

HOW DO WE KNOW?

THE EVIDENCE AUDIT FIELD GUIDE

A practical, printable toolkit for questioning claims, separating observation from inference, spotting hidden variables and changing your mind without losing your bearings.

8 frameworks / 4 worked cases / 12 exercises / reusable worksheets

START HERE

Certainty is not the assignment.

A short orientation for using this guide with intellectual care.

Most claims arrive as ready-made packages: a number, a story about what caused it, and an implied instruction about what to do next. The packaging can be persuasive even when the underlying evidence is thin. Inquiry begins by opening the package and examining the parts separately.

Four useful labels, not four identities

- **Observed report:** what a source says was seen, measured, or recorded. Reports can still be mistaken or manipulated; the label is not a certificate of truth.
- **Interpretation:** a possible meaning, mechanism, motive, or causal explanation for the observation.
- **Testable hypothesis:** a precise, disputable expectation about what a specified check could reveal.
- **Unsupported assertion:** a conclusion too broad or too confident for the supplied information. Unsupported does not necessarily mean disproven.

AN ETHICAL LIMIT

When studying people, do not treat ambiguous evidence as a licence to read minds. Consider context, privacy, consent and unequal access to information. An uncertain conclusion can be useful without being accusatory.

The examples in this guide are **fictional teaching cases** unless explicitly described otherwise. Numerical values are chosen to clarify reasoning, not to report real studies. Your goal is better judgment, not another metric for comparing yourself with strangers.

ASK YOURSELF / What belief would I prefer not to investigate? Why?

PROTOCOL / POCKET REFERENCE

Eight questions before you believe.

Use this checklist for a headline, workplace claim, research paper, or family disagreement.

01 / Define

What, exactly, is the claim? Can I restate it without adding implied certainty?

02 / Observe

Which facts were directly measured or reported, and by whom?

03 / Scope

Does the claim concern this sample, this time period or a broader population?

04 / Compare

What other explanations would predict similar observations?

05 / Audit selection

Who was excluded, did not respond or could not be measured?

06 / Find a counter-test

What plausible result would actually reduce my confidence?

07 / Calibrate

How strong is this evidence given the base rate and study design?

08 / Revise

What conclusion remains after saying what I do not know?

THE SLOW QUESTION

“What would the world look like if my preferred explanation were false?” If the answer is “exactly the same,” the observed pattern is not yet discriminating between explanations.

FRAMEWORK / 01 / OBSERVATION

The raw record is not the story.

The everyday situation

A team message was read at 9:20 a.m. and answered at 4:45 p.m. That timing is an observed event if the messaging platform records are accurate. “The person ignored me on purpose” is a separate interpretation. Even a precise timestamp cannot reveal an inner motive.

Why the obvious conclusion can fail

An observation is always shaped by how it was recorded. Was a read receipt generated accurately? Could somebody else have opened the device? Did we select only one message out of twenty? Precision is useful; it is not the same as completeness.

One practical experiment

Make a two-column note. In the left column, write only externally describable details: what, where, who, when, and the measurement method. In the right, write the story your mind supplied. Try two alternatives that do not require a villain.

A MORE DEFENSIBLE SENTENCE

“I have evidence of a six-hour response delay; I do not have evidence of why it happened.”

ASK YOURSELF / What can I say without pretending to know another person's mind?

Transfer challenge: find a recent headline or assertion where this same reasoning error could plausibly occur. Write one concrete limitation without claiming the entire story is false.

FRAMEWORK / 02 / CORRELATION

Association is a clue, not a verdict.

The everyday situation

In a fictional city, lemonade purchases and pool visits both increase as temperature rises. A plot of the two measurements would show a positive association even though buying lemonade does not send people to the pool. The hidden cause in our toy model is heat.

Why the obvious conclusion can fail

Possible causal explanations for a correlation include A causing B, B causing A, a third factor C causing both, selection bias creating an apparent link, and random coincidence. Several can occur at once. Avoid the opposite error of saying that correlation is worthless: in a sound design, an observed relationship may still support a causal account.

One practical experiment

Sketch three circles labelled A, B and C. Put an arrow $C \rightarrow A$ and another $C \rightarrow B$. Ask which measures would distinguish this story from $A \rightarrow B$. Often the answer involves time ordering, a comparison group, or variation in C.

A MORE DEFENSIBLE SENTENCE

"These patterns co-occur. The direction and cause are still open."

ASK YOURSELF / If I could measure one more variable, which would most improve the inference?

Transfer challenge: find a recent headline or assertion where this same reasoning error could plausibly occur. Write one concrete limitation without claiming the entire story is false.

FRAMEWORK / 03 / SAMPLING

Count the missing people.

The everyday situation

A fictional café sends a survey link only to newsletter subscribers and receives 40 replies. Thirty-two respondents approve of the new menu. It is accurate to say 32 of 40 respondents approved. It is not automatically accurate to say 80% of every customer approves.

Why the obvious conclusion can fail

A sample can be large and still unrepresentative. Selection bias concerns how people enter the study. Nonresponse bias concerns who answers once invited. Measurement bias concerns how questions were asked. These problems have different remedies and can point in different directions.

One practical experiment

Write the target population in one sentence. Then write the actual recruitment process. The difference between these two descriptions is where unearned generalization often begins. Next, ask whether respondents differed from nonrespondents in ways linked to the outcome.

A MORE DEFENSIBLE SENTENCE

"In this responding group, approval was 80%. We cannot yet generalize to all customers."

ASK YOURSELF / Who had no voice, and what would make them easier to include?

Transfer challenge: find a recent headline or assertion where this same reasoning error could plausibly occur. Write one concrete limitation without claiming the entire story is false.

FRAMEWORK / 04 / CONFOUNDING

What happened before the intervention?

The everyday situation

An invented company finds that employees who attended optional skills training received higher performance scores. Did the training cause the difference? Perhaps employees with supportive managers were both more likely to attend and more likely to receive strong performance reviews.

Why the obvious conclusion can fail

A confounder is a variable related to both exposure and outcome that distorts the causal comparison. You cannot always “control away” every difference with a spreadsheet. Some factors are unmeasured; some controls introduce new bias; and some questions are not suitable for random assignment. Good analysis makes assumptions explicit.

One practical experiment

Compare groups before the training. Did they differ in role, prior scores, workload, tenure, and manager support? Specify what a credible no-training comparison would look like. Resist claiming a causal effect if your design supports association alone.

A MORE DEFENSIBLE SENTENCE

“Training attendance was associated with higher scores in this dataset; causation is not established.”

ASK YOURSELF / Which pre-existing difference could plausibly account for both outcomes?

Transfer challenge: find a recent headline or assertion where this same reasoning error could plausibly occur. Write one concrete limitation without claiming the entire story is false.

FRAMEWORK / 05 / TESTABILITY

Let an idea risk being wrong.

The everyday situation

Suppose a shop owner believes a reminder message increases attendance at a workshop. “Reminders work” is too loose. A more testable claim states: “Among people registered by the same deadline, sending an SMS 24 hours ahead will increase arrival rate compared with sending none.”

Why the obvious conclusion can fail

Before looking at results, specify outcome, timing, comparison, exclusions, and the practical size of a meaningful difference. Post-hoc storytelling can make a weak claim look unusually convincing. Falsifiability matters, but testability alone does not guarantee importance, sound measurement or ethical design.

One practical experiment

Write a prediction whose result can disappoint you. What measured outcome would lower your confidence? What result would be inconclusive because the study was too small or inaccurate? Decide this before celebrating an outcome.

A MORE DEFENSIBLE SENTENCE

“If the attendance rates remain similar under a fair comparison, this version of the reminder claim loses support.”

ASK YOURSELF / Which result would make me update rather than explain it away?

Transfer challenge: find a recent headline or assertion where this same reasoning error could plausibly occur. Write one concrete limitation without claiming the entire story is false.

FRAMEWORK / 06 / BASE RATES

A good detector can still cry wolf.

The everyday situation

Imagine a factory where only 2% of parts are defective. The sensor detects 90% of true defects and correctly dismisses 95% of sound parts. Out of 1,000 items, expect 20 defects, 18 correctly flagged defects, and 49 false alerts among 980 sound parts.

Why the obvious conclusion can fail

There are about 67 positive flags. Only 18 concern true defects in this simplified model: $18 / 67 \approx 26.9\%$. The sensor is not useless; its meaning depends on prevalence. A rare event can make false flags outnumber true flags even with a decent instrument.

One practical experiment

Before reacting to a positive alarm, ask: How common is the condition in this relevant population? What are the false-positive and true-positive rates? Convert percentages into counts for a group of 1,000. Numbers are often easier to reason about than adjectives.

A MORE DEFENSIBLE SENTENCE

"The chance of a defect among positive flags is about 27% under these toy assumptions."

ASK YOURSELF / What population did the base rate come from, and does it match this decision?

Transfer challenge: find a recent headline or assertion where this same reasoning error could plausibly occur. Write one concrete limitation without claiming the entire story is false.

FRAMEWORK / 07 / MEASUREMENT

When the proxy becomes the prize.

The everyday situation

A customer service team is rewarded for shorter calls. After three months, average call time falls. Is service better? Maybe agents became more efficient. Or maybe they hurried callers into repeat contacts. An isolated metric cannot settle the question.

Why the obvious conclusion can fail

This is the family of problems often discussed as Goodhart's law: when a measure becomes a target, incentives can change the relationship between that measure and the goal. There is no reason to abandon metrics altogether; the answer is careful design, multiple signals and space for informed judgment.

One practical experiment

Name the real outcome: customers' issues resolved fairly. List a speed metric, a quality metric, a negative-outcome metric and a way to hear customer stories. Imagine how each could be gamed. Set review points before a metric dominates behavior.

A MORE DEFENSIBLE SENTENCE

"Call time fell. Resolution quality and customer experience still need to be checked."

ASK YOURSELF / Which valuable human reality goes missing when we reward only what is easy to count?

Transfer challenge: find a recent headline or assertion where this same reasoning error could plausibly occur. Write one concrete limitation without claiming the entire story is false.

FRAMEWORK / 08 / REVISION

Changing your mind is part of knowing.

The everyday situation

A study reports an intriguing difference, but does not provide question wording, exclusion rules or enough detail to reproduce its calculation. This does not prove the result is false. It does limit how independently the result can be checked.

Why the obvious conclusion can fail

Reproducibility asks whether the same analysis can be repeated with the same materials. Replication asks whether a result can emerge again in a fresh investigation. Different designs can give different answers because contexts differ, effects are noisy and the original estimate may be imprecise. One flashy study rarely closes a serious question.

One practical experiment

Keep a revision ledger: “What I believed; why I believed it; new evidence; what changed; what still remains open.” If evidence supports only a narrower claim, narrow it. If it challenges an appealing story, allow the story to change.

A MORE DEFENSIBLE SENTENCE

“My current best estimate has changed because the available evidence changed.”

ASK YOURSELF / What would I want my future self to know about the limits of this judgment?

Transfer challenge: find a recent headline or assertion where this same reasoning error could plausibly occur. Write one concrete limitation without claiming the entire story is false.

CASE 01 / THE MESSAGE

The unanswered message

A friend reads your invitation on Wednesday and replies on Friday: “Sorry, it has been busy.” You had been tempted to conclude that you matter less to them than you thought.

OBSERVATION

The invitation was sent, a read indicator appeared, and a response arrived two days later.

INTERPRETATION

The delay means the friend does not care about the relationship.

RIVAL STORIES

They were overloaded, unsure about their schedule, forgot to reply, lacked bandwidth, or were less interested than before.

DISCRIMINATING EVIDENCE

Look at patterns over time and whether they respond clearly when plans are discussed. A respectful question may help more than speculation.

LIMITED CONCLUSION

One response delay is weak evidence for a broad claim about commitment or motive.

WORDS TO TAKE AWAY

Ask: “Would you prefer that I check in closer to the day? No pressure if your schedule is full.”

Your turn: change one detail in this scenario. Does the best explanation change? Which additional observation would matter most?

CASE 02 / THE CUSTOMER POLL

The very convincing 80%

A café owner polls 40 newsletter subscribers; 32 click “Love it” about a new menu. A marketing post says: “Eight out of ten customers love the redesign.”

OBSERVATION

Among respondents, $32/40 = 80\%$ selected the favorable option.

INTERPRETATION

The result describes all customers equally well.

RIVAL STORIES

Newsletter readers may be more engaged, response selection may favor enthusiasts, or wording may influence choices.

DISCRIMINATING EVIDENCE

Invite a broader customer sample, track who declines, use neutral wording, and compare purchasing behavior without conflating it with approval.

LIMITED CONCLUSION

A strong favorable response within this self-selected group; population approval remains uncertain.

WORDS TO TAKE AWAY

Rewrite: “80% of the 40 newsletter respondents who answered selected ‘Love it.’”

Your turn: change one detail in this scenario. Does the best explanation change? Which additional observation would matter most?

CASE 03 / THE TRAINING PROGRAM

The workshop that “caused” success

In a fictional firm, volunteers who attended a new training course achieved higher average evaluation scores. A manager announces the course improved everyone’s performance.

OBSERVATION

In this dataset, attendees had higher observed scores.

INTERPRETATION

The training itself caused the difference for all employees.

RIVAL STORIES

Prior skills, supportive managers, better schedules and greater willingness to volunteer may predict both outcomes.

DISCRIMINATING EVIDENCE

Compare pre-training trends, use a credible control group, disclose missing data, and consider whether people had equal access.

LIMITED CONCLUSION

An association worth investigating; estimated causal impact requires a stronger comparison.

WORDS TO TAKE AWAY

Write: “We observed a promising difference, but volunteer selection prevents a clean causal estimate.”

Your turn: change one detail in this scenario. Does the best explanation change? Which additional observation would matter most?

CASE 04 / THE SOCIAL SCORE

The dashboard that swallowed the goal

A creative group begins counting posts and reactions to measure community health. Output rises, but people complain that conversations feel shallower and less welcoming.

OBSERVATION

Publishing counts and reactions increased; some participants reported poorer experiences.

INTERPRETATION

Higher activity means the community is healthy.

RIVAL STORIES

Incentives may reward quantity; reactions may reflect outrage; newcomers may receive less attention.

DISCRIMINATING EVIDENCE

Look at retention, newcomers' qualitative reports, constructive exchanges, moderation costs and participation distribution.

LIMITED CONCLUSION

Activity is one dimension of engagement, not a complete description of community value.

WORDS TO TAKE AWAY

Ask: "If the visible number falls but the quality of participation improves, are we succeeding?"

Your turn: change one detail in this scenario. Does the best explanation change? Which additional observation would matter most?

CALCULATION / NO MAGIC

The base-rate trap, shown.

A transparent calculation for a fictional factory sensor, using expected counts.

Suppose 1,000 items enter a quality-control line. A defect is present in 2% of parts. The sensor catches 90% of real defects (sensitivity) and correctly clears 95% of good parts (specificity).

CLASSIFICATION	CALCULATION	EXPECTED ITEMS
Defective items	$1,000 \times 0.02$	20
True positive flags	20×0.90	18
Sound items	$1,000 \times 0.98$	980
False positive flags	980×0.05	49
All positive flags	$18 + 49$	67
True among positive	$18 / 67$	$\approx 26.9\%$

The equation

$P(\text{defect} \mid \text{flag}) = \text{sensitivity} \times \text{prevalence} / [\text{sensitivity} \times \text{prevalence} + \text{false-positive rate} \times (1 - \text{prevalence})]$

With these numbers: $(0.90 \times 0.02) / [(0.90 \times 0.02) + (0.05 \times 0.98)] = 0.018 / 0.067 \approx 0.2687$. This is a model of long-run expected proportions, not an observed batch outcome.

TRY THIS

If prevalence rises to 10% with the same sensor, true flags become 90 and false flags become 45 per 1,000; $90 / 135 \approx 66.7\%$. The test did not improve. The context changed.

ASK YOURSELF / Have I confused “accuracy of the instrument” with “probability that this alert is correct”?

CAUSAL MAP / THOUGHT EXPERIMENT

Draw the arrows before the conclusion.

A toy world where two trends share a common driver.

Suppose temperature T affects both lemonade purchases L and pool visits P . Define fictional quantities for a day as $L = 15 + 3(T - 20)$, $P = 5 + 2(T - 20)$, with temperatures between 20°C and 40°C . These equations were invented for instruction.

20°C : Lemonade purchases = 15, pool visits = 5.

30°C : Lemonade purchases = 45, pool visits = 25.

40°C : Lemonade purchases = 75, pool visits = 45.

GRAPH IN WORDS

TEMPERATURE $\rightarrow$ LEMONADE PURCHASES AND TEMPERATURE $\rightarrow$ POOL VISITS. There is intentionally no arrow between lemonade and swimming in this model.

A four-step causal check

- Name the proposed cause and outcome. What would a genuine causal arrow mean?
- List at least three shared causes, selection processes, or plausible reverse pathways.
- Decide which evidence would distinguish the competing graphs. What would you measure or compare?
- Write the narrowest conclusion supported by the actual design. Do not use the toy model to claim real-world measurements.

In real systems, all three variables may themselves be influenced by geography, income, school holidays or local events. A diagram reveals assumptions; it does not automatically verify them.

ASK YOURSELF / If my preferred causal arrow disappeared, what other graph could explain the data?

PRACTICE / 1 OF 2

Classify the claim.

All case files below are invented. Choose: R = observed report; I = interpretation; H = testable hypothesis; U = unsupported assertion.

CASE 01 — A project board records seven missed deadlines.

“Seven tasks passed their recorded deadlines in April.”

YOUR ANSWER (R / I / H / U): _____

CASE 02 — Same board.

“The team missed deadlines because nobody cares.”

YOUR ANSWER (R / I / H / U): _____

CASE 03 — Same board.

“Adding a midweek scope check will reduce late handoffs next month.”

YOUR ANSWER (R / I / H / U): _____

CASE 04 — Same board.

“Every delayed task proves the entire team is incompetent.”

YOUR ANSWER (R / I / H / U): _____

CASE 05 — A fictional café receives 40 survey replies; 32 approve.

“Of the 40 people who replied, 32 selected approve.”

YOUR ANSWER (R / I / H / U): _____

CASE 06 — Same survey.

“Most customers definitely love the new menu.”

YOUR ANSWER (R / I / H / U): _____

A USEFUL REMINDER

An observed report can describe an invented measurement. It does not suddenly become a verified real-world fact. If a phrase has multiple meanings, use the given context and defend your reasoning.

PRACTICE / 2 OF 2

Classify the claim.

All case files below are invented. Choose: R = observed report; I = interpretation; H = testable hypothesis; U = unsupported assertion.

CASE 07 — Same survey.

"A shorter menu will reduce average ordering time in the next two weeks."

YOUR ANSWER (R / I / H / U): _____

CASE 08 — Same survey.

"These 40 replies prove that no dissatisfied customer exists."

YOUR ANSWER (R / I / H / U): _____

CASE 09 — A toy sensor flags 67 items in a batch of 1,000.

"The instrument flagged 67 items in the batch."

YOUR ANSWER (R / I / H / U): _____

CASE 10 — Same sensor.

"The factory must be deteriorating."

YOUR ANSWER (R / I / H / U): _____

CASE 11 — Same sensor.

"Recalibrating the sensor will lower false alarms in the next batch."

YOUR ANSWER (R / I / H / U): _____

CASE 12 — Same sensor.

"Every alert is a confirmed defect."

YOUR ANSWER (R / I / H / U): _____

A USEFUL REMINDER

An observed report can describe an invented measurement. It does not suddenly become a verified real-world fact. If a phrase has multiple meanings, use the given context and defend your reasoning.

ANNOTATED KEY / 1 OF 2

A reason matters more than a label.

Return to your answers, examine the alternatives, then decide whether you would revise any classification.

01 / OBSERVED REPORT

This reports what the stated board recorded. It does not, by itself, establish why deadlines were missed.

Next check: Check timestamps, definitions of a deadline, and missing entries.

02 / INTERPRETATION

The claim assigns an unseen motive to a documented pattern. Workload, unclear scope, and competing commitments are alternatives.

Next check: Ask about blockers and compare workload and planning records.

03 / TESTABLE HYPOTHESIS

This is a proposed intervention with a measurable future outcome. It needs a comparator and consistent definition.

Next check: Compare similar projects with and without the check.

04 / UNSUPPORTED ASSERTION

An absolute judgment about people is not established by seven task records.

Next check: Narrow the claim and separate delivery outcomes from personal worth.

05 / OBSERVED REPORT

A report of the stipulated sample responses, not an estimate of every customer's opinion.

Next check: Check who received the survey and who replied.

06 / INTERPRETATION

The result is generalized beyond the responding sample. Selection and nonresponse could alter the picture.

Next check: Design a representative survey or compare behavior over time.

DID AN ANSWER FEEL AMBIGUOUS?

That can be productive. In practice, a sentence may contain both a reported event and an interpretation. When possible, separate the clauses rather than forcing one label onto the entire sentence.

ANNOTATED KEY / 2 OF 2

A reason matters more than a label.

Return to your answers, examine the alternatives, then decide whether you would revise any classification.

07 / TESTABLE HYPOTHESIS

It predicts a measurable result after a proposed change.

Next check: Predefine ordering time and compare matched periods.

08 / UNSUPPORTED ASSERTION

A sample cannot establish the nonexistence of dissatisfaction among all customers.

Next check: Qualify the conclusion: "Most respondents approved."

09 / OBSERVED REPORT

This is a stipulated instrument count. It says nothing yet about the instrument's accuracy.

Next check: Inspect a reference sample and error rates.

10 / INTERPRETATION

It infers a trend and a cause without historical comparison or calibration checks.

Next check: Compare prior batches and instrument settings.

11 / TESTABLE HYPOTHESIS

This can be evaluated with an independent reference test and defined false-alarm rate.

Next check: Predefine the comparison and validate actual defect status.

12 / UNSUPPORTED ASSERTION

A positive signal is not logically equivalent to a true defect; false positives may exist.

Next check: Use a reference test before making that assertion.

DID AN ANSWER FEEL AMBIGUOUS?

That can be productive. In practice, a sentence may contain both a reported event and an interpretation. When possible, separate the clauses rather than forcing one label onto the entire sentence.

WORKSHEET / 01

A one-page claim audit.

Print this page and write in your own words. Silence and uncertainty are acceptable answers.

CLAIM IN ONE TESTABLE SENTENCE

WHAT I CAN DIRECTLY OBSERVE OR VERIFY

MY CURRENT INTERPRETATION

AT LEAST TWO SERIOUS ALTERNATIVES

A RESULT THAT WOULD CHANGE MY MIND

MY MORE DEFENSIBLE CONCLUSION

WORKSHOP / 02

Two people. Two explanations.

A reusable 15-minute exercise for teams, students or curious friends.

One participant proposes a defensible everyday claim. Both separately write an observation and at least two explanations. Swap papers. Each person tries to state the other's preferred explanation more charitably than the original author did. Then jointly design one fair check. No one has to surrender a belief on demand.

WHAT IS THE ACTUAL OBSERVATION?

INTERPRETATION A

INTERPRETATION B

ONE COMMON CAUSE OR SELECTION EFFECT

WHAT WOULD DISTINGUISH A FROM B?

A JOINTLY AGREED, CALIBRATED CONCLUSION

GROUP RULE

Challenge the claim, not the person. Do not force private disclosure, diagnostic speculation, confession or public scoring. Respect boundaries and the option to pass.

ENDNOTES / PRACTICAL COMPASS

Keep curiosity. Keep limits.

A small vocabulary, a few starting points and an invitation to use these tools gently.

Base rate: the prevalence of an event before new evidence. **Confounder:** a variable that affects causal comparisons. **Selection bias:** distortion from how observations enter a dataset. **Calibration:** whether expressed confidence matches outcome frequencies over many comparable cases. **Falsifiability:** whether evidence could in principle count against a claim. **Replication:** a fresh test of a previously reported finding.

Further study

- Douglas G. Altman, *Practical Statistics for Medical Research* (1991): introductory applied statistics and study design.
- Judea Pearl, *Causality*, 2nd ed. (2009): formal approaches to causal reasoning.
- Gerd Gigerenzer, *Risk Savvy* (2014): accessible reasoning about uncertainty and natural frequencies.
- The Cochrane Handbook for Systematic Reviews of Interventions: assessing bias and synthesizing evidence (www.cochrane.org).
- The Stanford Encyclopedia of Philosophy: entries on scientific method, probability and induction (plato.stanford.edu).

A note about evidence

Every situation, statistic and game scenario in this guide is a pedagogical construction unless a published work is explicitly cited. The resources above are suggestions for further study, not citations for any invented example. This material is educational, not individual legal, medical or financial advice.

ONE QUESTION TO KEEP

“What conclusion is fair to both the evidence I have and the evidence I do not have?”

A small invitation

The wider universe of projects and writing lives at sahiltayal.com. If you find an error, alternative interpretation or better question, write to sahil@sahiltayal.com. A field guide should be open to revision.

Prepared for the Sahil Tayal universe • 2026 • Free to download and keep.