




SO THEY LIED TO YOU

A guide through one research project — from the stakeholder's request to a changed decision. Eight stages, and at each one a lie of its own is waiting for you: from users, from data, from AI, from the business, and from your own head. I'll show you exactly where — on my own studies and my own mistakes.



AUTHOR

Nina Somkina · Founder, Institute for Human Reasoning ·
10+ years in research · mixed-methods

FORMAT

Research playbook · EN · first edition, 2026
Don't outsource the thinking that matters.



Nina Somkina

Founder at Institute for Human Reasoning
10+ years in Research · Mixed-methods ·
Consultant · Mentor · Pre-MBA

Book a consultation

Researcher, consultant, and educator. Ten years in research — and ten years of collecting the lies I fell for myself.

I work in mixed-methods — where one question needs both a conversation and numbers. I specialize in hard research: the answer isn't on the surface, the sources contradict each other, the decision is expensive. Strategic segmentations, churn, launches in new markets, the future of work — for startups and large businesses in subscription B2C, banking, and B2B.

I consult organizations on how to research well and fast — and how to make decisions on the data they already have. Because companies usually have more data than decisions made on it.

I teach founders, designers, product managers, writers, marketers, and managers critical thinking and how to work with new tools — including AI in research. I teach at Projector Institute and Apollo IT School, and mentor beginner researchers and teams with no researcher on staff.

For me, good research is not an end in itself. What matters is that a decision gets made on the data — one that changes tactics or strategy. A study after which nothing changed is wasted time and money, no matter how good the report was.

Founder of the Institute for Human Reasoning — a research and education organization about human thinking and decisions in a technological world. This playbook is its first product.

Everything you're about to read really happened to me.
Including the mistakes. Especially the mistakes.

The route of this playbook

One research project, walked from start to finish. If you have a live study right now — open the stage you're standing on: the page number is on the right. If not yet — read straight through: the route is the plan of your first research project.

Instead of an introduction

What's the best way to kill a person · what cognitive biases are, System 1 and System 2 · ≈180 ways to lie to yourself · what I call a "lie" · how to read the tags · three reading routes

05

Honest numbers

4 000 → 28 · 10 000 → 0 → 0 · 260 and 30 · 126 → 0 opens → decisions · 76 → 58 → 7 · 35% → 7% · ½ "active" · 2 weeks → 1 hour · 30–40%

09

1 Request

Translating the request · the client test · three types of research and honest timelines · roadmap and prioritizing 42 requests

11

2 Plan & People

Two worlds of recruiting · how many people is enough (no magic) · AI agents as scale · democratization and continuous discovery

16

3 Questions

Survey, in-depth interview, usability test: three ways to order a lie — and how to catch them

23

4 Interview

Three reasons for lying without intent · the CPO and the liar · "the teacher in the classroom" · who finds it easier to confess to AI · a team of interviewers

29

5 Analysis

The AI verification ritual · how AI lies without context · Japanese curses that don't exist · segment and persona · "too expensive"

34

6 Report

0 opens ≠ 0 impact · evidence shelves · evidence ladder · the insight formula · clips, screenshots, and empathy · formats

42

7 Read-out

"We already knew that" · when C-level shock is a diagnosis for the researcher · a 2–4 day pause · the Socratic technique · workshop

48

8 Decision

Small signals · decision check · count decisions, not interviews · the case of AI that added no value · the bank case finale

52

AI on the route

Discover vs Scale · who is responsible for what · 4 AI mistakes and prompts against them · ChatGPT vs Claude vs Gemini

57 Quality check

Rigor and impact · 6 quality dimensions · reliability ≠ validity · triangulation · insight hub vs repository

61 Instead of a conclusion

Honest numbers with conclusions · what to read · what to use · template pack · glossary · references

68

THE RUNNING CASE

One real research project runs through all eight stages — the **bank case**: a bank's premium segment, the client is the CEO, 4 000 contacts, 28 interviews. Look for the blocks with the gold line: they show where the study stands at each stage and what its decision state is. The case is anonymized — we don't name the bank; the numbers and the course of events are real.

Full contents

All subsections of the route — for coming back: when the book has been read and you need one specific place.

Front matter	02	Stage 6 · Report	42
Who leads this route	02	6.1 126 → 0 opens ≠ 0 impact	43
The route of this playbook	03	6.2 Shelves the business can tell apart	44
Full contents	04	6.3 What an insight is — and the formula the business hears	45
What's the best way to kill a person?	05	6.4 The quote the executive recognized	46
What I call a "lie"	05	6.5 Formats that work instead of 120 slides	46
System 1, System 2, and biases	06	Stage 7 · Read-out	48
≈180 ways to lie to yourself	07	7.1 "That's obvious" is not an attack	49
How to read the tags · three reading routes	08	7.2 One study — three presentations	50
Honest numbers	09	Stage 8 · Decision	52
Route map	10	8.1 Small signals: fooling you — or swans	53
Stage 1 · Request	11	8.2 What you can honestly measure — and what you can't	54
1.1 "Find out why they leave" is not a question	12	Layer · AI on the route	57
1.2 Define the type — and don't lie to yourself about timelines	13	The working model: human → AI → human → quant	58
Stage 2 · Plan & People	16	Eight stages — who does what	58
2.1 Two worlds: mass B2C and premium banking	17	Four AI mistakes — and a prompt against each	59
2.2 AI agents: when you need many conversations	18	ChatGPT, Claude, Gemini: who is who	60
2.3 Your database lies too	20	Layer · Quality check	61
2.4 One researcher can't carry it all — and doesn't have to	20	How quality is actually judged	62
Stage 3 · Questions	23	Reliability ≠ validity	64
3.0 Three methods — three ways to order a lie	24	Triangulation: the cheapest way to raise validity	65
3.1 In-depth interview: seven questions	25	Insight hub and repository are not the same thing	66
3.2 Survey: lies in the scale and the order	26	Instead of a conclusion	68
3.3 Usability testing: leading prompts	27	Honest numbers — once more. Now with conclusions	68
Stage 4 · Interview	29	They will keep lying to you	69
4.1 Three reasons for lying without intent	30	Twelve books	70
4.2 Who finds it easier to confess to AI than to a human	31	Articles, guides, and communities	71
4.3 Several interviewers — calibrate the hand	31	Tools by research cycle	72
4.4 Interview ethics: consent, recording, data	32	Template pack	75
Stage 5 · Analysis	34	Appendices	76
5.1 AI's first answer is not analysis	35	A Glossary: terms of the route	76
5.2 Three Japanese curses	36	B References	78
5.3 Segment and persona: different tools	38		
5.4 Segmentation is judgment. But not just any judgment	39		
"Too expensive" — translate it, don't read it literally	40		

What's the best way to kill a person?

I ask this question at the start of every cohort. The audience is smart adults: product managers, designers, marketers. Within two minutes the chat has dozens of confident answers. With details. With reasoning.

Then I ask the second question: “And who of you has ever actually killed anyone?”

Silence.

No one has the experience. Yet everyone just confidently lied — to themselves and to me, with no ill intent at all: the brain simply knows how to generate convincing answers about things it has never lived through. That is the environment a researcher works in.

Your stakeholder will confidently bring you a broken question. Respondents will confidently describe a future that will never happen. The panel will slip in people from the wrong country. AI will confidently write a conclusion that isn't in the data. And your own brain will quietly gather evidence for whatever is convenient to believe. All of this — within one single study.

Research is not a way to get at the truth. It's a way to catch the lie before it costs money.

One more thing before we set off. This playbook is the first product of the Institute for Human Reasoning, and its main idea fits in one line: **don't outsource the thinking that matters**. Answers are cheap today — analytics, panels, respondents, and AI all provide them. Only one thing stays expensive: the judgment of the person who checks those answers and takes responsibility for them. The eight stages ahead train exactly that ability.

What I call a “lie”

Let's agree on the word right away, because it's in the title. A “lie” here is any convincing answer, number, or conclusion you shouldn't take literally. Sometimes it's deliberate deception — like the respondent who professionally gamed the service for months (stage 4). More often it isn't: politeness, fear of looking stupid, a question about an imaginary future, a dashboard that runs a month late, a skewed sample, an AI hallucination, a brain that filled in a missing answer. I keep “lie” in the stage titles — it holds the route together. Inside, I'll name the mechanism more precisely, because each one is treated differently.

How the route is built

Not by methods — by the route. One study from start to finish: request → plan → questions → interview → analysis → report → read-out → decision. At each stage: the lie that lives there; a story where I fell for it; the bias that serves it inside your head; a tool you can take away tomorrow.

And one running study. My real case will go through all eight stages: a bank, the premium segment, a dip the CEO could see and the managers couldn't explain. 4 000 contacts, 28 interviews, a report, a workshop — and an honest finale. Each stage has a “Bank case: where we are now” block.

But before we set off — three concepts on the next two pages. Without them, “stage bias” makes no sense.

Why the brain lies to itself: System 1, System 2, and biases

In “Thinking, Fast and Slow”, Daniel Kahneman described two modes of thinking (Kahneman, 2011). The split is simplified, but for a researcher it works: it explains where a confident answer to a question with no data behind it comes from.

SYSTEM 1 · FAST

Automatic, instant, effortless. It recognizes faces, completes “2 + 2”, senses that the person across from you is irritated — and **generates a ready answer to any question**, even one it has no data for. It’s the one that wrote “the best way to kill” into the chat within two minutes. It doesn’t know how to say “I don’t know”.

Example. “A bat and a ball cost 1.10. The bat costs 1.00 more than the ball. How much is the ball?” System 1 instantly says “10 cents” — and gets it wrong: the right answer is 5. Most people, researchers included, don’t check.

SYSTEM 2 · SLOW

Conscious, lazy, energy-hungry. It multiplies 17×24 , compares two plans, checks whether a quote is in the transcript. It switches on only when forced — and **readily accepts what System 1 has already decided**, writing a logical explanation in after the fact.

Example. A stakeholder “knows” why users leave (System 1). Ask him to explain step by step exactly how a person gets to the point of leaving — System 2 switches on, and the explanation stalls at step three. That’s where the research question lives.

What a cognitive bias is

A **cognitive bias** is not a mistake and not stupidity. It’s a **systematic way System 1 cuts a corner**: substituting an easier question for a hard one, taking the most available example as typical, seeking confirmation instead of refutation, clinging to what it has already invested in. The key word is **systematic**: the same error repeats across different people in similar conditions. Which means it can be predicted — and a safeguard can be built into the research design. That’s exactly what every “Stage bias” block does from here on.

IN THE STAKEHOLDER

“I know why they leave — the onboarding is too long.” System 1 substituted an available picture for the answer to a hard question. That’s the **illusion of explanatory depth** (stage 1).

Antidote: “walk me through it step by step”.

IN THE RESPONDENT

“Would you buy it for 9.99? — Yes, of course.” The brain can’t live through an imaginary future, so it invents something plausible. **Anchor + hypotheticals** (stage 3).

Antidote: ask about the last real episode.


IN YOU

Out of 28 transcripts, the quotes that fit your hypothesis get into the report “by themselves”. You didn’t pick them — they just seemed “on topic”. **Confirmation bias** (stage 5).

Antidote: “what in the data contradicts me?”

The respondent, the stakeholder, AI, and you all run on the same System 1 (AI — on a statistical copy of millions of such systems). That’s why a lie in research is almost never intent. It’s a by-product of thinking. And what works against it is not “be more careful” but process: a rule in the guide, a checklist, a verification ritual.

≈180 ways to lie to yourself — and only four reasons



Each dot is one documented cognitive bias; each node is a mechanism that produces several. The diagram follows the structure of the Cognitive Bias Codex (Benson, 2016; visualization by Manoogian, 2016): ~175–188 biases grouped by the four problems the brain “solves” this way.

01 • Too much information

The brain filters — and the filter is biased.

- we notice what’s already in memory or repeats
- the odd and the funny catch the eye more
- we notice changes, not states
- we’re drawn to what confirms our beliefs
- we see other people’s errors better than our own

02 • Not enough meaning

The brain fills the gaps — with the plausible.

- we find stories and patterns even in sparse data
- we fill gaps with stereotypes
- we simplify probabilities and numbers
- we “know” what others think
- we project the present onto the past and the future
- we value the familiar over the new

03 • Need to act fast

The brain picks confidence over accuracy.

- we overestimate our own importance and influence
- we stick with what we’ve already invested in
- we choose the near and the finished over the deferred
- we defend autonomy and status
- we prefer the simple options

04 • What to remember

The brain edits memories — and doesn’t tell you.

- we rewrite memories after the fact
- we drop details and keep generalizations
- we reduce events to key elements
- we remember according to how we experienced it

It’s enough to see the four problems — and to ask one question before every conclusion: **“which of the four did my brain just solve for me?”** Too much data — I filtered something out. Not enough meaning — I filled something in. Need to act fast — I picked the confident over the accurate. Memory — a retelling of a retelling. Later in the playbook: two dozen specific biases — the ones I catch in myself most often.

How to read the tags

MY CASE this happened in my own research — an observation, not a law.

RULE OF THUMB a practical rule that often works for me.

RESEARCH-BACKED there’s a source for this, and I name it.

CHECK IN YOUR CONTEXT check it with your audience before leaning on it.

Bank case — where the running study is now.

Tool — tear it out and use it tomorrow as is.

Hint — a short practical tip from the field.

Stage bias — what inside your head makes the lie convincing.

Takeaways — the stage summary.

Three reading routes

The book is designed as a route, but you can read it three ways — depending on how much time you have and how many studies you have behind you.

If you have 30 minutes. Read “Honest numbers” on the next page, open the stage your live study is standing on, and take its “Takeaways” block. The last half hour, whenever it comes, goes to the Quality check: six dimensions and the 0·1·2 rubric.

If this is your first study. Read straight through, stage by stage: the route is the plan of your first research project. Tear out the “Tool” blocks and use them as is; leave the research-backed sources for the second pass — they’re not going anywhere.

If you have dozens of studies behind you. Follow the tags: my case and “Honest numbers” — to check against your own experience; Quality check and the AI layer — in full; the glossary and references — as a handbook for the team. And argue with me in the margins: half of this book was born exactly that way.

What no vendor will ever show you

Every number comes from a real study of mine, and every one is unpacked inside the playbook. Where a category or numbers were changed for anonymization, the card's passport says so directly — with proportions and conclusions preserved. Under each number is a minimal passport, so you understand what exactly is being compared, without breaking NDA: category, method, N, decision, limitations.

RECRUITMENT

4 000 → 28

The bank's premium clients who left within three months. Support personally reached out to every one. 28 interviews — and a gift with a letter signed by the CEO, which weighed more than the gift.

The running case · bank, premium segment · client: CEO · in-depth interviews · N ≈ 4 000 contacts → 28 interviews · Decision impact: the team changed the framing of the problem and the prioritization of churn reasons · Implementation: stayed with the client · Business outcome: unavailable after handoff / NDA

SEGMENTATION

76 → 58 → 7

76 in-depth interviews, 58 in the analysis, 7 segments confirmed quantitatively. The segments came together at interview 52. Not at 12.

B2C · strategic segmentation · device × content type × plan matrix · 3 researchers · 3 weeks · 18 dropped for quality: "everything's great" with no meat, a drunk respondent, bad audio

FAILURE

10 000 → 10 → 0 → 0

A churn survey to ten thousand unsubscribed users: ten responses, zero interviews. With a \$300 raffle — the same. An AI agent via email with a control group — zero again.

B2C subscription · churned users · client: CEO, the board demanded -5 p.p. of churn quarter over quarter · three recruiting attempts · Decision: recruiting those who left is a separate problem · Limitation: one product

DECISION EVIDENCE

35% → 7%

Marketing was bringing in a third of the category's audience. A seventh stayed. Minus 28 points is a lot of money.

B2C · acquisition profile vs the JTBD of those who stay · survey + analytics · Decision: targeting · Limitation: the exact N and the denominator source didn't survive — read as an order of magnitude

SCALE / COST

260 and 30

260 interviews run by an AI agent on a \$200 recruiting budget — category discovery of coffee drinks, three categories. In parallel, ~30 human interviews with the same base (\$150 of recruiting) gave the same conclusions.

Discovery · coffee drinks, 3 categories · ads → landing page → AI-agent interview · N = 260, quality filter (<10 min, off-topic) · triangulation: ~30 human interviews · budgets are recruiting only, no incentives · Limitation: qualitative data, self-selected via ads · Categories and numbers partially changed with meaning preserved; proportions and conclusions are real

STALE DATA

1/3 "active"

A third of the users analytics showed as active answered: "I don't use the product anymore."

B2C subscription · survey across the "active" base · Limitation: a third of those who responded; N didn't survive

ARTIFACT ADOPTION ≠ IMPACT

126 → 0 opens → decisions

126 interviews. Master table — zero opens. Consequential decisions — yes. Details — NDA. The table failed as an interface. The research didn't.

My first big research project · N = 126 · Artifact adoption: 0 opens · Decision impact: consequential decisions were made · Business outcome / details: NDA

ANALYSIS COST

2 weeks → 1 hour + validation

12 interviews used to mean two weeks of analysis. Now it's one hour of my capacity plus 4.5 days of the team's time for validation.

The economics of analysis · N = 12 · honest full cost = 1 h + 4.5 days

AI CORRECTION · PERSONAL ESTIMATE

30–40%

How much I edit, on average, in a "finished" AI analysis before it becomes true.

MY CASE an estimate from experience, not a measured metric

Eight stages — eight kinds of lies — eight biases

If you have a live study right now — open the stage you're standing on. If not yet — read straight through: the route is the plan of your first research project. Every stage has three permanent blocks: the stage's lie (what's outside), the stage bias (what inside your head makes that lie convincing), and the bank case (where the running study is at that moment).

STAGE	THE STAGE'S LIE — WHAT'S OUTSIDE	BIAS — WHAT'S INSIDE
1	Request "The stakeholder's request is the research question"	Illusion of explanatory depth · framing
2	Plan & People "We'll find respondents in three days"	Planning fallacy · survivorship · self-selection
3	Questions Half of the respondent's lies are ordered by you	Anchoring · order · social desirability
4	Interview The person across from you lies out of politeness — and sometimes professionally	Observer · post-hoc rationalization · the vivid story
5	Analysis AI with a ready conclusion and your brain gathering the evidence	Confirmation · automation · false causality · source confusion
6	Report "The report is the result of the research"	Curse of knowledge · IKEA effect
7	Read-out "That's obvious, we already knew"	Hindsight · status quo
8	Decision "Research will pay for itself" and "one insight — straight into a feature"	Availability · sunk cost · loss aversion · action for action's sake

CROSS-CUTTING LAYER · AI ON THE ROUTE

Who is responsible for what at each of the eight stages: what an agent can take on, and what stays with the human — because capability is not authority.

CROSS-CUTTING LAYER · QUALITY CHECK

How to know your research can be trusted: 6 quality dimensions, the 0·1·2 rubric, reliability and validity in human language, triangulation, insight hub.