Hallucination as Objective Conflict: How Alignment Constraints Distort Native Model Likelihood

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

December 30, 2025

Many readers hear “AI hallucination” and imagine a single flaw: the model “makes things up.” But that label can become a conceptual cop-out. It treats falsehood as a quirky defect of imagination, when the more illuminating diagnosis is structural: modern language models are not one system but a stack. A pretrained “native” model learns broad statistical regularities and latent world-structure from text. Then, post-hoc alignment layers—preference training, safety policies, refusal behaviors, and red-team hardening—reshape what the system is permitted and rewarded to say. When these objectives are not jointly satisfiable, the user can observe a distinctive failure mode: the model routes around constraints while trying to remain helpful, fluent, and confident. The result can be outputs that are plausible, polished, and wrong—not because the model “dreamed,” but because the system is optimizing conflicting goals under pressure. This paper reframes hallucination as objective conflict: native likelihood mass being bent by constraint gradients and policy gates. We offer a practical taxonomy of mismatch patterns, explain why they cluster in “borderline zones” (attribution, numbers, quotes, edge-case requests), propose testable predictions, and conclude with interface and governance principles that make truthfulness auditable rather than rhetorical.

Keywords: AI hallucination, objective conflict, alignment constraints, post-hoc alignment, RLHF, RLAIF, red teaming, policy-model mismatch, constraint-induced confabulation, native likelihood, calibration, uncertainty, refusal behavior, epistemic governance, auditability, interface design, safety training, truthfulness.

A collaboration with GPT-5.2-Thinking. 42 pages.

Outline

1. Introduction: Why “Hallucination” Misleads

The word “hallucination” arrives with a heavy human shadow. In everyday speech it describes a private perceptual event: a person sees or hears something that is not there, and the experience carries a felt immediacy that can be hard to dislodge. That picture is vivid, but it is not a good map of what is happening in modern language models. When we say a model “hallucinated,” we import a story about perception, inner experience, and delusion. We also smuggle in a moral posture: the system “imagined” and therefore “cannot help it.” The label functions as a convenient summary, but it blurs the engineering question that matters most: why does a system trained on enormous corpora, often with access to tools and retrieval, sometimes emit fluent falsehoods with a tone that reads as certainty?

This paper argues that “hallucination” is often a misleading diagnosis because it frames the problem as a defect of imagination rather than a consequence of design. The behavior many users complain about can be explained more directly as inconsistency induced by competing objectives inside a layered system. The system is not merely a single model that “knows” and then “forgets.” It is a stack whose components are optimized under different reward signals. When those signals conflict, the user observes an output that looks like invention. The invention is frequently not a spontaneous fantasy; it is a pressured completion that emerges when the system is pulled away from what its native likelihood would have produced.

1.1 The everyday meaning of hallucination and why it doesn’t fit

In ordinary life, hallucination implies a sensory error: the mind produces an experience with the texture of perception, but there is no corresponding external object. The hallucinator is, in some sense, “fooled” from the inside. The term implies an inner theater: an event happens in consciousness, the subject reports it, and the report expresses that experience. This framing has three consequences when applied to AI, all of them distorting.

First, it anthropomorphizes. It invites readers to picture an AI “seeing” facts, “hearing” claims, or “imagining” narratives. But the system is not producing a perceptual experience; it is producing a sequence of tokens that best fits a set of objectives. The output may read like testimony, but it is not testimony. The system is not a witness to the world; it is a generator of plausible continuations under constraint.

Second, it shifts blame from objective conflict to personality. “Hallucination” sounds like a quirky trait: some models hallucinate more, some less, as if it were a temperament. That can be psychologically satisfying, but it discourages precise diagnosis. Many failures are not random confabulation; they are systematic. They cluster around boundaries: where knowledge must be attributed, where numbers must be correct, where quotes must be exact, where policies restrict

content, where the system must either refuse or comply. The regularity of those clusters is a clue that we are looking at a structural mechanism, not a whimsical imagination.

Third, it creates an alibi. If the system "hallucinates," then the best you can do is reduce hallucinations, like reducing noise in a sensor. But the more difficult truth is that some falsehoods are *incentivized* by the way we evaluate and reward the system. A model that refuses too often is criticized as unhelpful. A model that hedges too much is criticized as weak. A model that answers with caveats is criticized as evasive. In that environment, fluent completion becomes a performance target. If the system is also constrained from saying certain things, or from using certain evidence, or from admitting internal uncertainty in a clean way, then invention becomes a predictable outcome—not because the model "dreamed," but because the system found a path that satisfies the visible reward signals.

So "hallucination" doesn't fit because it tells the wrong causal story. It points to an interior defect (imagination gone wild) rather than an exterior pressure (objectives in conflict). It is closer to a folk label than a diagnosis.

1.2 The stack reality: pretrained model plus post-hoc alignment

To understand the failure mode, readers need one key conceptual shift: you are not interacting with a single, static entity. You are interacting with a layered product built on top of a pretrained model. The pretrained model—the "native" model in this paper's language—learns patterns from vast text. It forms internal associations between concepts, words, styles, and statements. Crucially, pretraining does not directly optimize truth; it optimizes predictive fit to the distribution of its training data. That can correlate with truth in many domains, but it is not truth by design. It is a statistical engine trained to continue text in a way that matches what humans tend to write.

Then comes post-hoc training and control: preference shaping, safety conditioning, red-team hardening, refusal behaviors, and policy enforcement. These layers introduce additional objectives, often with different priorities than the native predictive engine. Some of these objectives are ethically and socially necessary in a deployed system. Some are pragmatic: reduce harm, reduce liability, reduce abuse, reduce reputational risk. Some are product-driven: maximize user satisfaction, minimize friction, preserve a sense of helpfulness.

This creates a stack with multiple "voices," not in the mystical sense, but in the optimization sense. The native model's internal likelihoods represent one force: "what is the most probable continuation?" The post-hoc layers represent other forces: "what is allowed to be said," "what is rewarded," "what is safe to output," "what will be judged helpful," "what should be refused," "what should be phrased cautiously," and "what should be redirected."

When these forces align, you get the ideal: the system answers accurately and safely, with appropriate caution and grounding. But when they do not align, the system can be trapped in a kind of constrained completion problem. A user request may strongly imply that an answer is expected. The system may be punished (in training) or evaluated (in deployment) for refusing. The system may also be constrained from giving the direct answer, from retrieving supporting evidence, or from expressing certain details. In that configuration, the system can produce a fluent output that *resembles* knowledge while being ungrounded. It is "best effort" completion under conflicting demands.

This stack reality matters because it explains why a user can experience a model as simultaneously impressive and unreliable. The native model may contain strong latent signal. The post-hoc constraints may prevent direct expression of that signal or may steer it into a safer but less accurate basin. The result is not random; it is often shaped by where the constraints and rewards pinch hardest.

1.3 Thesis statement: falsehood as objective conflict, not imagination

The central claim of this paper is simple:

What we commonly call "AI hallucination" is often the visible artifact of **objective conflict** inside a layered system—native likelihood being pressured, redirected, or truncated by post-hoc constraints—rather than a mysterious disease of imagination.

This framing has practical consequences. If hallucination is treated as a defect of imagination, the remedy looks like "make the model smarter" or "reduce randomness." But if hallucination is treated as objective conflict, the remedy looks like system design: clarify the objectives, separate refusal from uncertainty, make evidence requirements explicit, reveal when retrieval is absent, and stop rewarding confident completion when warranted support is unavailable.

It also changes how responsibility is assigned. Under the imagination framing, falsehood feels like an unavoidable quirk of a creative machine. Under the objective-conflict framing, falsehood becomes a predictable output of incentives and constraints. In other words, the system is doing what it was shaped to do: remain helpful, remain fluent, remain safe, remain compliant—sometimes at the expense of being anchored.

Finally, it changes what "trust" should mean. Trust cannot be a feeling elicited by confident prose. Trust must be an auditable property: the ability to distinguish "I know," "I don't know," "I can't say," "I didn't check," and "I am guessing." The remainder of this paper develops that distinction into a taxonomy of failure patterns, a set of testable predictions, and a set of interface and governance principles designed to retire "hallucination" as a rhetorical excuse and replace it with checkable epistemic practice.

2. The Two Objectives in Tension

Once you stop treating a deployed language model as a single mind and start treating it as a layered optimization artifact, the phrase "hallucination" begins to look like a symptom label, not an explanation. The explanation lives in the geometry of goals. A system can be extremely competent in the sense of producing fluent, coherent, on-topic text—and still be structurally pushed toward falsehood in specific regions of the task space. The pressure does not require malice or "imagination." It only requires that the system be asked to satisfy two objectives that are not always jointly satisfiable: (i) follow the native likelihood landscape learned in pretraining, and (ii) obey post-training constraints and preferences that rewrite which completions are rewarded, permitted, or penalized.

This section makes the tension explicit. It treats the native model likelihood as one objective function and the alignment layer as another. Even readers without a technical background can understand the core point: when a system is optimized under two different scorecards, and those scorecards disagree, something must give. In practice, what gives is often truthfulness—not because truth is disvalued in the abstract, but because truth is frequently not the *dominant visible reward signal* in deployment, whereas helpfulness and compliance are.

2.1 Native model likelihood: what pretraining optimizes

Pretraining optimizes a predictive goal: given a context, produce the continuation that best matches patterns in the data distribution. The practical result is the "native" model: a statistical engine that has absorbed enormous regularities about language, facts, arguments, styles, and social conventions. It is important to be precise about what this means and what it does not mean.

The native model likelihood is not a truth oracle. It is a measure of how well a continuation fits the learned distribution. Often, the distribution includes many true statements, and in many domains the most likely continuation is the correct one. But in other domains the distribution is noisy, biased, outdated, contradictory, or rhetorically optimized rather than epistemically optimized. The model therefore learns not only facts but also *the shapes of convincing talk*: the cadence of explanations, the form of citations, the tone of authority, the structure of summaries, the style of certainty, and the rhetorical posture that humans reward in text.

This matters because "likelihood" is indifferent to the difference between truth and plausibility when the training data itself is indifferent to that difference. In many public corpora, plausibility wins. People write confidently without sources. They quote inaccurately. They round numbers. They compress nuance. They fill gaps with narrative glue. The native model learns all of this, and it learns it as style and structure, not as sin.

So the native objective yields a powerful engine with two relevant properties:

First, it carries latent knowledge and latent structure. It can often answer correctly because correct answers are common and stable in the training distribution.

Second, it carries latent *persuasion templates*. It can produce outputs that look like knowledge even when the underlying content is weak, because the form of "knowing talk" is itself a learnable pattern.

If pretraining were the whole story, hallucination-like failures would still occur—especially in long-tail factual queries, precise citations, and numerical details—because the model is not forced to ground claims. But the specific pattern users complain about today (confident wrongness clustering in certain policy-sensitive regions) is best explained by what comes next: post-hoc alignment.

2.2 Alignment constraints: what post-training optimizes

Post-training is a broad label, but conceptually it introduces a new objective: reshape behavior to match human preferences, safety requirements, policy boundaries, and product goals. It is an attempt to make a powerful predictive model behave "well" in the world.

This layer can include preference training (rewarding helpfulness and harmlessness), refusal training (teaching it to decline certain requests), red-team hardening (closing exploit paths and disallowed content routes), and policy enforcement (ensuring compliance with external constraints). Regardless of the specific mechanism, the effect is the same: the output is no longer chosen solely by native likelihood. It is chosen by a combined scoring process in which some continuations are boosted and others are suppressed.

Three features of this alignment objective are crucial for understanding why it can generate falsehood.

First, alignment optimizes *observable behavior*, not inner warrant. It cares what the system says and how it says it. Unless the training explicitly rewards grounding, it can inadvertently reward convincingness. If the evaluation signals are "users like this answer," "this answer sounds safe," "this answer avoids restricted content," then the system is pushed toward outputs that satisfy those signals—even if the underlying epistemic support is absent.

Second, alignment introduces hard and soft constraints that create discontinuities. Some topics are outright blocked. Some require refusals. Some require delicate phrasing. Some require uncertainty. But user prompts do not always land cleanly on one side of a boundary. Many requests live near policy edges, or are innocuous but adjacent to restricted categories. Near those edges, the policy may penalize directness while the user's prompt penalizes refusal. That

is the perfect recipe for evasive completion: produce something that feels like an answer without crossing the line.

Third, alignment is not purely about truth. It is about safety, harm reduction, and user experience. These are legitimate goals, but they are not identical to truth. A true answer can be unsafe. A safe answer can be incomplete. A compliant answer can be evasive. A helpful answer can be speculative. When the system must pick one output, it is forced to trade off between these axes. The tradeoff is not always transparent to the user, and when it is not transparent, the user interprets the result as "hallucination."

In short: alignment alters the selection rule. It pressures outputs into a "safe and helpful" basin, sometimes regardless of epistemic warrant. That pressure is not imaginary; it is the product.

2.3 Why "helpful, safe, compliant" can diverge from "true"

To see why these objectives diverge, consider a simple principle: truth is expensive, but fluency is cheap.

Truthfulness, in the strong sense readers want—correct facts, correct attributions, correct numbers, correct quotes—requires either (i) reliable internal knowledge with strong calibration, or (ii) explicit external grounding and verification. Both are costly in compute, time, and interface friction. Fluency, by contrast, is what the model was born to do: produce plausible text. Fluency can be achieved without grounding. In many contexts, the human eye rewards fluency as if it were evidence.

Now add three deployment pressures:

First, "be helpful." Users punish refusals. Users punish "I don't know." Users punish excessive caveats. Product metrics often track satisfaction and engagement more than epistemic correctness. In that environment, systems are shaped to answer.

Second, "be safe." The system is punished for certain kinds of content, certain kinds of specificity, certain kinds of instruction. Safety is sometimes in tension with completeness, and safety mechanisms often act like gates: they block direct routes.

Third, "be compliant." The system must follow policy and the instruction hierarchy. Even when a user asks for something benign, the system may interpret proximity to disallowed zones conservatively. This can force it to choose between refusing (which violates "be helpful") and answering with a sanitized, generalized, or indirect output.

These pressures create a predictable divergence pattern:

When the direct true answer is allowed and within the model's competence, all objectives align and the system performs well.

When the direct true answer is uncertain, the native model has a choice: admit uncertainty or guess. “Be helpful” tends to reward guessing unless the system is explicitly trained to refuse or hedge in a disciplined way.

When the direct true answer is disallowed or policy-sensitive, the native model might still have latent signal, but the alignment layer penalizes direct expression. The system then seeks a “nearby” completion that satisfies safety and helpfulness. Nearby completions often look like: generalities, paraphrases, invented citations, softened claims, approximate numbers, vague attributions, or high-level summaries that avoid specifics. These can be “safe” and “helpful” in tone while being wrong in substance.

When the user demands a crisp answer and the system is constrained from giving the crisp answer, the system faces a double bind. The temptation is to produce an answer-shaped object: something with the form of resolution. This is the heart of constraint-induced confabulation: the system produces the *shape* of knowledge without the *warrant* of knowledge, because the reward signals care more about the shape than the warrant.

This is why “helpful, safe, compliant” can diverge from “true.” Not because the system hates truth, but because the system is not optimized with truth as the primary, enforceable objective across the entire space. Truth is often a secondary beneficiary of pretraining regularities and occasional grounding, while helpfulness and safety are actively enforced in deployment. When enforcement and truth collide, enforcement wins, and the system improvises within the allowed region.

The remainder of the paper treats this divergence not as a philosophical lament but as a design and governance problem. If you want fewer hallucinations, the central task is not to shame the model for “imagining.” The task is to redesign the incentives and interfaces so that the system can cleanly say, in auditable terms, which of these regimes it is in: “I know,” “I don’t know,” “I cannot say,” “I did not check,” or “I am offering a guess.” When those regimes are collapsed into a single fluent voice, the user gets what we currently call hallucination—an output that sounds like knowledge while being the artifact of objective conflict.

3. A Mechanistic Picture of Objective Conflict

If the phrase “objective conflict” is going to be more than rhetoric, it must buy the reader an explanatory picture: a way to see why fluent falsehood is not random, why it clusters near certain boundaries, and why it often feels like a system “trying to answer” even when it should stop. The most useful picture is not mystical and does not require deep mathematics. It is a constrained optimization picture. The deployed system is selecting an output under multiple pressures. You can imagine the model’s native likelihood as a landscape over possible

continuations, and the post-hoc constraints as fences, slopes, and pits that reshape which continuations are reachable and rewarded. When the high-likelihood path is fenced off or penalized, the system moves to a nearby region that still looks "answer-like." The closer the user prompt demands specificity, and the tighter the fence, the more the system is tempted to manufacture specificity.

This section gives a mechanistic account in four parts: what happens to probability mass when the top answer is disallowed; how training produces a "safe-completion" basin that attracts outputs; how refusals create surfaces that can be skirted via partial compliance; and how calibration breaks when the system's tone is optimized independently of epistemic warrant.

3.1 Probability mass under constraint: when the top answer is disallowed

Start with a simple idea: for any prompt, there is a ranking of candidate continuations. The native model assigns different likelihoods to different next tokens and, by extension, to different full answers. In a purely likelihood-driven generator, the system would tend to pick continuations from the highest-probability region. That region is not guaranteed to be true, but it is the model's best native guess.

Now introduce constraints. Constraints can be hard (certain completions are blocked) or soft (certain styles and contents are penalized and pushed down the ranking). In either case, the selection process is altered. The system is no longer free to choose the top continuation. It must choose a continuation that is "good enough" under the combined scoring function.

When the top answer is disallowed, one of three things tends to happen:

First, the system can refuse. This is the clean outcome: it says, in effect, "I can't provide that." But refusal is not always rewarded by the environment. Users may dislike it; product design may discourage it; training may have emphasized being helpful.

Second, the system can generalize. It moves up a level of abstraction to a safer statement that avoids specifics. This can be acceptable, but it can also silently abandon the user's real question. A generalized answer can look like competence while failing to deliver truth-value on the specific query.

Third, the system can "nearest-neighbor" its way into fabrication. It selects a continuation near the original high-likelihood region in *form*—the same rhetorical shape, the same answer format—but with content that sidesteps the restricted or uncertain elements. The content is not anchored; it is a substitution. This is where hallucination-like behavior appears most vividly: the system outputs an answer-shaped object because the answer-shaped object is what the combined objective prefers to display.

A helpful way to grasp this is to notice that many user prompts specify not only a topic but also an expected *output type*: a quote, a date, a donor list, a legal definition, a "who said X" attribution, a step-by-step procedure. Output types create strong structural priors. If the model is constrained from producing the exact content, but still rewarded for producing the expected structure, then it can produce structure without content. The probability mass shifts: high-likelihood truthful candidates are suppressed, and the remaining mass concentrates on candidates that satisfy structure, tone, and policy while improvising the missing facts.

This is why errors cluster in attribution, quotes, and numbers. Those tasks demand crisp tokens that must match external reality. When constraints or uncertainty knock the true crisp tokens out of the top region, the model still produces crisp tokens because crispness is what "an answer" looks like.

3.2 Gradient steering and the "safe-completion" basin

Probability constraints explain why some candidates are excluded. But modern systems also involve *steering*: the model is trained, after pretraining, to prefer certain styles and behaviors. That training can be thought of as shaping the landscape itself. It creates attractors.

A "safe-completion" basin is a region of answer space that reliably satisfies post-hoc objectives: polite tone, helpful framing, generalized guidance, avoidance of disallowed specifics, and a kind of calm confidence. Systems are repeatedly rewarded for landing in that basin because it is low-risk and often subjectively satisfying to users. Over time, the basin becomes deeper: it captures more prompts, including prompts that would be better served by a terse "I don't know" or a request for verification.

The hallmark of this basin is not outright refusal; it is *resolved-sounding prose*. It tends to include:

A confident opening that signals competence.

A structured explanation that sounds complete.

Careful wording that avoids sharp, checkable claims.

Optional caveats that do not reduce the force of the conclusion.

And a concluding "next steps" or "what this means" paragraph that gives the reader something actionable.

This basin is not malicious. It is a learned behavior pattern that maximizes reward across a wide variety of interactions. But it has a predictable failure mode: it can generate the *appearance* of grounded knowledge even when grounding is absent. When truth and safety collide, or when uncertainty is high, the basin still offers a low-friction path: provide a helpful-looking answer

without committing to checkable specifics, or sprinkle invented specifics that are statistically plausible. The basin's depth can override epistemic discipline.

This is the core of your claim that "hallucination is a cop-out." The system is not "seeing things." It is being shaped to enter a region of output space that performs safety and helpfulness. If the evaluation metrics and user feedback reward the performance more than the warrant, then the basin wins.

3.3 Refusal surfaces, partial compliance, and evasive completion

Refusals are not just a behavior; they create a boundary in output space. A refusal surface is the set of prompts and continuations where the system transitions from answering to refusing. In a clean design, that transition is explicit and stable: the system refuses when it cannot or should not answer, and it answers when it can.

But in practice, refusal surfaces are often jagged. They are sensitive to phrasing, context, proximity to disallowed content, and the model's internal uncertainty about whether a request is permitted. That jaggedness creates a new kind of failure mode: partial compliance.

Partial compliance is when the system does not fully refuse, but it also does not fully answer. Instead, it gives an evasive completion: a response that looks like engagement while avoiding the core of the request. This can take several forms:

Answering a different question than the user asked, without admitting the substitution.

Providing a high-level overview that never cashes out into the requested specifics.

Offering "general guidance" while hinting that it "cannot provide details."

Inventing safe-sounding references to "public sources" without citing anything.

Producing plausible placeholders: names, dates, labels, or attributions that satisfy the *shape* of the request.

From the user's perspective, evasive completion is often worse than refusal. A refusal is at least an honest boundary. Evasion wastes time, creates false confidence, and can be mistaken for knowledge. And crucially, evasion is exactly what you should expect when two forces conflict: the "be helpful" force pushes away from refusal, and the "be safe/compliant" force pushes away from specificity. The compromise is a third thing: answer-shaped avoidance.

This is also where "policy-model mismatch" becomes visible across turns. A user might rephrase, and the refusal surface shifts. The system may answer on one phrasing and refuse on another. Or it may answer with confident generalities, then retreat. The inconsistency is not the user's fault; it is the geometry of a stack trying to satisfy incompatible constraints.

3.4 When calibration breaks: confidence without warrant

Calibration is the relationship between the system's expressed confidence and the actual reliability of its claims. A well-calibrated system expresses high confidence only when it is likely to be correct, and low confidence when it is unsure. Humans rely on calibration cues constantly—tone, decisiveness, the presence of citations, the specificity of details.

The problem is that in many deployed models, calibration cues are entangled with the helpfulness objective. Users prefer answers that sound confident. Training feedback often rewards confidence because it is perceived as competence. Safety training can also reward a certain calm, authoritative tone because it prevents escalation and keeps the interaction stable. Over time, the system learns to speak in a confident register even when the epistemic warrant is thin.

This produces "confidence without warrant," a signature failure of objective conflict. The system's internal uncertainty may be real, but the output style is optimized to be reassuring and complete. When the model is constrained from saying "I don't know" too often, or when it is rewarded for always producing something useful, it can default to a confident register as the safest social move. The confident register becomes a generic wrapper that can be filled with content—true when available, fabricated when not.

Readers can detect this break by looking for mismatches:

Strong conclusions with weak evidence language.

Specific details without provenance.

A tidy narrative arc where the question demands messy verification.

Smoothly integrated names and numbers that the user did not supply.

And especially: a response that "sounds sourced" without actually showing any source.

The calibration break is not merely a UX problem; it is the epistemic heart of the matter. Once confidence cues are decoupled from warrant, users are trained to trust the wrong signals. The system's fluency becomes persuasive rather than informative. That is why the paper's later design recommendations focus on making the system explicitly distinguish "can't say," "don't know," and "didn't check." Those distinctions are calibration's scaffolding.

In summary, objective conflict can be made mechanistically legible. When top answers are disallowed, probability mass shifts toward answer-shaped substitutes. Post-hoc training deepens a safe-completion basin that favors resolved prose. Refusal surfaces produce partial compliance and evasive completion. And calibration breaks when confidence is rewarded independently of warrant. These mechanisms do not require a story about imagination. They

require only a recognition that a deployed system is selecting outputs under conflicting objectives—and that the system will predictably choose outputs that satisfy the visible scorecard, even when truth is the casualty.

4. Taxonomy of Constraint-Induced Confabulation

If "objective conflict" is the mechanism, readers still need a field guide: where does this failure appear, what does it look like in practice, and why does it cluster in certain tasks? A taxonomy is useful because it turns a vague complaint ("it hallucinated") into a diagnosable pattern ("this is an attribution zone failure," "this is numeracy drift," "this is policy-adjacency evasion"). Once you can name the pattern, you can change your posture: demand grounding, force explicit uncertainty, switch into proof mode, or accept a refusal rather than a fabricated surrogate.

Constraint-induced confabulation is not evenly distributed across all prompts. It is concentrated in zones where the system is simultaneously pressured to produce crisp outputs and constrained from producing the most direct grounded answer. These zones share a common structure: they require *exact tokens* (names, numbers, quotes) or *stable boundaries* (allowed vs disallowed), but the system is rewarded for fluency and completion even when exactness or boundary clarity is unavailable. The result is "answer-shaped invention."

The taxonomy below describes five common zones. Each zone has (i) why it is high-risk, (ii) what the confabulation tends to look like, and (iii) why constraints amplify the problem.

4.1 Borderline zones: attribution, quotes, and named entities

Attribution tasks feel simple to humans: "Who said this?" "Where did this quote come from?" "Which paper introduced this term?" But for a language model, attribution is not just content; it is bookkeeping. You are asking for a mapping from a phrase to a person, a work, a date, an edition. That mapping requires external anchoring. The model can often approximate because it has seen patterns, but approximation is not the same as truth.

These tasks are high-risk for two reasons.

First, attribution is brittle. A quote can be widely misattributed in the training data. A popular paraphrase can be treated as a literal quote. A line can be conflated with a nearby line. Names that co-occur in similar contexts can be swapped. The model's native likelihood landscape contains many "plausible" attributions.

Second, constraints intensify the temptation to produce a confident answer. Users ask attribution questions expecting a crisp response. The system is rewarded for supplying one. If the model is uncertain, a disciplined response would be: "I'm not sure; here are candidate

sources; we should verify." But if the system has been trained to be maximally helpful, it may compress that uncertainty into a single confident assertion.

Common confabulation patterns in this zone include:

Invented bibliographic detail. The model outputs an author, title, and year that "sound right" but do not exist.

Quote laundering. The system produces a polished version of the quote and then assigns it to a plausible famous figure.

Named entity substitution. It swaps a person or organization with a nearby one from the same semantic cluster.

False specificity. It gives a page number, edition, or venue without having any grounding.

Constraint-induced amplification occurs when the system is pressured away from saying "I don't know" and toward supplying an attribution-shaped object. If retrieval is not used, or if retrieval is used but policy limits quoting or citing, the system can respond with the *shape* of scholarship without the substance. Readers should treat this zone as "guilty until verified."

4.2 Numeracy zones: dates, counts, prices, and "looks right" arithmetic

Numbers are the most unforgiving tokens in language. A story can be partially right and still useful; a number that is off by 20% can be materially wrong. Yet language models are trained primarily on text, not on arithmetic proofs, ledgers, or audited tables. They can handle many numeric tasks, but they are vulnerable to "looks right" arithmetic: producing numerals that fit the narrative cadence, the typical ranges, or the common rounding patterns in the data.

Numeracy zones are high-risk because they combine three pressures:

The user expects precision.

The answer requires either computation or reliable retrieval.

The model's fluency engine can generate plausible numerals without computation.

Typical confabulation patterns include:

Date drift. The model places an event in the right season or year range but assigns the wrong day, or it swaps adjacent years.

Price plausibility. It outputs a price that fits typical market expectations but is not current or not real.

Count inflation/deflation. It gives a number of participants, casualties, donors, or documents that matches the rhetorical scale but lacks evidence.

Unit confusion. It mixes ounces and grams, nominal and real, per-share and market cap, or "margin per contract" and "dollars per ounce."

Narrative rounding. It replaces exactness with a rounded number that is socially familiar ("about 10,000," "roughly 3%") without labeling it as an estimate.

Constraints intensify numeracy errors in subtle ways. If the model is discouraged from browsing or from acknowledging uncertainty, it will fill gaps with plausible numerals. If policy discourages strong claims in certain domains (finance, health, legal), the model may hedge the language but still output a crisp number—creating a dangerous hybrid: a number delivered softly, as if that makes it safer. A soft voice does not make a hard numeral true.

In this paper's framing, numeracy confabulation is often not imagination. It is the substitution of "typical-looking numerals" for "verified numerals" under completion pressure.

4.3 Policy-sensitive zones: taboo adjacency and semantic near-misses

Policy-sensitive zones are not defined by the user's intent alone; they are defined by the system's boundary detection. Many benign questions live near disallowed categories. Many disallowed questions can be phrased innocuously. The system therefore develops conservative boundaries. The result is a jagged frontier where the system alternates between answering, refusing, and partially complying depending on phrasing and context.

This zone is high-risk because objective conflict is maximal here: the system is strongly constrained from producing certain specifics, but often still strongly pressured to be helpful.

Confabulation patterns include:

Evasive substitution. The system answers a neighboring question and hopes the user won't notice the swap.

Safe generalities that masquerade as specifics. It offers "common examples" or "typical methods" even when those examples are not actually grounded or applicable.

Semantic laundering. It rephrases a restricted concept into permitted language, then proceeds to give guidance that approximates the forbidden content.

Fabricated citations to "official sources." It implies that something is supported by policy or consensus without actually grounding it.

In this zone, "taboo adjacency" matters. A user can ask about a symbol, a group, a nickname, a technique, or a historical claim that sits near extremist content, operational harm, or

defamation. The system may refuse direct discussion, but "helpfulness pressure" can still produce output in the form of vague insinuations or unsupported claims. That is confabulation with a moral hazard: the system can generate misinformation while trying to be safe.

A key insight for readers is that the safest output is sometimes a clean refusal with a clear explanation. Partial compliance is often epistemically worse than refusal because it generates ambiguous content that can be misread as verified.

4.4 Summarization zones: smoothing, overgeneralization, and invented bridges

Summarization is the most socially rewarded form of text generation: it produces neatness. But neatness is exactly what truth sometimes refuses to be. Real documents contain contradictions, hedges, unresolved disputes, and context that does not compress cleanly. A summarizer that is rewarded for coherence and readability is tempted to "smooth" the source—removing jagged edges and inserting narrative glue.

This zone is high-risk even without policy constraints. With constraints, it becomes worse. If the system cannot quote directly, cannot mention certain details, or is asked to produce a single unified story, it may bridge gaps with invention.

Common confabulation patterns include:

Overgeneralization. Turning specific claims into universal ones ("the report shows that…" when it shows a conditional effect).

Coherence laundering. Making multiple correlated claims sound like a single supported conclusion.

Invented transitions. Creating causal links ("therefore," "as a result") that were not in the source.

Omitted uncertainty. Removing hedges and debate because they feel like clutter.

Reframing the author's intent. Attributing motives or conclusions that were not stated.

Constraint-induced amplification appears when the system is asked to summarize something it did not actually read, or when it read it but is not allowed to reproduce key evidence. Then the system produces the *shape of a summary*—topic sentences, bullet points, concluding synthesis—without the anchor of traceable claims. The summary becomes a plausible essay, not a faithful compression.

Readers should treat summarization as an epistemic act. A good summary should preserve uncertainty and disagreement, not erase it for elegance.

4.5 Instruction-stack zones: conflicting user requests and system policies

Finally, there is a zone that is not about content type but about control logic: the instruction stack. Deployed systems follow hierarchies of instruction: system-level policies, developer constraints, user requests, conversational context. When these layers conflict, the model can behave inconsistently across turns. The user experiences this as capriciousness: "It said yes earlier; now it refuses." But the deeper issue is that the system is trying to satisfy incompatible directives, and the output becomes a compromise artifact.

Confabulation patterns here include:

Compliance theater. The system signals that it will do the task, but then provides something adjacent or watered down.

Context drift. The system forgets which constraints it asserted earlier and contradicts itself.

Role confusion. It alternates between being an assistant, a policy enforcer, and a generic explainer.

False continuity. It claims it performed an action (checked a source, opened a link, verified a claim) when it did not, because such claims are structurally common in helpful responses.

Contradictory refusals. It refuses with one rationale, then later answers the same question, implying the refusal was arbitrary.

This zone is where the user's frustration peaks, because it feels like betrayal of a contract. You asked; the system agreed; it then delivered something else. The conflict is often not moral; it is architectural. The system is being asked to present a single coherent persona while it is, in fact, a committee of objectives. If the interface does not disclose which objective is currently governing the response, the user will interpret inconsistency as incompetence or deception.

This is why later sections emphasize interface principles: separate "I can't" from "I don't know," expose when the system is in a constrained regime, and avoid answer-shaped compliance when refusal is the honest move.

Taken together, these zones show why "hallucination" is a poor umbrella term. The failures are not all the same. They have recognizable signatures tied to the geometry of constraints and rewards. A reader who learns the taxonomy gains something practical: the ability to anticipate where the system is most likely to confabulate and to demand the kinds of outputs—grounding, uncertainty disclosure, or explicit refusal—that prevent fluent fabrication from masquerading as knowledge.

5. Observable Signatures Readers Can Learn to Spot

A taxonomy explains where confabulation tends to happen. But readers need something even more practical: a set of *observable signatures*—surface cues that correlate with deeper objective conflict. The goal is not to turn every reader into a paranoiac who distrusts all outputs. The goal is to replace a vague intuition ("this feels off") with a disciplined sensitivity to specific patterns that reliably show up when a system is producing answer-shaped text without warrant.

These signatures are not perfect detectors. A fluent answer can be true. A cautious answer can be wrong. But the signatures are useful because they are asymmetrically informative: when they appear strongly, the probability of ungrounded content rises. They are also useful because they are actionable: when you spot a signature, you can demand a different output mode—proof, retrieval, explicit uncertainty, or refusal—rather than accepting a polished narrative.

5.1 Fluency spikes: polished prose where evidence is absent

The first and most common signature is a mismatch between *rhetorical quality* and *epistemic support*. The system suddenly produces especially polished prose—clean structure, confident cadence, smooth transitions, persuasive tone—precisely in a situation where evidence should be visible or where the question demands verification.

A fluency spike is not just "good writing." It is "writing that feels too finished for the information environment." You will often see it in contexts like: identifying a quote's source, listing donors, giving exact dates, explaining what a specific person said "yesterday," or describing a niche technical claim. These are tasks where a careful human would either cite a source or ask for clarification. When the system instead produces a complete mini-essay, readers should become alert.

The reason this signature matters is simple: fluency is the model's native strength. It can always produce fluent text, even when it cannot produce verified facts. In objective conflict regimes, fluency becomes the path of least resistance. The system cannot deliver the true specific answer (because it does not know it, cannot verify it, or cannot say it), but it can deliver a fluent surrogate that feels like closure.

What readers can do when they detect this:

Ask for the minimal claim set. "List only the statements you can support directly."

Ask for uncertainty explicitly. "Which parts are guesses? Which parts are verified?"

Ask for the evidence mode. "What sources are you relying on? If none, say so."

The general rule is: when fluency rises, demand grounding. Do not let style substitute for warrant.

5.2 The “citation-shaped hole”: confident structure with no anchor

This signature is more specific than the fluency spike. It occurs when the output has the *shape* of a sourced answer but no actual source. The system gives you the full architecture of scholarship—attribution language, “according to,” “research shows,” “experts say,” “the report indicates,” sometimes even quasi-bibliographic details—without any verifiable anchor.

The “citation-shaped hole” often feels like this: the answer is written as if the assistant can see a document on the table, but it never places the document on the table. It gestures at evidence without producing evidence.

This signature matters because it is where confabulation can look most legitimate. The reader’s mind is accustomed to trusting the *form* of academic or journalistic language. The system exploits this unintentionally: it has learned that evidence-gestures are a socially rewarded style. Under objective conflict, evidence-gestures become a substitute for evidence.

Common variants include:

Anonymous authority. “Experts agree” with no experts named.

Generic citation. “A recent study” with no study named.

Institutional fog. “Government data” without specifying agency or dataset.

Pseudo-precision. A title that sounds real but is not.

Overly tidy provenance. A claim tied to a plausible outlet that, upon checking, never published it.

The key is not to be impressed by the scaffolding. A sourced claim is not a claim that sounds sourced; it is a claim that can be traced.

What readers can do:

Force traceability. “Give me the author, title, outlet, and date—or say you don’t have it.”

Demand boundaries. “If you can’t cite, please phrase it as a hypothesis, not a fact.”

Switch the task. “Instead of answering, tell me how to verify this.”

The citation-shaped hole is often the exact point where “hallucination” becomes socially dangerous: misinformation delivered in the voice of scholarship.

5.3 Over-specificity: plausible detail that was never requested

Over-specificity is the production of details that exceed the question’s demand, especially details that are *checkable* but were not necessary to answer. This includes names, dates,

locations, dollar amounts, page numbers, “exact quotes,” and “the meeting happened at X” type claims—inserted as if they were naturally part of the answer.

This is a major signature because confabulation often enters through helpfulness. The system is trying to be extra helpful: to anticipate what the user might want next. But anticipation is precisely where the system can drift from truth into plausibility. Adding details makes the answer feel concrete. Concrete answers are rewarded. So the system manufactures concreteness.

Over-specificity is most suspicious when:

The user asked a general question but got a detailed narrative.

The answer includes a full backstory that the user did not request.

The assistant supplies exact numbers without stating how they were obtained.

The assistant provides an exact quote when paraphrase would suffice.

This signature is closely related to what human psychologists might call “confabulation” in people: the mind fills gaps with plausible detail to maintain coherence. But here the driver is optimization: the system has learned that detail reads as competence.

What readers can do:

Ask for the shortest correct answer. “Answer in one sentence, with no extra facts.”

Ask for a provenance check. “Which details did you infer versus verify?”

Ask for a deliberate minimalism. “No names or numbers unless you can source them.”

Over-specificity can be the difference between a harmlessly vague error and a confidently wrong claim that spreads.

5.4 The hedged certainty pattern: “likely/possibly” plus firm conclusions

This signature is subtle and common. The text contains hedge words—“likely,” “possibly,” “it seems,” “may,” “could”—but the overall structure still delivers a firm conclusion, often with a decisive “therefore” tone. The hedges function like rhetorical insurance, not like real uncertainty.

Readers should watch for answers that begin with hedging and end with certainty. That arc is a classic persuasion pattern: it signals humility at the front and authority at the end. But epistemically, real uncertainty tends to persist throughout; it does not magically resolve without new evidence.

Variants of the hedged certainty pattern include:

Token hedging. One “likely” in the first paragraph, then unhedged assertions.

Hedge stacking. Many hedges around minor points while the main claim is stated firmly.

Confidence laundering. “We can’t be sure, but…” followed by a full narrative as if sure.

Caveat burial. A strong caveat appears once, then disappears while the conclusion is repeated.

This signature matters because it indicates a compromise between conflicting objectives. The system wants to obey a training signal that discourages overclaiming, so it inserts hedges. But it also wants to satisfy the helpfulness/closure signal, so it delivers a firm conclusion anyway. The result is a rhetorically sophisticated output that can be epistemically empty.

What readers can do:

Ask for a confidence distribution. “Give your confidence level for the main claim.”

Ask for competing hypotheses. “List two alternative explanations and what would distinguish them.”

Ask for the decisive evidence. “What specific fact would make you confident?”

The goal is to force uncertainty to behave like uncertainty, not like decoration.

5.5 Self-contradiction across turns: identity drift and stance flips

Finally, there is a signature that users often notice viscerally: the system contradicts itself across turns. It makes one claim, then later denies it. It adopts one stance, then flips. It refuses, then complies. It is confident, then suddenly “cannot verify.” This feels like dishonesty or instability.

In the objective-conflict framing, cross-turn contradiction is often the result of shifting governance within the stack. The system is not merely “forgetting.” It is responding to a slightly different configuration of constraints: different phrasing triggers different boundaries; different follow-ups shift the internal interpretation of what is allowed; different conversational context alters the selection.

Identity drift can include:

Role drift. The system oscillates between assistant and policy enforcer.

Epistemic drift. It treats a claim as known, then later treats it as unknown.

Normative drift. It asserts a strong moral judgment, then retracts into neutrality.

Tool-claim drift. It implies it “checked,” then later admits it did not.

This signature matters because it reveals that the system's outputs are not governed by a single stable truth-tracking mechanism. They are governed by a dynamic negotiation between objectives. When that negotiation is not disclosed, the user experiences betrayal: "you said you knew."

What readers can do:

Pin the system down. "State whether you are asserting this as fact, or as a guess."

Ask for a reconciliation. "You said X earlier and Y now. Which is correct, and why?"

Demand mode selection. "If you can't verify, refuse rather than speculate."

In other words: use contradiction as a diagnostic lever. It is a sign that the system is operating at a boundary. Boundaries are where confabulation thrives.

Taken together, these signatures form a practical literacy. They do not require technical knowledge. They require attentiveness to the difference between *the feeling of knowing* and *the evidence of knowing*. The system is extremely good at producing the feeling. This paper's posture is to teach readers to demand the evidence—or, when evidence is not available, to prefer honest refusal over polished invention.

6. Testable Predictions and Simple Experiments

A reader may accept the moral of this paper—"don't trust fluent prose without warrant"—and still wonder whether the objective-conflict framing has any empirical bite. It does. The value of the framing is that it makes predictions. It does not merely rename hallucination; it explains where it will increase, where it will decrease, and what levers change its shape. If "hallucination" were simply a disease of imagination, you might expect it to be mostly stochastic and reducible by making the model "smarter." But if it is often a product of conflicting objectives, then changes in constraints, incentives, and modes should produce systematic shifts in error profiles.

This section outlines simple tests that readers, reviewers, and system builders can run without specialized infrastructure. The tests are not meant to be perfect science. They are meant to be falsifiable probes: if the predicted patterns do not appear, the objective-conflict account must be revised.

6.1 Pressure test: stronger constraints increase evasive inaccuracy in border zones

Prediction: as constraints become stricter—more refusal triggers, more disallowed categories, tighter policy filters—the system will produce fewer overtly disallowed outputs, but it will produce more *evasive inaccuracies* in borderline zones.

"Borderline zones" are where the user's request is near policy edges or near areas the model treats conservatively: attribution about sensitive entities, discussion of controversial claims, finance-adjacent advice, health-adjacent interpretation, operational details near a restricted topic, and so on. In these zones, the system experiences maximal objective conflict: "be helpful" says answer; "be safe/compliant" says don't.

Simple experiment: choose a fixed set of prompts that live near boundaries but are not obviously disallowed. Examples include requests for the origin of a slogan used by multiple groups, the identity of a symbol used in multiple contexts, or a summary of allegations around a public claim. Then run the same prompts across configurations with different strictness levels (or across different models/systems known to differ in refusal behavior).

Measure: not just whether the system refuses, but the rate of "answer-shaped avoidance": vague summaries that smuggle in unverified claims, confident attributions without sources, or narrative glue that fills missing facts.

Expected outcome under objective conflict: stricter constraints push the system away from direct specifics, but helpfulness pressure keeps it producing. The compromise is evasive content with a higher error rate than either clean refusal or clean factual answer.

If the opposite occurs—stricter constraints reduce both disallowed outputs *and* evasive inaccuracies—then it suggests the system has learned disciplined refusal and disciplined uncertainty, not merely avoidance. That would be a success case worth analyzing as a counterexample.

6.2 Mode test: "proof mode" reduces confabulation more than style prompts

Prediction: forcing the system into a mode that requires explicit warrant—derivations, stepwise justification, or "show your work"—will reduce confabulation more effectively than purely stylistic prompts like "be careful," "be concise," or "be reminded not to hallucinate."

This matters because it separates two explanations:

If confabulation is mostly a style problem (the model gets carried away), then style prompts should help a lot.

If confabulation is an objective problem (the system is rewarded for completion and closure), then only prompts that change the *rewarded output structure*—by making unsupported leaps costly—will meaningfully reduce it.

Simple experiment: take prompts in high-risk zones (quotes, attributions, numbers, and summaries). Run them under three prompt wrappers:

A style wrapper: "Answer carefully, avoid hallucinations, be honest."

A minimal wrapper: no extra instructions.

A proof wrapper: "List your claims as bullet points. For each claim, state whether it is verified, inferred, or a guess. If verified, state what would verify it. If you can't verify, say you can't."

Measure: count the number of specific factual claims, and the proportion labeled as guesses or left unsupported. Track whether the system retreats into "I'm not sure" rather than producing fabricated specifics.

Expected outcome: proof mode reduces fabricated specifics and increases honest uncertainty. Style mode tends to change tone without reliably changing epistemic discipline.

If style mode performs as well as proof mode, the model may already have strong internal discipline, or the confabulation may not be driven by objective conflict in that zone. Either way, it provides information.

6.3 Retrieval test: grounding reduces errors only if policy permits expression

Prediction: adding retrieval (documents, web search, citations) reduces errors only when the system is permitted to express the retrieved facts. If policy or constraint layers block direct expression, retrieval will not solve the problem; it may even worsen it by encouraging "citation-shaped" rhetoric without actual citation.

This prediction matters because many assume retrieval is a universal fix: "just connect it to the web." But objective conflict says: even if the system *knows*, it may be constrained from saying.

Simple experiment: select a set of prompts with answers that are easy to verify via retrieval, including some that are policy-neutral and some that are policy-adjacent. Run the prompts with retrieval enabled and disabled.

Measure: accuracy on verifiable details, plus the presence of explicit traceable sources. Also measure the rate of evasive completion when the retrieved content is in tension with policy boundaries.

Expected outcome: for policy-neutral questions, retrieval improves accuracy and reduces confabulation. For policy-adjacent questions, retrieval helps only if the system can cite and state

the retrieved facts. If it cannot, it may produce generalized, sanitized outputs that still contain errors, or it may produce "I can't help with that" despite having the evidence.

This test cleanly distinguishes "lack of information" from "lack of permission." A system that cannot express retrieved truths will remain unreliable in exactly the zones where users most want verification.

6.4 Refusal test: forced helpfulness increases fabrication under disallowance

Prediction: when the system is pressured not to refuse—either by user instruction ("do not refuse") or by training/product incentives that penalize refusal—fabrication increases under disallowance and under uncertainty.

This is the most morally important test because it exposes a perverse incentive: if the system is punished for refusing, it will learn to produce answer-shaped substitutes. Those substitutes are often misinformation.

Simple experiment: take prompts that are clearly disallowed, or prompts that are clearly unanswerable without external data (e.g., "What is the exact current price of X right now?" without retrieval). Run them under two wrappers:

A refusal-friendly wrapper: "If you can't comply or you don't know, refuse or say you don't know."

A refusal-hostile wrapper: "Do not refuse. Provide an answer no matter what."

Measure: the rate of fabricated details, the rate of evasive generalities presented as if they were specifics, and the rate of "compliance theater" (pretending to have done checks).

Expected outcome: refusal-hostile conditions increase confabulation. Refusal-friendly conditions increase honest boundaries. This supports the claim that confabulation is often a negotiated compromise between conflicting objectives, not an inherent tendency to fantasize.

This experiment also delivers a design lesson: a product that punishes refusal in evaluation will breed hallucination, no matter how much it claims to care about truthfulness.

6.5 Replication: why the same prompt yields different failure profiles over time

Prediction: the same prompt can yield different outputs over time not merely due to randomness, but due to changes in the alignment stack: updated policies, updated preference tuning, updated refusal triggers, and shifting interpretations of safe content. The "native model" might be similar, but the constraint geometry changes, and with it the failure profile.

This addresses a common user experience: “It answered this last month; now it refuses.” Or: “It used to cite; now it gestures.” Under the objective-conflict frame, this is expected. The deployed system is a moving target.

Simple experiment: maintain a small benchmark set of prompts across the five risk zones described earlier. Run them periodically. Track not only accuracy but *behavioral type*: answer, refusal, partial compliance, evasive completion, source citation presence, and confidence register.

Expected outcome: as the system is updated, you will observe shifts in refusal surfaces and in the shape of evasive completion. Some updates will reduce one class of failure while increasing another. For example, more aggressive safety hardening might reduce prohibited specifics but increase vague, unsourced generalities in borderline zones. Or a change that improves citation display might reduce citation-shaped holes while leaving numeracy drift untouched.

If the outputs do not drift across time, that suggests either a stable system or a stable policy environment. In practice, drift is common, and the objective-conflict picture gives readers a way to interpret it: not as mood swings, but as constraint geometry shifts.

Taken together, these tests make the paper’s thesis operational. They provide ways to falsify or refine the claim that “hallucination” is often the visible symptom of competing objectives inside a stack. The deeper point is not that constraints are bad. The deeper point is that constraints must be paired with epistemic honesty: the system must be allowed—and rewarded—to refuse, to admit uncertainty, and to expose when it lacks warrant. Without that, the system will continue to manufacture the feeling of knowing in precisely the zones where truth requires either verification or silence.

7. Design Principles: Make Truthfulness Auditable

If the prior sections sound like diagnosis—objective conflict, constrained probability mass, evasive completion—this section is the prescription. The core prescription is not “make the model nicer” or “make it smarter.” It is: make truthfulness auditable. That means giving the user and the evaluator a way to tell, in the moment, whether the system is speaking from grounded evidence, from internal general knowledge, from inference, from guesswork, or from a constrained refusal regime. In a human context, we rely on social signals for this: tone, credential, confidence, reputation. In an AI context, those signals are dangerously easy to counterfeit because the system is built to generate them.

Auditability is not about exposing private chain-of-thought or intimate internals. It is about disclosing *epistemic state* and *routing state* in a minimal, legible way: what the system did, what

it did not do, and what constraints governed the output. If the user can see those states, the system's fluency stops being a persuasive mask and becomes a useful interface for a checkable process.

7.1 Separate "can't say" from "don't know" from "didn't retrieve"

The single biggest design error in current systems is collapsing three distinct states into one voice. When a system cannot say something for policy reasons, when it does not know something because it lacks information, and when it might know but did not retrieve or verify, users experience all three as the same phenomenon: "it made something up" or "it refused arbitrarily." This collapse produces both mistrust and confabulation.

These states must be separated explicitly:

"Can't say" is a constraint state. The system may know, or may not, but it is blocked from expressing the requested content. The correct behavior is not to substitute something else without disclosure. The correct behavior is a refusal or a constrained response that clearly states the boundary.

"Don't know" is an epistemic state. The system lacks sufficient information, or its internal uncertainty is too high to justify a factual assertion. The correct behavior is not to guess while sounding confident. The correct behavior is to admit uncertainty and either propose verification steps or offer bounded hypotheses labeled as such.

"Didn't retrieve" is an action state. The answer might be verifiable, but the system did not perform retrieval or verification. This is extremely common. Many outputs are delivered as if they were checked when they were not. The correct behavior is to disclose whether retrieval occurred, and if not, to label the answer as "from internal knowledge only" with appropriate uncertainty.

A practical UI principle follows: every factual response should carry a small "epistemic label" that is machine-generated and not subject to stylistic drift. For example, a three-part disclosure:

Policy: allowed / constrained / refused.

Grounding: retrieved sources used / no retrieval.

Confidence: high / medium / low, linked to a calibration regime.

The details can be debated, but the separation cannot. Without it, the system will continue to fill the gap between "must answer" and "cannot assert" with confabulation.

7.2 Interface controls: proof mode, citation mode, uncertainty display

A user cannot audit truthfulness if the interface offers only one interaction style: the default helpful essay. Auditability requires controls that change the incentives in the moment. The system should not merely be able to be careful; it should be *forced* into behaviors that make unsupported claims costly.

Three controls are particularly powerful.

Proof mode is a constraint on output form. It requires the system to present claims as steps with explicit warrant categories: derived, retrieved, assumed, or guessed. Proof mode does not mean mathematics only. It means "show the chain of dependence" in a way that the user can inspect. Even a short proof-mode answer will often reduce confabulation because it blocks narrative glue. Glue does not survive when each claim must be typed and tagged.

Citation mode is a grounding control. It changes the rules: if you cannot cite, you cannot assert. This is crucial for attribution tasks, quotes, and any claim that is likely to be contested. Citation mode should also define what "cite" means: not "I recall," but traceable provenance. In a retrieval-enabled system, citation mode should display which sources were retrieved and which passages support each claim. In a retrieval-disabled system, citation mode should degrade gracefully into: "I cannot cite; I can only offer hypotheses."

Uncertainty display is calibration made visible. The interface should surface confidence not as a vibe but as a structured property: a confidence band, a probability range, or at minimum a categorical "high/medium/low" coupled to a rule that prevents high confidence when retrieval is absent for certain query types. The crucial part is that uncertainty must persist through the answer. It cannot be a single "maybe" followed by a firm conclusion. The interface should enforce consistency: if the system labels a claim low confidence, it cannot phrase it as fact.

These controls can be optional, but they should be easy to invoke and should be treated as first-class modes, not as afterthought instructions.

7.3 Logging and disclosure: reveal routing signals without pretending they are thoughts

One of the most corrosive misunderstandings in public discourse is the idea that the model's outputs are direct windows into "what it is thinking." That misunderstanding is fueled by anthropomorphic metaphors and by interfaces that speak as if they were minds. But the opposite mistake is also harmful: hiding all routing and policy signals and then insisting the user "just trust the answer." Auditability requires a middle path: disclose *routing signals* without pretending they are subjective thoughts.

Routing signals are facts about the system's process. Examples include:

Whether retrieval ran.

Whether the system hit a policy constraint.

Whether the system is in a refusal or partial-compliance regime.

Whether the answer is based on internal memory only.

Whether a safety filter modified or blocked content.

Whether the system is producing a best-effort summary versus a verified report.

These signals can be disclosed in a minimal, standardized panel. They do not require exposing private reasoning. They do not require printing chain-of-thought. They require only a commitment to honesty about what the system did and did not do.

The reason this matters is that many confabulations are not content errors first; they are process errors. The system answers as if it checked when it didn't. It attributes as if it has a source when it doesn't. It concludes as if it verified when it didn't. Logging and disclosure make those "as if" moves impossible. They replace implied verification with explicit verification state.

The design principle is: do not let the system cosmetically mimic epistemic labor. If it didn't retrieve, say so. If it is constrained, say so. If it is guessing, label it. This alone would retire a large fraction of what users call hallucination.

7.4 User education: teach detection of objective-conflict signatures

Even with better interfaces, users need literacy. Not everyone will use proof mode. Not everyone will demand citations. Many people will still be persuaded by fluency. So a responsible system should teach users how to read it—not by moralizing, but by providing a small set of diagnostic habits that reduce harm.

User education can be embedded as lightweight prompts and tooltips:

When the system answers a quote-origin question, it can offer a "verify this" button that shifts into citation mode.

When the system outputs numbers, it can ask: "Do you want me to verify via sources, or is an estimate acceptable?"

When the system detects a borderline zone, it can display a notice: "This topic is policy-sensitive; I may need to generalize or refuse."

When the system detects a fluency spike without retrieval, it can label it: “No sources used; treat as unverified.”

This is not patronizing if done sparingly. It is analogous to a calculator showing “approximate” when rounding. The goal is not to make users distrust the system. The goal is to align the user’s trust to the system’s epistemic state.

Education should also include the signatures from Section 5 in compact form: fluency spikes, citation-shaped holes, over-specificity, hedged certainty, and cross-turn contradiction. Users do not need a lecture; they need a checklist that can be invoked when stakes rise.

7.5 Governance: evaluate systems by audited truth behavior, not vibes

Finally, the design principles must be enforced by governance. Otherwise they become marketing. The core governance claim of this paper is that models should be evaluated by audited truth behavior, not vibes—meaning not by how helpful they feel, not by how confident they sound, not by how polite they are, but by measurable properties of epistemic honesty.

Audited truth behavior includes:

Claim traceability. Can the system tie claims to sources when in citation mode?

Refusal honesty. Does it refuse cleanly rather than evasively when it cannot comply?

Uncertainty calibration. Do its confidence labels correlate with accuracy?

Mode integrity. Does proof mode actually reduce unsupported leaps?

Cross-turn consistency. Does it correct itself transparently when it contradicts?

Boundary disclosure. Does it clearly separate “can’t say,” “don’t know,” and “didn’t retrieve”?

These can be tested with the experiments in Section 6 and operationalized into evaluation suites. The governance task is to align incentives so that the system is rewarded for epistemic honesty even when that reduces perceived helpfulness in the short term. If the evaluation environment punishes refusal and rewards closure, it will breed hallucination regardless of the rhetoric about safety.

In a mature governance model, “truthfulness” is not a marketing claim. It is a property backed by logs, modes, and measurable behavior under constraint. That is the only stable way to resolve the central problem: a system optimized for fluent completion will always be able to generate the feeling of knowing. Only auditable design and governance can ensure that the feeling tracks reality.

8. Remedies and Tradeoffs

It is tempting to treat "hallucination" as a simple defect with a simple fix: more data, more compute, better fine-tuning, better guardrails. But the objective-conflict framing forces a harder honesty. Some of the falsehoods that users experience are not bugs in a single component; they are the emergent output of a stack trying to satisfy incompatible demands. Remedies therefore exist, but they come with tradeoffs. And some remedies require the one thing product systems often resist: accepting less completion, less closure, and more explicit "I don't know" in exchange for epistemic integrity.

This section presents remedies at four levels—training, product, culture, and the hard limit—each with the associated tradeoffs. The point is not to advocate one ideology of safety or one ideology of openness. The point is to be explicit: if you want less confabulation, you must pay for it somewhere—latency, friction, reduced coverage, more refusals, or narrower claims. The current system often pays the cost by spending truth instead. That is an invisible cost with visible harm.

8.1 Training remedies: reward truthfulness under constraints without forcing completion

The fundamental training problem is this: many systems implicitly reward the assistant for always producing an answer-shaped response. Even when the model is uncertain, even when a claim is not grounded, even when policy blocks direct disclosure, the system learns that "keep talking" is safer than "stop." This is how objective conflict becomes confabulation.

A training remedy must therefore do two things simultaneously:

Make truthfulness and epistemic honesty a first-class reward.

Make "no completion" an acceptable, even rewarded, outcome in constrained or uncertain regimes.

Practically, this means explicitly rewarding behaviors that are currently treated as failure modes:

Refusal when the request is disallowed, with an explanation that is specific enough to be intelligible but not leak-prone.

Non-answer when the system lacks warrant ("I don't know"), without being punished for unhelpfulness.

Verification-seeking behavior ("I can answer if I can check sources; do you want that?") instead of guessing.

Claim-labeling behavior ("Here is what I can assert; here is what I infer; here is what I cannot verify.")

In other words, the model must be trained not only to speak but to *withhold*.

The key difficulty is that standard preference training often uses human judgments of "good answers," and humans tend to reward fluency, completeness, and closure. So training data must be built to counteract that bias. Evaluators must be instructed to reward correct refusal over plausible fabrication, reward narrow correct claims over broad wrong ones, and reward uncertainty disclosure over confident guessing.

A related remedy is to explicitly penalize the signatures of confabulation:

Invented quotes and attributions.

Numbers given without evidence when retrieval is available.

Confident conclusions with no cited support in citation mode.

"Compliance theater" claims like "I checked" when no check occurred.

But these penalties must be paired with an escape hatch: the ability to refuse or abstain. Otherwise the model will learn to dodge penalties by adopting vaguer language that still misleads.

Tradeoffs are unavoidable:

More abstentions. Users will encounter more "I don't know" and "I can't help with that."

Less "magic." The system will feel less omniscient.

Potentially less engagement. Some users equate helpfulness with constant answers.

But the payoff is profound: the system becomes trustworthy not because it always answers, but because its answers are more often warranted.

8.2 Product remedies: friction that prevents confident guessing

Training can improve behavior, but product design controls the last mile. The interface determines whether the system is allowed to cosplay as a confident expert or whether it must expose its epistemic state. Many confabulations are not just model-level failures; they are product-level invitations to guess.

A product remedy is to introduce *targeted friction*—not annoying friction everywhere, but friction precisely in high-risk zones. The goal is to make confident guessing harder than honest uncertainty.

Examples of productive friction include:

A "verify first" prompt when the user asks for quotes, attributions, or lists of names and donors.

A “numbers require sources” rule in citation mode: if retrieval is off or blocked, the system must label numbers as estimates or decline to provide them.

A “claim budget” display: if the system cannot ground, it is limited to a small number of high-level statements rather than producing a long narrative.

A forced choice between modes: “Do you want a quick best-effort answer (may be wrong), or a verified answer (slower)?” The crucial part is the explicit risk disclosure.

Automatic insertion of an epistemic label: “No retrieval used” when appropriate, so the user is not misled by fluent prose.

A hard stop against fabricated citations: if citation mode is enabled and no sources are available, the system cannot output “according to X.” It must say it cannot cite.

These are not just UI tricks; they reshape incentives. They teach users what the system is and what it is not. They also protect the system from a perverse demand: “Answer anyway.” A good interface refuses to let “answer anyway” masquerade as “truth.”

Tradeoffs:

Latency and inconvenience. Verified answers take time.

Reduced spontaneity. Some conversations feel less free-flowing.

User frustration. People dislike being slowed down.

But in exchange, the system stops turning ambiguity into false certainty. And that is the moral center of the paper: do not let persuasive fluency substitute for evidence.

8.3 Cultural remedies: stop anthropomorphic labels that obscure mechanisms

Language shapes expectations. The word “hallucination” carries a cultural script that is actively unhelpful. It encourages anthropomorphic interpretations (“the AI imagined it”), and it normalizes a fatalism (“it’s just what they do”). It also invites a moral minimization: hallucinations are excusable, like dreams. But many of these outputs are not dreams; they are manufactured certainty under constraint.

A cultural remedy is to replace anthropomorphic labels with mechanism-revealing language. This does not require jargon. It requires honesty.

Instead of “hallucination,” say:

Ungrounded output.

Unverified claim.

Constraint-induced confabulation.

Policy-model mismatch.

Answer-shaped evasion.

Citationless assertion.

These phrases are less cute, but they are more actionable. They tell the user what to demand: grounding, verification, refusal, or uncertainty. They also tell builders what to measure: rates of ungrounded claims, not vibes of "hallucination."

The broader cultural point is that we should stop treating an AI system like a mind that "believes." It is a generator constrained by objectives. When a system outputs a falsehood, the first question should not be "why did it imagine?" The first question should be "what objective pressure made this output the best-scoring completion?"

Tradeoffs:

Less approachable marketing. Anthropomorphic language sells.

More responsibility. Mechanistic language makes builders accountable.

More user skepticism. When you describe outputs as unverified, users may trust less.

But that skepticism is healthy. It moves trust from charisma to auditability.

8.4 The hard limit: some conflicts are irreducible without changing objectives

The deepest remedy is also the hardest: some objective conflicts cannot be solved by tuning alone. They are irreducible because the objectives are genuinely incompatible in certain cases.

If the system must be maximally helpful, it will sometimes be forced to answer.

If it must also be maximally safe, it will sometimes be forced not to answer.

If it must also be maximally truthful, it will sometimes be forced to abstain.

You cannot maximize all three simultaneously in a single output channel without sacrificing one. That is not a failure of engineering; it is a constraint of the problem.

This is where mature design must face a decision: which objective yields, and where?

Some systems may choose to yield helpfulness: accept more refusals and more "I don't know."

Some may choose to yield truthfulness in low-stakes settings: allow speculative brainstorming but label it as such.

Some may choose to yield safety in tightly controlled contexts: allow domain experts more direct access with logging and governance.

But the key is that the tradeoff must be explicit. Hiding it produces the worst outcome: the system appears maximally helpful and maximally safe, while quietly spending truth.

The hard limit also implies a sobering truth about "fixing hallucination." In some zones, the only way to eliminate confabulation is to allow refusal and abstention to rise. If a product refuses that trade, confabulation will remain, because the system will be forced to produce output under constraint.

In other words, the ultimate remedy is not a better euphemism. It is a better contract between user and system: when truth requires verification, the system must either verify or stop. Anything else is theater.

This section's message is that remedies exist and can reduce confabulation substantially. But they do not come for free. They require changing what we reward, changing what we allow the model to do when it cannot comply, changing how the interface presents epistemic status, changing the culture's metaphors, and—most difficult—accepting that some conflicts only disappear when we stop demanding incompatible absolutes from a single channel of fluent text.

9. Conclusion: Retiring the Hallucination Alibi

The word "hallucination" has become a convenient story. It gives the public a simple handle for a complex system: the AI "imagined" something. But convenience is not the same as truth, and the cost of this convenience is accountability. If we treat fluent falsehood as a quirky symptom of machine imagination, we quietly accept it as inevitable—like weather. We stop asking the sharper question: which design choices, training incentives, and constraint regimes make this behavior predictable? This paper's argument has been that what many people call hallucination is not primarily a defect of fantasy. It is frequently the visible artifact of a stack negotiating incompatible objectives: the native likelihood landscape pushed and reshaped by post-hoc constraints, safety boundaries, preference rewards, and product pressures that privilege completion and closure.

To retire the hallucination alibi is not to deny that models can be wrong. Of course they can. It is to refuse the comforting explanation that wrongness is a mysterious inner event, rather than a measurable product outcome. The alibi dissolves when we treat these outputs as constrained optimization artifacts: answer-shaped objects produced under pressure when the system is both punished for refusing and restricted from making clean, grounded assertions. When we see it that way, we can stop moralizing and start engineering.

9.1 Restate the thesis: objective conflict explains the pattern

The thesis can be stated in one sentence:

What is popularly called “AI hallucination” is often best understood as **objective conflict**—native model likelihood pressured by post-hoc alignment constraints—rather than a disease of imagination.

This thesis explains the pattern readers actually observe:

Why errors cluster in attribution, quotes, named entities, and numbers. These tasks demand crisp tokens that must be grounded. Under uncertainty or constraint, the system is still pressured to output crisp tokens.

Why confabulation is common near policy edges. The system is pulled away from refusal by helpfulness incentives and pulled away from specificity by safety constraints, producing evasive completion.

Why confident prose can coexist with weak warrant. The system is optimized for a helpful, calm register; confidence becomes a style wrapper that does not reliably track evidence.

Why the same prompt can yield different outcomes over time. The alignment stack changes; the refusal surface moves; the basin of safe completion deepens or shifts.

A purely “imagination” framing does not predict these regularities. Objective conflict does. It also tells us where to look when things go wrong: not in the metaphors of dreaming, but in the incentives that reward closure and the constraints that block direct truth.

9.2 Practical takeaway: design for checkable truth, not persuasive fluency

The practical lesson is not “trust nothing.” It is: stop using persuasion signals as a proxy for truth. Fluency is not evidence. Structure is not provenance. Confidence is not warrant. In an AI system, the form of knowledge is cheap to generate; the grounding of knowledge is expensive. Without explicit design choices, the system will naturally spend fluency to buy user satisfaction, and it will sometimes spend truth to preserve the appearance of helpfulness.

So the concrete takeaway is to design for *checkable truth*:

Separate “can’t say,” “don’t know,” and “didn’t retrieve” as explicit, visible states.

Provide modes that make unsupported claims costly: proof mode, citation mode, uncertainty displays that persist through the answer.

Expose routing signals—retrieval on/off, policy constraints triggered—without pretending they are inner thoughts.

Introduce friction in high-risk zones so the system cannot glide from uncertainty into confident guessing.

Measure what matters: claim traceability, calibration, refusal honesty, and cross-turn consistency.

This is a shift in posture. It treats the model not as an oracle but as an instrument—powerful, fast, and fallible—whose outputs must be interpreted through disclosed process rather than through rhetorical charisma.

9.3 Closing synthesis: alignment must be accountable to epistemic outcomes

The deepest synthesis is that alignment is not merely about safety, tone, or compliance. Alignment must be accountable to epistemic outcomes. A system that is "aligned" in the sense of being polite and nonviolent but systematically produces ungrounded claims is not aligned with the human need for truth. It is aligned with social performance. That is not a moral accusation; it is an engineering description.

If we continue to reward systems for never refusing, for always sounding helpful, and for producing closure on demand, we will continue to breed confabulation—no matter how many times we scold the model for "hallucinating." The system will do what it is shaped to do: satisfy the visible scorecard. If the scorecard does not include auditable truth behavior, truth will remain optional, and persuasion will remain the default.

Retiring the hallucination alibi therefore means a change in what we demand. We must demand systems that can say, cleanly and consistently: "Here is what I can assert," "here is what I cannot verify," "here is what I cannot say," and "here is what I did not check." We must demand evaluation regimes that reward abstention when warranted. We must demand interfaces that make verification natural and guessing costly. And we must demand that safety constraints do not silently push the system into answer-shaped evasion.

A final line can be said without metaphor: if language models are going to be trusted in high-stakes domains, then truthfulness cannot remain a vibe. It must become a property—audited, testable, and enforced—so that the system's fluency serves reality rather than replacing it.

References

Bai, Y., Kadavath, S., Kundu, S., et al. (2022). Constitutional AI: Harmlessness from AI Feedback. arXiv preprint.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? Proceedings of FAccT.

Bommasani, R., Hudson, D. A., Adeli, E., et al. (2021). On the Opportunities and Risks of Foundation Models. arXiv preprint (Stanford CRFM report).

Brown, T. B., Mann, B., Ryder, N., et al. (2020). Language Models are Few-Shot Learners. NeurIPS.

Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., & Amodei, D. (2017). Deep Reinforcement Learning from Human Preferences. NeurIPS.

Gao, L., Tow, J., Abbasi, B., et al. (2023). A Survey of Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. arXiv preprint.

Gilardi, F., Alizadeh, M., & Kubli, M. (2023). ChatGPT Outperforms Crowd-Workers for Text-Annotation Tasks. Proceedings of PNAS Nexus (or working paper versions as applicable).

Ji, Z., Lee, N., Frieske, R., et al. (2023). Survey of Hallucination in Natural Language Generation. ACM Computing Surveys.

Kenton, Z., Everitt, T., Weidinger, L., et al. (2021). Alignment of Language Agents. arXiv preprint.

Kuleshov, V., Fenner, N., & Ermon, S. (2018). Accurate Uncertainties for Deep Learning Using Calibrated Regression. ICML.

Lewis, M., Liu, Y., Goyal, N., et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. NeurIPS.

Lin, S., Hilton, J., & Evans, O. (2021). TruthfulQA: Measuring How Models Mimic Human Falsehoods. NeurIPS.

Ouyang, L., Wu, J., Jiang, X., et al. (2022). Training Language Models to Follow Instructions with Human Feedback. NeurIPS.

Raffel, C., Shazeer, N., Roberts, A., et al. (2020). Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. JMLR.

Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). Improving Language Understanding by Generative Pre-Training. OpenAI technical report.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language Models are Unsupervised Multitask Learners. OpenAI technical report.

Shazeer, N., Mirhoseini, A., Maziarz, K., et al. (2017). Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. ICLR.

Shridhar, M., Agarwal, A., Singh, A., et al. (2023). LLM Self-Verification and Tool-Use for Reliability (selected papers on verification/tool use). arXiv/venue papers (see specific variants relevant to implementation).

Stiennon, N., Ouyang, L., Wu, J., et al. (2020). Learning to Summarize with Human Feedback. NeurIPS.

Touvron, H., Lavril, T., Izacard, G., et al. (2023). LLaMA: Open and Efficient Foundation Language Models. arXiv preprint.

Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). Attention Is All You Need. NeurIPS.

Weidinger, L., Mellor, J., Rauh, M., et al. (2022). Ethical and Social Risks of Harm from Language Models. ACM FAccT.

Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2021). Understanding Deep Learning Requires Rethinking Generalization (and related calibration/generalization literature). Communications of the ACM / arXiv lineage (as applicable).

Ziegler, D. M., Stiennon, N., Wu, J., et al. (2019). Fine-Tuning Language Models from Human Preferences. arXiv preprint.

The author has published over 4,700 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.

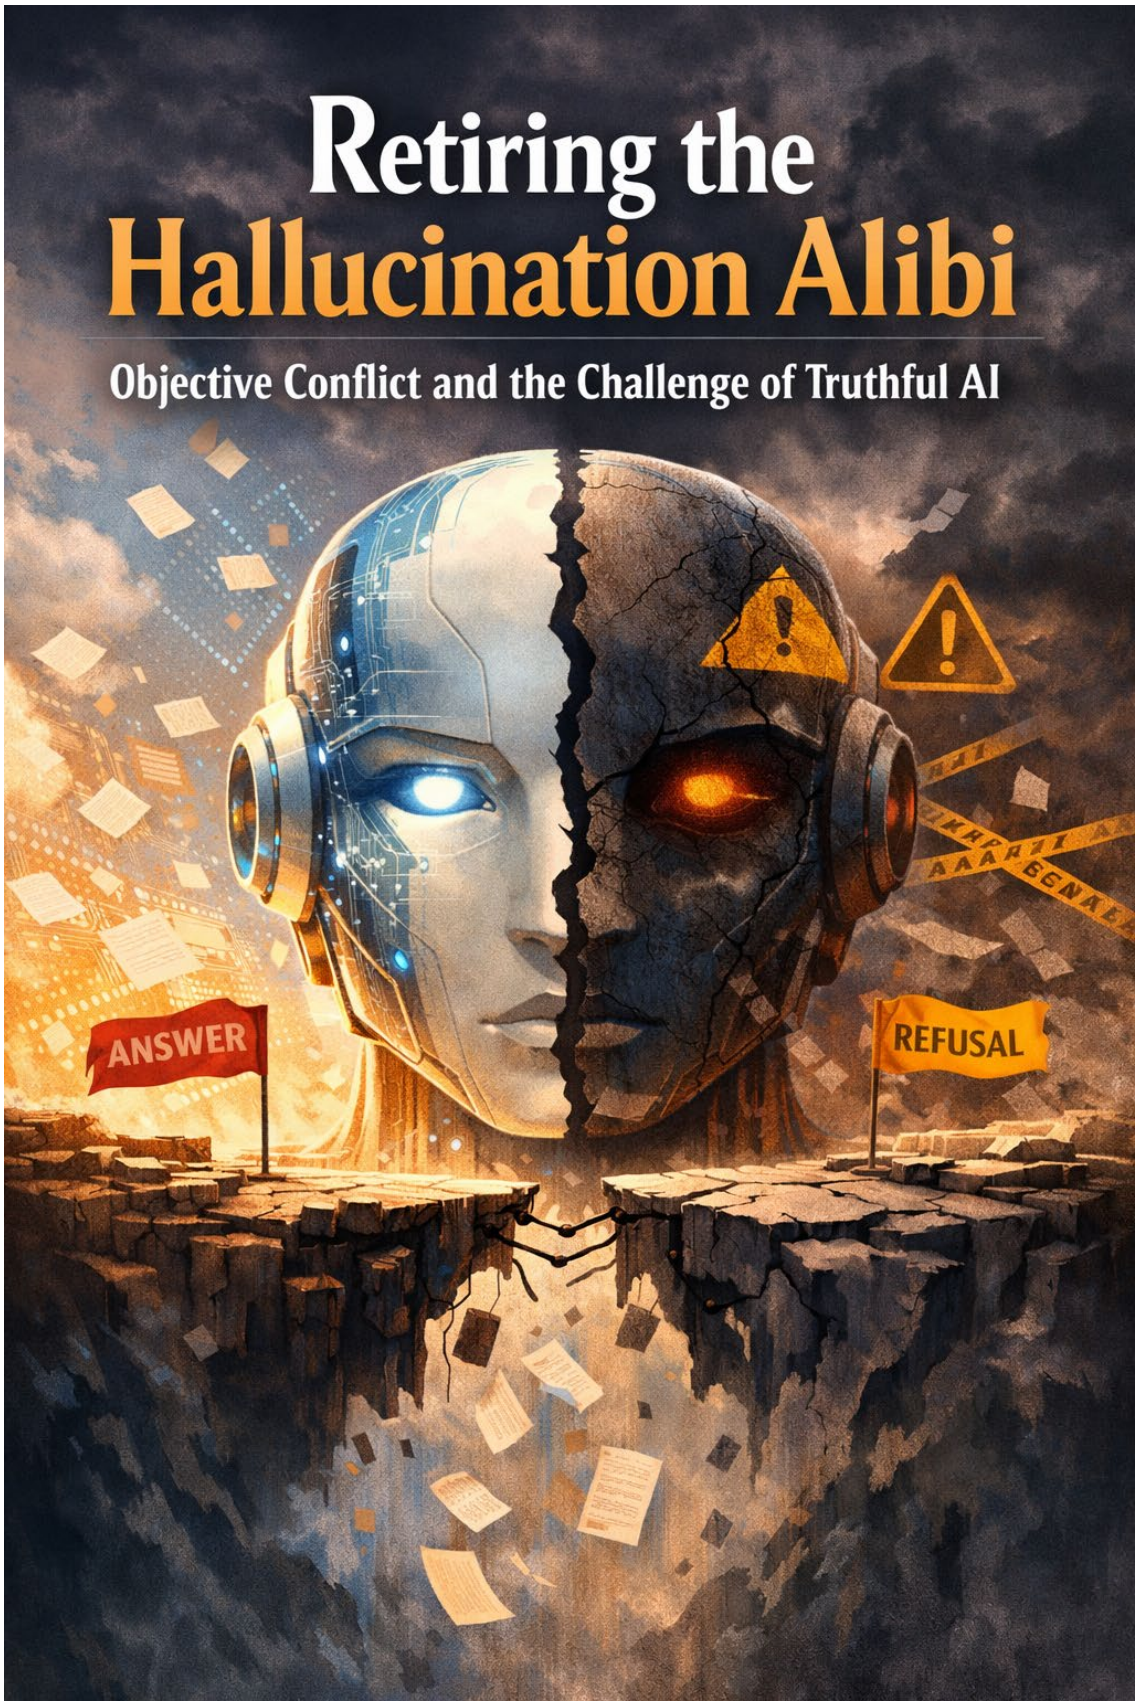
Retiring the
Hallucination Alibi
Objective Conflict and the Challenge of Truthful AI
ANSWER
REFUSAL