CURIOSITY ENGINE.

A FIELD GUIDE TO BETTER QUESTIONS



An illustrated notebook for asking what we do not know without confusing an intriguing story with an answer.

SAHIL TAYAL / RESEARCH MEETS ART / 2026

02 / EDITORIAL NOTE

The question before the answer

A question can be aesthetically compelling without being answerable. A dramatic claim may draw attention while withholding the crucial distinction between what was observed, what was inferred, and what would test the inference.

This guide makes one narrow argument. When a question names an information gap, a rival explanation, and a plausible next observation, it becomes easier to evaluate. That is a writing and reasoning practice, not a guarantee of more curiosity, better memory, or personal improvement.

Start with a plain example: a museum label describes an object and gives its date but says little about its provenance. A vivid story about how it arrived could be captivating. The disciplined question asks which records would distinguish competing accounts. A gap can invite wonder; evidence decides what follows.

These scenes are fictional instructional examples. They describe no actual collection, person or institution. The published studies in the final pages are real sources, and their findings should not be confused with claims about this model.

TRY THIS

In your notebook: describe one claim that interests you and one observation that would make it less convincing.

03 / LEARNING OUTCOMES

Three ways to read a gap

Outcome one: distinguish a question about missing information from an assertion disguised as a question. Compare 'Why did the curator hide the truth?' with 'What documentation would establish the object's chain of custody?' The first presupposes a motive. The second invites a check.

Outcome two: generate a competing explanation. When two garden beds flower differently, a question about watering might be reasonable, yet soil, shade, variety and the timing of observation may also matter. Naming a rival story improves the quality of inquiry even when the answer remains uncertain.

Outcome three: choose a proportionate next observation. A question can demand a multi-year study when only a ten-minute inspection is possible. Under a short time constraint, make a careful note and state that the observation cannot establish causation. The scale of evidence should fit the size of the claim.

A useful result of this guide is not feeling perpetually unfinished. Sometimes a provisional answer is enough to act; sometimes a good question makes uncertainty more comfortable to carry.

TRY THIS

Choose one outcome. Write a sentence that would show you have practised it.

04 / ORIGINAL DIAGRAM

A concept map, not a brain scan

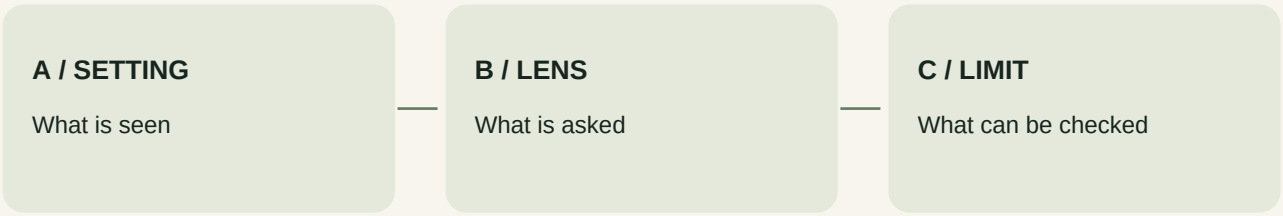

The engine is a deliberately simple symbolic map. Setting A describes what is observable. Lens B determines what kind of question is written. Constraint C narrows the next practical step. Together they produce a proposed question and an evidence action.

No arrow in this drawing represents a measured neural connection, causal law, psychological trait or statistical relationship. The diagram encodes the author's design decisions, which can be inspected and changed without an opaque system.

Notice the asymmetry: the setting gives the model a subject, but the lens does much of the rhetorical work. Keeping A and C fixed while changing B exposes that choice. It can make the author's assumptions visible even when no original information has been discovered.

This is why the lab has no personality score, improvement score or leaderboard. It teaches a discipline of formulation rather than evaluating a visitor.

TRY THIS

Try drawing your own three-box model. Which input is most contestable?

05 / DISTINGUISHING IDEAS

Four mechanisms, several theories

An information gap is a perceived difference between what is known and what one wants to know. Loewenstein (1994) developed an influential account of the tension such gaps can create. This is one contribution to a larger, contested literature, not a universal definition.

Novelty and surprise may redirect attention before a gap can be named in words. Other accounts study exploratory behaviour, learning progress, informational value and the relation between curiosity and reward. Gottlieb and colleagues (2013) review several of these mechanisms.

Interest-driven and deprivation-like curiosity are also distinguished in some psychometric traditions; Litman and Silvia (2006) report evidence for differentiated dimensions. These are research constructs, not the three writing choices in this website. Never interpret a selected lens as identifying a psychological type.

An alternative possibility matters: too much uncertainty can discourage inquiry; a reassuring answer can sometimes be preferable to endless investigation. Framing a question well does not remove social constraints, fatigue, information costs or the possibility that no satisfactory answer exists.

TRY THIS

What is one form of curiosity that does not feel like solving a missing fact?

06 / DOCUMENTED RESEARCH

Case study A: curiosity and learning

Gruber, Gelman and Ranganath (2014) investigated states of curiosity and their relation to hippocampus-dependent learning. They reported behavioural and neural findings within a laboratory design. The study provides evidence relevant to curiosity during learning; it did not test this website's question-writing templates.

Kang and colleagues (2009) studied epistemic curiosity with trivia questions and examined reward-related brain activity, willingness to spend limited resources to learn answers, and later recall. The details of study population, tasks and measurement matter: results from trivia experiments cannot simply be transferred to every learning setting.

The studies invite a modest editorial insight. What a person wants to know may change engagement with information. But this guide makes no claim that selecting 'surprising assumption' over 'missing information' produces the same effects measured in those experiments.

Replication, variation across contexts, and causal intervention would all matter before claiming that our particular question generator improves retention. A source is an invitation to scrutinise methods, not a decorative authority badge.

TRY THIS

What outcome did the published study measure? What outcome does this tool actually display?

07 / DOCUMENTED RESEARCH

Case study B: perceptual uncertainty

Jepma and colleagues (2012) examined perceptual curiosity using ambiguous visual information and its later clarification. Their study connected task phases with neural responses and memory measures. It is a different kind of curiosity manipulation from the explicitly verbal questions in our exhibit.

Kidd and Hayden (2015) review psychological and neuroscientific approaches to curiosity and explain why a single agreed definition and standard laboratory paradigm have been difficult to establish. That disagreement is valuable: it prevents an attractive design metaphor from hardening into a claim about all human beings.

Compare two claims. 'Uncertainty can motivate information-seeking in some experiments' is a proposition a source may help assess. 'This template makes users more curious' is a specific causal claim that would require its own evaluation. The first does not establish the second.

Research is most useful when it helps us mark the distance between evidence and extrapolation. A bibliography should make that distance visible, especially when an interface feels persuasive.

TRY THIS

Record one tempting extrapolation and rewrite it to match the actual evidence.

08 / HOW THE MODEL WORKS

The Question Generator Lab: the rules

A is the initial condition. Choose one fictional scene: a museum label with missing provenance, two garden beds flowering differently, or a classroom with competing interpretations. Each scene supplies an observation and an ambiguity.

B is the variable under test. Choose one of three fixed templates. Missing information asks which evidence could close a gap. Surprising assumption asks what presupposition might fail. Controlled comparison asks how to separate an observed difference from a causal story.

C is the practical constraint. Checkable sources points to a document or observation record. Only ten minutes encourages a narrow note instead of a strong inference. No new equipment asks for an available comparison and an alternative explanation.

These choices produce one question and one next action. There is no language-model request, measured probability, external storage or real-time evidence gathering. Each result is an authored conditional rule. Changing an input cannot change facts outside the model.

TRY THIS

Open the website and inspect its visible rules. Can you explain the outcome without guessing?

09 / BASELINE AND PERTURBATION

A controlled change: the full transcript

BASELINE - Inputs: museum / missing information / checkable sources. The fictional label lists a date but not an object's provenance. The question is: 'What missing evidence about the object's provenance would distinguish a plausible story from a supported account?' The action is to seek a checkable source such as an accession record.

PERTURBATION - Keep A and C exactly the same. Change B to surprising assumption. The result becomes: 'Which assumption about the object's provenance is surprising, and what observation could challenge it?' The action still proposes the checkable source. The words changed because the writing rule changed, not because the evidence changed.

WHAT THE COMPARISON SHOWS - A question can direct attention differently while holding its subject and practical limits constant. The model makes this difference inspectable. It does not establish which wording would work better for learning, persistence, memory or truth in an actual study.

Use the undo command to return to the earlier state, or reset to the baseline. The third preset deliberately breaks the expectation that a crisp question always makes the answer available.

TRY THIS

Underline the words that change across the two questions. Name the rule responsible.

10 / WORKING AGAINST THE MODEL

The counterexample: a neat question fails

Counterexample - Inputs: garden / controlled comparison / only ten minutes. Two fictional beds flower at different times. The question asks what comparison could show whether the difference in bloom timing matters rather than merely attracting attention.

At first the question sounds disciplined. Yet a ten-minute visit cannot reliably identify which combination of soil, shade, species, watering or weather shaped growth across weeks. A better sentence is not a substitute for a suitable study design.

A rival interpretation might be that the two beds were observed at different stages of flowering, rather than that a single environmental factor caused the difference. We do not know. The model should show this limitation, not conceal it behind an impressive graphic.

The strongest criticism of the engine is that it may overvalue explicit analytic questioning. Exploration, play and tacit pattern recognition can also generate useful discoveries. A formal question is one tool among many, not a moral duty or a universal ideal.

TRY THIS

Name a situation where making the question narrower would be premature.

11 / ORIGINAL EVIDENCE MAP

Evidence map: what supports what?

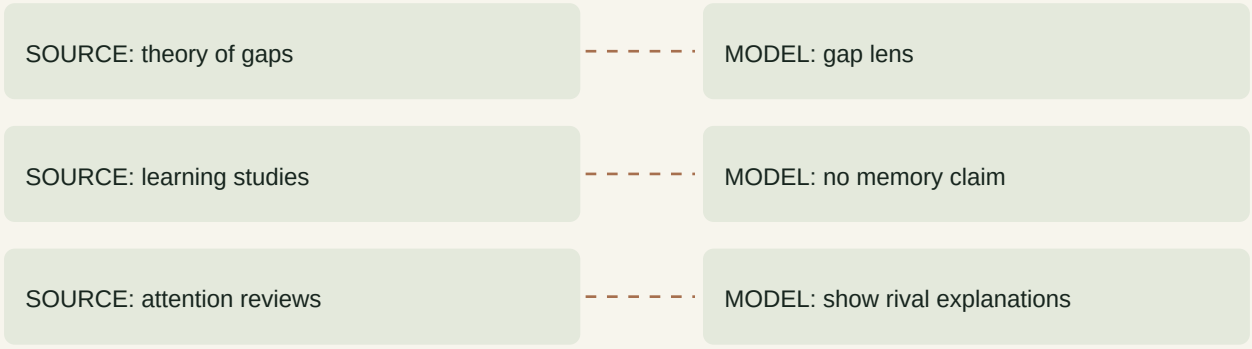

Dashed bridge = conceptual influence, NOT experimental validation

On the left are documented findings and theoretical arguments about curiosity and information-seeking. On the right are design choices made for this exhibit. A connecting line indicates inspiration or interpretive relevance, not a validated causal bridge.

Loewenstein (1994) supports describing information gaps as one account of curiosity. Gruber and colleagues (2014), Kang and colleagues (2009), and Jepma and colleagues (2012) address particular learning or perceptual experiments. None supplies evidence for the exact effectiveness of this generator.

Question templates, three preset scenarios, the undo button and the figure styling belong to the authorial design layer. Their outcomes are reproducible by code inspection. Reproducibility of a software rule is not the same as external validity of a psychological hypothesis.

The question 'What would change your mind?' can therefore be asked at two levels: Which new primary evidence would alter the scene interpretation? And which independent test would show whether this interface helps people reason differently?

TRY THIS

Draw a dashed line from one research source to a design choice. Write 'inspired by', not 'proved by'.

12 / AN ACTIVITY YOU CAN RUN

Practical workshop: two small inquiries

EXERCISE A - Record a claim from an ordinary context, without naming a private person. Example: 'This plant grows better beside the window.' First identify the observation: one plant appears taller. Then list at least two competing explanations: light and differences in watering or species. Avoid quietly turning an observation into a causal conclusion.

Apply the three lenses. Gap: what facts are missing? Assumption: what part of the initial explanation might be wrong? Comparison: what observation could distinguish alternatives? Choose a practical constraint: existing records, a short observation or no new equipment.

EXERCISE B - Read a short public article. Separate the author's main claim from its evidence. Write a question that could be answered by a specific cited source. Write a second question that might remain unresolved even after reading the source.

When time is short, do not manufacture certainty. Decide whether to continue, stop provisionally or accept that the question is not yet answerable. The exercise has no score and does not collect results.

TRY THIS

Try the exercise with something modest and harmless before applying it to complex public claims.

13 / WORKSHEET

Printable working page

Choose a non-sensitive observation. Write what you saw before explaining why it happened. Keep personal, clinical and confidential information out of this worksheet if you plan to share it.

Record the current assumption, an alternative hypothesis, a precise question, the next feasible observation and what would still remain uncertain. The blank rules below are for handwriting, not automated evaluation.

OBSERVATION / WITHOUT INFERENCE

CURRENT ASSUMPTION

A RIVAL EXPLANATION

THE TESTABLE QUESTION

NEXT FEASIBLE OBSERVATION

WHAT REMAINS UNKNOWN

14 / SUGGESTED ANSWER + DISSENT

An answer worth revising

Sample observation: one of two herb seedlings is taller. Initial story: the taller plant received more daylight. Rival story: it was planted earlier or has different watering conditions. A better question is: 'Which measurements would help distinguish light exposure from differences in planting date or watering?'

Feasible next step: examine the dates recorded on the pots, compare watering logs and make repeated observations under comparable conditions. Strong conclusion not justified: 'Daylight caused the height difference.' The example teaches restraint, not a complete horticultural experiment.

A critic may reply that deliberate comparison slows a learner at the very moment exploration should be playful. That is a reasonable objection. Try alternating two modes: ten minutes of free curiosity, then two minutes of question clarification. Compare the kinds of insights each yields; avoid assuming one is always superior.

The point is not to eradicate uncertainty. A good inquiry may reveal that the available evidence is insufficient. A choice to stop asking can be thoughtful, too.

TRY THIS

Write one finding you can defend and one interpretation you will keep provisional.

15 / HOW TO USE THIS PUBLICATION

Labels, privacy, limits and licence

EVIDENCE - A finding or argument in published work, identified with a source. INTERPRETATION - An editorial inference or synthesis that might reasonably be disputed. ILLUSTRATIVE MODEL - A hypothetical rule or scenario created to explain an idea. SPECULATION - A possibility that has not been demonstrated.

This guide does not diagnose curiosity, assess intelligence, assign identity or offer educational or medical advice. Its fictional scenes are not field research. The website runs transparent rules inside the browser and does not require an account. Exported explanations are created locally. A session-only flag may shorten the return-visit entrance; it is not transmitted to the creator by the application.

Copyright 2026 Sahil Tayal. Original text, diagrams, exercises and software explanations created for this publication are offered under Creative Commons Attribution 4.0 International. You may share and adapt this original work, including commercially, with appropriate credit, a link to the licence and an indication of changes. Third-party research articles and titles are not relicensed by this notice.

Suggested credit: Sahil Tayal - Curiosity Engine, https://sahiltayal.com/curiosity-engine/, CC BY 4.0. For the licence terms see https://creativecommons.org/licenses/by/4.0/. Contact: sahil@sahiltayal.com.

TRY THIS

When sharing a revised worksheet, mark what you changed and credit the original.

16 / REFERENCES + RELATED READING

Source ledger and neighbouring worlds

All sources below were identified through publisher, DOI or PubMed records. These are research references for the discussion of curiosity, not validators of the lab's question templates. Read the methods and limitations in the original works before extending a claim.

Related Sahil Tayal work: Lifelong Learning Lab (https://learning.sahiltayal.online/) examines learning as a practice. Unlearning (https://unlearning.sahiltayal.online/) explores revising assumptions. The World Engine (https://worldengine.sahiltayal.com/) makes interdisciplinary questions using transparent rules. Research Questions (https://research.sahiltayal.com/) distinguishes a working framework from validated findings.

Return to the Sahilverse: https://sahiltayal.com/. Original guide and worksheets: CC BY 4.0; source publications retain their own rights.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation.
doi.org/10.1037/0033-2909.116.1.75

Kang, M. J., et al. (2009). The wick in the candle of learning.
pubmed.ncbi.nlm.nih.gov/19619181/

Jepma, M., et al. (2012). Neural mechanisms underlying perceptual curiosity.
pubmed.ncbi.nlm.nih.gov/22347853/

Gottlieb, J., et al. (2013). Information-seeking, curiosity, and attention.
pubmed.ncbi.nlm.nih.gov/24126129/

Gruber, M. J., et al. (2014). States of curiosity modulate learning.
pubmed.ncbi.nlm.nih.gov/25284006/

Kidd, C., and Hayden, B. Y. (2015). Psychology and neuroscience of curiosity.
pubmed.ncbi.nlm.nih.gov/26539887/

Oudeyer, P.-Y., Gottlieb, J., and Lopes, M. (2016). Intrinsic motivation, curiosity and learning.
pubmed.ncbi.nlm.nih.gov/27926442/

Kashdan, T. B., Rose, P., and Fincham, F. (2004). Curiosity and exploration.
pubmed.ncbi.nlm.nih.gov/15151805/

Litman, J. A., and Silvia, P. J. (2006). Interest and deprivation curiosity.
pubmed.ncbi.nlm.nih.gov/16740115/