervals, and Forking Paths

A p-value below 0.05 means the observed pattern would be uncommon if there were no effect-it does not prove a big or important effect. Look for the effect size and confidence interval. Results shrink on replication for good reason: flexible choices (p-hacking, multiple comparisons, peeking) inflate surprises (Ioannidis, 2005; Gelman & Loken). Strong signals survive preregistration and fresh data.

Capstone - Decision Hygiene & Personal Playbooks

Capstone - decision hygiene & personal playbooks this is where the parts click together.

- Decision Hygiene in One Page

Treat decisions like labs treat measurements. Break the task into independent judgments (each rater scores criteria alone), use structured scales with examples, and aggregate before discussion to reduce noise (Kahneman, Sibony & Sunstein, 2021). Add a quick base-rate check and write a one-line change-my-mind condition. It feels slower; it's faster than cleaning up avoidable mistakes later.

- Build Your Personal Playbook

Keep a decision journal: date, options, your prediction, confidence, and why. Do a premortem for big bets ("It failed-why?"), then fix what you can now (Klein, 2007). Standardize recurring moves with a short checklist (Gawande, 2009). Once a month, calibrate: did your 70% calls land near 70%? Tighten or loosen your gut accordingly (Tetlock & Gardner, 2015; Lichtenstein & Fischhoff, 1977).

- Team Guardrails and Cool Heads

Make dissent cheap: assign a red team or rotate the devil's advocate. Collect independent estimates before meetings and reveal them at once. Use a cool-off window for costly, irreversible moves; sleep beats sprint when stakes are high. For pricing and timelines, start from reference classes-what happened to projects like this?-then adjust with specifics.

How mastery is tested

Every unit ends with practice questions in three formats (109 multiple choice - 64 fill in the blank - 47 order the words). Wrong answers are automatically recycled in later sessions until the learner proves mastery. Try one from each chapter - click an answer to test yourself:

109 Multiple-Choice

64 Fill-in-the-Blank

47 Word-Order

Wrong answers are recycled automatically in later sessions until the learner proves mastery at 90% - the course is designed so that every child gets there.

Supporting your learner at home

This course takes learners on a guided tour of Mental Debugging - every topic paired with a memorable story and hands-on practice. A few habits make the learning stick:

- Preview and prime: before a chapter, show your child the unit names here and ask which one sounds most interesting. Start there.
- Ask for the story: every unit is built around a memorable analogy or real-world story. Ask your child to retell it - the story is the hook the lesson hangs on.
- Labs are conversations: the practice units turn knowledge into applied challenges. Try the same question at dinner and let your child be the expert.
- Look for the practical lessons: the ideas in each unit are designed to be tried at home - run one experiment together this week.
- Trust spaced repetition: wrong answers are not failures, they are queue items. The platform resurfaces them until mastery, so encourage retries without shame.
- Keep it short and steady: units are designed for short sessions. Fifteen focused minutes beat an hour of fatigue.
- Celebrate the proof: each unit completed at mastery mints a collectible certificate - let the collection be the motivation.

Key vocabulary

Terms highlighted inside the lessons, gathered for quick reference. If your child can define each one, they have mastered that lesson's core idea.

phoneme restoration	blind spot	anchoring
gorilla suit	shifts	cocktail party effect
inattentional blindness	change blindness	door study
attentional blink	multitasking	task switching
switch costs	load	attentional residue
Ponzo	hollow-mask	McGurk
sound-induced flash	hit	smashed
bed, pillow, dream, nap	sleep	source monitoring
source	forget the base rates	feel
ease	like	low risk and high benefit
dislike	affect heuristic	95% accurate
1,000	990	HHHTTT
HHHHHT	law of small numbers	suggestive
Linda	bank teller	bank teller and active in the feminist movement
cannot	conjunction fallacy	contains
save 200 lives	result in 400 deaths	absolute frequencies
9 winning red tickets out of 100	1 red out of 10	ratio bias
2,000	200,000	baseline denominator
\$50	\$100	loss of \$50
loss aversion	gain	risky

Decode reality. Master the future.

almanac.academy