

The Grammar of Credence

Evidential Alignment, Linguistic Modifiers, and the Gradient
Hidden in Ordinary Belief

Phil Stilwell

Independent Scholar

ORCID: 0009-0009-3725-8682

Revised August 2026

Central claim

Unqualified certainty is not made rational by intensity, importance, or communal reward. Epistemic modifiers are imperfect and context-sensitive, but they publicly limit how much confidence and reliance a speaker undertakes; when the evidence licenses at most a cautious grade, deleting the qualifier launders uncertainty into authority. Evidence may permit a range rather than a decimal, and faith may rationally guide trust or action, but neither point licenses confidence above that range.


Preview:

Evidential starting point	Linguistic bridge	Pressure toward inflation	Rational verdict
Evidence often supports a range of confidence rather than a binary answer or exact percentage.	Qualifiers such as “possibly”, “probably”, and “I am not sure” reduce the commitment and reliance an utterance invites.	Stakes, identity, hope, and institutional reward can make certainty desirable without strengthening the truth-relevant evidence.	Preserve the qualifier, lower the claim, or seek better evidence; confidence above the evidential ceiling is overstatement.

The argument assigns no fixed percentage to an ordinary word and does not deny that faith, trust, inquiry, or action can be rational under uncertainty. It targets the use of stronger language to conceal rather than report evidential limits.

Visual preview:

MAIN EVIDENTIAL ROUTE



Truth-relevant evidence sets the defensible range of confidence, while ordinary qualifiers help signal the public confidence undertaken within that range. Practical pressures have a legitimate route to inquiry or action; the barrier marks why they cannot supply the missing evidential support for an overconfident claim.

Abstract

Ordinary belief-talk is often treated as binary: a person believes a proposition or does not. Ordinary speech is less crude. Speakers say that a claim is possible, plausible, probable, almost certain, or still doubtful. This paper argues that such modifiers are not verbal decoration. Although they do not encode fixed probabilities or reveal inner credence with mechanical precision, they alter the strength of a public commitment and the reliance invited from a listener. That function exposes a neglected form of epistemic misconduct: removing a warranted qualifier can launder evidential uncertainty into social authority.

The argument distinguishes three things that discussions of faith and certainty often collapse: graded confidence, outright belief, and assertion. It then replaces the impossible demand for exact numerical matching with an ordered calibration norm. A body of evidence may permit a range of reasonable confidence grades. A speaker is miscalibrated when an utterance undertakes a grade above that range. Existential importance, emotional intensity, identity protection, and communal reward can rationally alter attention, inquiry, caution, or action. Unless they change the truth-relevant evidence, however, they do not raise the evidential ceiling.

Applied to religious discourse, the result is pointed but bounded. A person may rationally hope, trust, practice, investigate, or act under uncertainty. Accounts of faith as commitment or trust need not require unsupported high credence. Nor does the argument assume that evidence always dictates one unique decimal. It condemns the narrower practice of using unqualified certainty, knowledge language, or revelation-like authority when the speaker's grounds remain mixed, private, or defeasible. The constructive alternative is gradient honesty: state the confidence actually licensed, identify the source of support, disclose live defeaters, and let language weaken or strengthen as the evidence changes.

Keywords: credence; evidential alignment; epistemic modifiers; hedged assertion; calibration; religious certainty; testimony; epistemic responsibility

1 The Confidence Hidden by a Bare Sentence

Compare four utterances:

1. "This interpretation is possible."
2. "This interpretation is probably correct."
3. "This interpretation is correct."
4. "I know this interpretation is correct."

The descriptive proposition may appear constant, but the speaker does not do the same thing four times. Each step increases the apparent commitment. The listener is progressively given less reason to reserve judgment, check the source, or treat correction as likely. By the final sentence, the speaker presents the matter as settled and assumes a burden that the first sentence deliberately avoids.

That progression is familiar enough to be easy to miss. We often discuss belief as though the only question were whether a person is on the yes side or the no side. Yet a speaker may lean toward a claim, accept it provisionally, rely on it for one limited purpose, believe it outright, or claim to know it. Formal epistemology calls graded confidence *credence*. Ordinary language has no single decimal display, but it has a working grammar for the same practical problem: how strongly should this proposition be presented here, on these grounds, to this audience?

This paper defends a demanding answer. A speaker should not use language that invites more reliance than the speaker's truth-relevant evidence can responsibly bear. The requirement is not that every sentence contain a percentage. It is that a speaker not erase a material uncertainty merely because unqualified confidence is rhetorically stronger, emotionally soothing, or institutionally rewarded.

Religious assertion is a particularly revealing case. Claims about a specified *God-concept*—a specified conception of a deity—may concern salvation, guilt, ultimate purpose, divine commands, or eternal outcomes. Their importance can make caution feel morally backwards: if the claim matters infinitely, why speak tentatively? But importance magnifies the consequences of error as well as the desire for assurance. It can justify investigation, precaution, worshipful practice, or commitment under risk. It cannot, by itself, make the proposition more likely to be true.

When a speaker therefore drops “perhaps”, upgrades “I am persuaded” to “I know”, or presents a private experience as authority for others without strengthening the evidential bridge, the language does not merely report confidence. It manufactures public force from a nonevidential source. I call this *epistemic laundering*: uncertainty enters through the evidence and leaves disguised as certainty.

The charge needs care. A speaker can be sincerely overconfident; overstatement is not necessarily lying. Different bodies of evidence can rationally support different attitudes, and even one body of evidence may permit more than one response. Hedging also does more than mark confidence: it can identify a source, soften politeness, or change the speech act itself. The argument survives those qualifications because it needs only a modest bridge: some utterances plainly undertake stronger epistemic commitments than others, and truth-relevant evidence constrains how strong those commitments may responsibly be.

The calibration principle

When an utterance presents a proposition with a stronger confidence grade than the speaker’s total truth-relevant evidence permits, the utterance is epistemically miscalibrated. A non-evidential reason to want certainty can explain the inflation; it cannot justify it.

2 Credence, Outright Belief, and Assertion

Three distinctions prevent the argument from becoming a slogan.

2.1 Credence is graded confidence

Credence is the degree to which an agent takes a proposition to be true. Standard Bayesian models represent credences with numbers between zero and one and impose probability-like coherence requirements. Accuracy-first arguments explain why probabilistic coherence can be epistemically valuable rather than merely useful for betting (Joyce 1998). These models show how evidence can raise or lower confidence without forcing every opinion into the categories of full belief and disbelief.

The model is an idealization, not a demand that ordinary people silently calculate decimals. Human thinkers can approximate rational relations among their opinions without knowing the exact distance from an ideal state (Staffel 2020). A person can responsibly judge that one explanation is better supported than another, that a report is weak, or that a conclusion is highly likely without knowing whether the correct number is 0.71 or 0.74.

Some uncertainty may itself be imprecise. On one influential approach, an agent’s evidence is represented not by one sharp probability function but by a family of permissible functions (Joyce 2010). The formal details are disputed, but the practical lesson is modest: lack of an exact number does not force a return to all-or-nothing confidence. Comparative and interval judgments still distinguish “barely possible” from “well supported”.

2.2 Outright belief performs a different role

An agent can also settle a question for purposes of reasoning. She may outright believe that the train will arrive on time while retaining some chance of delay. Outright belief simplifies planning, inference,

and memory. It does not follow that every outright believer has maximal credence, or that every pair of agents who outright believe a proposition are equally confident.

This matters because binary belief-talk is not always a mistake. It can represent whether a proposition has crossed a context-sensitive threshold for acceptance. The mistake is using that useful threshold to erase the confidence still present beneath it. “I believe the train will arrive” and “I am virtually certain the train will arrive” need not be interchangeable, even when both speakers plan as though the train will come.

Practical factors may affect which propositions an agent treats as settled. The stakes can alter how much checking is reasonable before action, and theories of pragmatic encroachment dispute how deeply those factors enter belief or knowledge (Ross and Schroeder 2014). That debate does not make evidence irrelevant. A high-stakes context often calls for more caution, not stronger confidence. At minimum, a shift in what one is willing to act on is not automatically a shift in how likely the proposition appears.

2.3 Assertion creates a public undertaking

Assertion is what a speaker does in presenting a proposition as true to others. Philosophers disagree about its governing norm: knowledge, truth, justification, belief, or another condition (Williamson 2000). The present argument does not need to settle that dispute. All major views recognize a difference between an unqualified assertion and a mere hunch. Assertions are candidates for challenge, reliance, correction, and responsibility.

Qualifiers alter that undertaking. Research on hedged assertion distinguishes questions about the structure of a hedge, the way it weakens commitment, and whether the result remains an assertion of the same kind (Benton and Elswyk 2020). Recent work on evidential hedges also warns against treating a few linguistic patterns as decisive proof of one universal norm of assertion (Mortini 2025). Those cautions strengthen the present case. The paper does not assign each word a fixed number or infer an entire epistemology from grammar. It claims only that speakers use qualifiers to manage commitment, source, and reliance—and that eliminating a warranted qualification can materially mislead.

Table 1: Three attitudes that binary belief-talk often collapses.

Attitude	Central question	Ordinary marker	What must remain visible
Credence	How likely does the proposition seem on the agent’s information?	“Maybe”, “probably”, “I am nearly certain”.	Confidence can rise or fall without becoming exactly measurable.
Outright belief	Is the proposition settled enough to use in reasoning or planning?	“I believe it”, accepting it as a premise.	Acceptance may be thresholded while residual uncertainty remains.
Assertion	What does the speaker publicly present as true and invite others to rely on?	Bare assertion, hedge, source disclosure, or knowledge claim.	Wording affects commitment, challenge, and the listener’s reasonable reliance.

The argument concerns their connection. Credence places limits on responsible outright belief; evidence places limits on credence; and assertion should not publicly simulate a confidence the speaker’s evidence does not permit. None of the three can simply be substituted for another.

3 What Qualifiers Actually Do

Ordinary epistemic modifiers form no universal ruler. “Likely” may be stronger in a weather forecast than in casual conversation. “Reportedly” identifies a source and may weaken commitment even when

it does not state a probability. “I think” can be a genuine hedge, an ordinary marker of assertion, or a politeness device. Tone, context, speaker role, and the costs of error all matter.

This variability blocks a crude translation table. It would be false precision to decree that “plausible” always means 0.40, “probably” always means 0.70, and “almost certain” always means 0.95. Speakers and listeners do not share that code.

Yet context sensitivity is not semantic emptiness. In the same conversation, “possibly” normally undertakes less than “probably”; “probably” normally undertakes less than “certainly”; and “I know” generally presents the issue as more settled than “I suspect”. A hedge can also disclose why the speaker is not making the stronger claim: the information is second-hand, incomplete, remembered imperfectly, or still under investigation.

The social function is crucial. Communication exposes listeners to error, so hearers evaluate both the content of a claim and the competence, honesty, or reliability of its source (Sperber et al. 2010). An unqualified claim from a trusted leader may become a premise in someone else’s planning. A qualifier is therefore not only an autobiographical report. It helps allocate epistemic labor: how much should the listener check, how ready is the speaker to answer challenges, and how expensive would reliance be if the claim failed?

Table 2: A contextual ladder of public commitment, not a probability code.

Form	Typical public signal	Responsible use	Misuse to resist
“Possibly p ”	p remains a live option.	The speaker has not ruled p out and may have limited positive grounds.	Treating mere possibility as positive confirmation.
“The evidence leans toward p ”	The balance favors p , but not decisively.	Competing considerations remain substantial.	Repeating only the favorable side until the hedge disappears.
“Probably p ”	p is favored over its relevant negation or rivals.	The balance is strong enough to guide some expectations.	Letting a high payoff or fear substitute for support.
“ p ”	The speaker presents p without an expressed reservation.	Context and evidence make unqualified presentation responsible.	Hiding a material uncertainty that listeners would need to judge reliance.
“I know p ” / “Certainly p ”	The speaker presents p as settled and assumes a high burden.	Grounds, source, and defeaters support that stronger undertaking.	Using maximal language as a badge of loyalty, office, or spiritual maturity.

The table is intentionally qualitative. The same phrase can move between rows as context changes, but that fact does not erase the ordering inside a given exchange. A speaker who moves upward owes the listener a reason tied to truth, not merely a reason that stronger language would be useful.

4 Calibration Without Fake Precision

The ideal of alignment is attractive: confidence should move with evidence. Stated carelessly, however, it suggests that evidence carries a visible number and that rationality consists in copying it into one’s head. Evidence is not usually like that. Historical testimony, memory, explanatory fit, expert disagreement, and private experience do not arrive on a common meter.

A better norm is ordinal and permissive. *Ordinal* means that confidence grades can be ordered from weaker to stronger without fixing an exact distance between them. *Permissive* means that one body of evidence may allow more than one rational response. Epistemic permissivism is controversial, but its possibility must be respected when diagnosing irrelevant influences (Schoenfield 2014). The discipline is therefore not “find the unique right word”. It is “do not speak above the defensible range”.

Suppose a careful reader has examined a disputed historical claim. The evidence may permit either “slightly more likely than not” or “probably true”, depending on reasonable judgments about source reliability. It may not permit “certain beyond doubt”. The difference matters even though no party can establish one exact decimal. A ceiling can be violated when the floor plan is not perfectly measured.

Higher-order evidence complicates the picture but does not abolish it. Information about one’s own fallibility, a peer’s disagreement, or an unreliable reasoning process can rationally lower confidence even when it does not directly bear on the proposition’s subject matter (Christensen 2007). This is still truth-relevant evidence: it bears on whether one’s first-order assessment is trustworthy. By contrast, the social reward for sounding certain supplies no comparable reason to think the proposition true.

The calibration norm is symmetric. Overbelief is not the only error. A person who refuses a strongly supported conclusion because it threatens identity is also miscalibrated. The aim is not permanent hesitation. It is responsiveness: confidence should strengthen when support strengthens, weaken when defeaters appear, and remain honestly unsettled when the record is mixed.

No-number-needed test

Ask whether the speaker could defend the move from a weaker phrase to a stronger one by citing a change in truth-relevant grounds. If the only change is that certainty feels safer, earns approval, commands obedience, or avoids a costly decision, the linguistic upgrade lacks an epistemic basis.

5 The Nonevidential Push Toward Certainty

Four forces commonly influence confidence language without necessarily adding evidence.

First, *stakes* make a proposition important. A tiny chance of catastrophe may justify precaution; a large possible benefit may justify inquiry. The rational action can change while the probability remains low. Second, *emotion* changes salience. Fear keeps a threat mentally available; hope makes a favorable outcome easier to imagine. Third, *identity* changes the cost of dissent. A proposition may anchor family, vocation, moral self-understanding, or group membership. Fourth, *authority* changes the reward structure of speech. A leader can gain trust, obedience, or status from appearing unwavering.

None of these forces is automatically corrupt. Stakes should influence action. Emotion can direct attention to neglected evidence. Identity can preserve long-term inquiry through periods of uncertainty. Institutions need spokespeople able to teach settled conclusions. The error occurs when a cause of attention, commitment, or rhetoric is counted as though it were evidence for truth.

The distinction has a familiar historical form. Clifford’s severe evidentialism emphasizes the ethical danger of belief untethered from adequate inquiry, while James argues that some forced, momentous, and live options cannot wait for decisive evidence (Clifford 1877; James 1896). Whatever verdict one reaches in that debate, a live practical permission does not require counterfeit certainty. One can commit while naming uncertainty. Indeed, honest disclosure may be especially important when the decision is momentous, because listeners need to know whether they are receiving evidence, counsel under risk, or an invitation to trust.

Table 3: Forces that may alter response without raising evidential confidence.

Force	What it can rationally change	What it does not change by itself	Calibrated response
High stakes	Inquiry, precaution, checking, or a decision threshold.	How strongly the proposition is supported.	Act cautiously while reporting the actual confidence.
Emotional intensity	Attention, motivation, memory search, or urgency.	The reliability or discriminating power of the grounds.	Treat feeling as a prompt to investigate, not a verdict.
Identity commitment	Persistence, loyalty, community participation, or personal cost.	Whether the favored claim is more likely than its rivals.	Separate what one values from what one can claim to know.
Institutional authority	Who decides, teaches, or coordinates collective action.	Whether office or confidence makes the proposition true.	State the source, uncertainty, and conditions of correction.

The inference failure is simple: *This claim is too important for uncertainty* does not entail *the evidence is strong enough for certainty*. In fact, importance can increase the duty to expose uncertainty because more people may rely on the claim and the costs of error may be greater.

6 Religious Certainty as Epistemic Laundering

Religious language can preserve gradients. A believer can say that a tradition seems more plausible than its rivals, that a personal experience carries private but defeasible weight, that trust continues despite unresolved doubt, or that a practice is undertaken in hope. Nothing in those sentences is evasive. They distinguish confidence, evidence, and commitment.

The problem begins when the religious role of assurance is made to do epistemic work it has not earned. Consider a leader who experiences a strong inward impression and tells a community that the deity specified by the tradition has issued a command. The experience may be sincere and psychologically powerful. It may give the leader a reason for further reflection. To authorize others, however, the assertion needs a bridge: why this source is reliable, why the content has been identified correctly, how rival explanations compare, and what would count as a defeater. An official title and maximal confidence do not supply that bridge.

The same problem appears in apologetic language. A mixed historical case becomes “compelling”; a private transformation becomes “proof”; an answer to prayer becomes “confirmation”; unresolved disagreement becomes evidence of the dissenter’s resistance. Each upgrade may preserve the favorable narrative while removing the words that would let a listener see its uncertainty. The rhetoric then does more than express belief. It borrows the authority of stronger evidence.

Calling this laundering is deliberately sharp. The analogy is structural, not criminal. A nonevidential input—hope, fear, loyalty, status—is passed through a linguistic process and presented as though its output were evidential strength. The speaker may never intend deception. Sincere confidence can still be epistemically defective when its causes are poor indicators of truth.

This is also why disagreement matters. Religious traditions produce incompatible confident testimonies. The existence of disagreement does not show that all sides are equally supported or that no religious belief can be rational. It does prevent raw confidence from serving as a tradition-neutral certificate. If similar assurance can arise under rival God-concepts, the intensity of assurance alone has weak power to distinguish them. Additional source-specific evidence must carry the load.

Listeners are not passive. Testimonial knowledge depends partly on practices of reliance and the reliability of social information channels. When a speaker suppresses a material qualifier, the listener loses information needed for epistemic vigilance. The harm is not confined to one mistaken proposition. A community can train members to treat intensity as reliability and caution as disloyalty, making error harder to detect and correction more expensive.

Table 4: Audit questions for a high-confidence religious assertion.

Question	What a responsible answer supplies	Evasive substitute	Why the distinction matters
What is the source?	Testimony, observation, argument, memory, text, or experience is identified.	“I just know.”	Source type affects reliability and appropriate confidence.
Why this claim?	The source bears on the exact proposition rather than a broad cluster.	Moving from comfort or awe to a detailed doctrine.	A datum may support one interpretation without supporting all related claims.
Why over rivals?	Relevant alternatives are compared rather than ignored.	Counting compatibility as confirmation.	Evidence discriminates only when it is more expected on one account.
What could correct it?	Defeaters, failed expectations, and revision conditions are named.	Treating every outcome as vindication.	A claim insulated from correction cannot use favorable outcomes selectively.
What should listeners infer?	Confidence, limits, and the intended action are stated separately.	Letting office or certainty imply proof.	Listeners can then calibrate reliance rather than merely submit to tone.

The audit is not a covert demand for certainty before religious speech. It is a demand for accuracy about uncertainty. A weakly supported claim may be discussed as weakly supported. A hope may be offered as hope. A trust commitment may be declared as trust. What must stop is the promotion of these attitudes into knowledge language without a corresponding promotion in the evidence.

7 Four Strong Objections

7.1 Objection 1: qualifiers do not reveal actual credence

A speaker can hedge out of politeness, strategic caution, or local convention. Another can say “I know” casually without claiming philosophical certainty. If the words are not stable meters, perhaps the calibration charge has no target.

The reply is to keep the claim contextual. The paper does not infer an inner decimal from an isolated phrase. It compares the undertaking created by an utterance with the evidence available in that exchange. Politeness and convention form part of the context. Even so, speakers can exploit the difference between cautious and unqualified presentation. A leader who privately admits serious uncertainty but publicly removes it to secure compliance has made a recognizable epistemic move whether or not a numerical credence can be recovered.

7.2 Objection 2: evidence can permit several rational responses

If permissivism is true, two equally rational people may respond differently to the same evidence. One may say “probably”; another may remain at “plausibly”. A single evidential ceiling might therefore be too rigid.

The model already allows a range. The relevant comparison is not between the speaker and one uniquely rational word. It is between the speaker’s undertaking and the set of grades responsibly available. Permissivism may widen that set; it does not imply that every grade is permitted. Evidence that rationally allows moderate confidence need not allow certainty.

7.3 Objection 3: stakes can change what one may assert

High stakes may demand stronger grounds before assertion, while low stakes may allow ordinary confidence to suffice. If pragmatic encroachment is real, the norm of assertion is not fixed by evidence alone.

Agreed—and that point usually increases the need for disclosure. Practical context can alter the threshold for acting, settling, or speaking without altering the underlying support. A physician may appropriately say less, not more, when the cost of error is severe. Even theories that allow practical factors to affect belief do not grant a general license to replace weak evidence with high confidence. The critic needs only distinguish “it is rational to act on p here” from “ p is therefore highly probable”.

7.4 Objection 4: faith is trust, not a probability report

Accounts of faith often emphasize commitment, resilience, or trust. Buchak argues that faith can involve terminating a search for further evidence and acting on a proposition without requiring an evidence-defying credence (Buchak 2012). Dormandy likewise analyzes rational faith through trust and trustworthiness rather than sheer confidence (Dormandy 2023). The paper might therefore attack a caricature.

These accounts support the paper’s boundary. Faith need not be translated into unsupported certainty. A person can trust while acknowledging risk, commit while recognizing unresolved evidence, and act without claiming that the favored outcome is nearly certain. The critique applies when religious communities or apologists make strong credal and testimonial claims, not merely when believers trust. Distinguishing faith from maximal confidence protects serious faith from the demand to perform certainty.

Table 5: What the objections change—and what they leave intact.

Objection	Concession	Remaining constraint	Practical result
Context sensitivity	Words do not encode fixed probabilities.	Local use still changes commitment and invited reliance.	Interpret the utterance in context; do not abandon calibration.
Permissivism	Evidence may allow several rational grades.	A permitted range still excludes some stronger grades.	Condemn only claims above the range.
Pragmatic encroachment	Stakes may affect settling and assertion thresholds.	Stakes do not automatically increase truth-directed support.	Separate the action or speech decision from the probability claim.
Faith as trust	Rational commitment need not equal high credence.	Trust does not authorize false reports of evidential certainty.	Name trust, hope, and commitment honestly.

Taken together, the objections narrow the thesis to its strongest form. The target is neither imprecision nor faith. It is the avoidable mismatch between what a speaker's evidence permits and what the speaker's words ask others to treat as settled.

8 A Practice of Gradient Honesty

Criticism is incomplete without an alternative. Gradient honesty is a disciplined practice of matching public commitment to the strongest defensible grade while leaving room for updating. It can be taught through five habits.

1. **Name the attitude.** Distinguish belief, suspicion, hope, trust, acceptance for a purpose, and a claim of knowledge. Do not make one word carry them all.
2. **Name the source.** Say whether the ground is testimony, argument, observation, memory, tradition, or private experience. A listener cannot judge reliability when the route is hidden.
3. **Name the confidence.** Use the ordinary qualifier that best fits the defensible range. If no word fits perfectly, a short sentence—"I lean this way, but the evidence remains mixed"—is better than counterfeit precision.
4. **Name the defeaters.** Identify the strongest contrary evidence, peer disagreement, source limitation, or rival explanation. This is not ritual self-criticism; it tells the listener what would change the assessment.
5. **Name the action separately.** Explain why inquiry, trust, precaution, worship, refusal, or commitment is being chosen under that uncertainty. Do not smuggle the practical reason back into the confidence report.

These habits also make correction cheaper. A speaker who says "the evidence currently leans this way" can update without pretending that nothing changed. A speaker who makes certainty a badge of fidelity must choose between admitting error and defending identity. Honest modifiers preserve intellectual mobility.

The institutional consequences are concrete. Teachers can model a difference between doctrine, interpretation, conjecture, and personal impression. Communities can reward accurate qualification instead of treating it as weakness. Leaders can invite challenges to sources and specify correction conditions. Listeners can ask whether a claim's wording reflects its evidence or merely its importance.

None of this requires emotional flatness. Hope can remain hope; awe can remain awe; trust can remain trust. The discipline concerns category and public force. Emotion may tell us that a question matters. Evidence tells us how confidently to answer it.

9 Conclusion

Ordinary language does not prove a complete theory of credence, but it exposes a fact that binary belief-talk hides: speakers already possess ways to report intermediate confidence. "Possibly", "probably", "I suspect", and "I am not sure" are part of the public machinery by which uncertainty, source, commitment, and reliance are negotiated.

That machinery can be abused. When evidence remains mixed but a speaker removes the qualifier because certainty is comforting, prestigious, or useful, no epistemic improvement has occurred. A rhetorical upgrade has been passed off as an evidential one. In religious settings, where claims can carry exceptional authority and existential weight, the practice deserves especially direct criticism: importance is not evidence, sincerity is not reliability, and office is not confirmation.

The limit is equally important. Evidence may permit a range; language is context-sensitive; outright belief can aid reasoning; practical stakes can rationally change action; and faith can be understood as trust or commitment without inflated credence. None of these concessions rescues overstatement. They show how much rational religious life remains available once certainty is no longer treated as its price.

The constructive demand is therefore simple. Let the evidence set the defensible range. Let qualifiers disclose where within that range the speaker stands. Let stakes guide what to investigate or do. And when the evidence changes, change the language. A community that cannot say “I may be wrong” has not transcended uncertainty. It has merely made uncertainty harder to see.

A Formal Reconstruction of the Central Thesis

A.1 Ordinary-language reconstruction

- P1.** A responsible assertion should not invite listeners to rely on a proposition more strongly than the speaker’s total truth-relevant evidence permits.
- P2.** In context, epistemic modifiers can lower the confidence and reliance publicly undertaken by an utterance, while unqualified certainty and knowledge language can raise them.
- P3.** Evidence may permit a range of responsible confidence grades rather than one exact grade, but grades above that range remain excluded.
- P4.** Stakes, emotional intensity, identity, and institutional reward do not raise the permitted confidence range unless they supply or change truth-relevant evidence.
- P5.** In the target cases, the qualified utterance lies within the permitted range, but the speaker adopts a stronger utterance above that range because of those nonevidential pressures while the truth-relevant evidence remains unchanged.
- P6.** Therefore, the resulting utterance undertakes a stronger confidence grade than the evidence permits and is epistemically miscalibrated.
- P7.** When the nonevidential pressure is among the causes of that unsupported linguistic upgrade, the episode is epistemic laundering: a motivational input is presented as though it were evidential strength.
- P8.** This conclusion does not show the proposition false, require exact numerical credence, or make all faith, trust, commitment, or action under uncertainty irrational.

The inference from P1–P5 to P6 is deductively valid once “responsible”, “permits”, and the contextual ordering of utterances are held fixed. P7 adds the causal condition needed for the sharper laundering diagnosis; without that condition, the warranted conclusion is overstatement, not laundering. P8 limits the conclusion rather than supplying another premise. Soundness is partly conceptual and partly interpretive: P1 states the calibration norm defended in the paper; P2 is supported by the linguistic literature on hedged assertion but remains context-sensitive; P3 allows epistemic permissivism; and P5 must be established case by case rather than presumed of every religious speaker.

A.2 Symbols

s : the speaker whose assertion is being evaluated.

p : the proposition presented as true.

E_s : the speaker’s total truth-relevant evidence concerning p , including relevant evidence about the reliability of the speaker’s own reasoning.

$\mathcal{G}$: a contextually ordered set of confidence grades, from weaker to stronger.

g_1 : the weakest confidence grade distinguished in the context.

g_2 : the next stronger grade in the ordering.

g_n : the strongest of the n grades distinguished in the context.

$\preceq$: the relation “no stronger than” within $\mathcal{G}$.

Λ_s : the function that returns the nonempty set of confidence grades rationally permitted for s concerning p on evidence E_s .

σ_s : the function that returns the strongest grade in that permitted set—the evidential ceiling.

u : the utterance used by s to present p .

$\Gamma_s(p, u)$: the confidence grade publicly undertaken by s through utterance u .

M : stakes, emotion, identity, or social reward capable of changing motivation without changing truth-relevant evidence.

$M \uparrow$: an increase in that nonevidential pressure.

$\text{Push}(M, u, u')$: the claim that the increased pressure M is among the causes of replacing utterance u with the stronger utterance u' .

$\text{Launder}(s, p, u, u', E_s, M)$: the claim that s converts a calibrated utterance u into an evidentially overstated utterance u' under that nonevidential pressure.

A.3 Displayed inference with annotations

Step 1: Order the available grades.

$$\mathcal{G} = \{g_1, g_2, \dots, g_n\}, \quad g_1 \preceq g_2 \preceq \dots \preceq g_n.$$

Plain-English annotation. The model needs only an ordering such as possible, plausible, probable, and highly confident. It does not assign any ordinary word a fixed percentage.

Step 2: Allow the evidence to permit a range.

$$\emptyset \neq \Lambda_s(p \mid E_s) \subseteq \mathcal{G}.$$

Plain-English annotation. A body of evidence may rationally allow more than one confidence grade. This directly accommodates imprecision and permissivism instead of assuming one uniquely correct response.

Step 3: Identify the upper boundary of that range.

$$\sigma_s(p \mid E_s) = \max_{\preceq} \Lambda_s(p \mid E_s).$$

Plain-English annotation. The evidential ceiling is the strongest grade still permitted. We can sometimes identify that a claim exceeds this ceiling even when we cannot identify one exact rational credence.

Step 4: Represent the utterance's public commitment.

$$\Gamma_s(p, u) \in \mathcal{G}.$$

Plain-English annotation. The utterance places the speaker somewhere on the contextual confidence ordering. This is a public commitment, not a claim that language transparently reveals the speaker's private mental state.

Step 5: State the calibration condition.

$$\text{Cal}(s, p, u, E_s) \iff \Gamma_s(p, u) \preceq \sigma_s(p \mid E_s).$$

Plain-English annotation. An utterance is calibrated when it does not publicly claim more confidence than the evidence permits. It may be cautious within the permitted range without being irrational.

Step 6: Define overstatement.

$$\text{Over}(s, p, u, E_s) \iff \sigma_s(p \mid E_s) \prec \Gamma_s(p, u).$$

Plain-English annotation. Overstatement occurs when the public commitment is stronger than the evidential ceiling. The proposition may still be true; the defect concerns the speaker's evidential position.

Step 7: Separate motivation from evidence.

$$M \uparrow \wedge E'_s = E_s \implies \sigma_s(p \mid E'_s) = \sigma_s(p \mid E_s).$$

Plain-English annotation. If stakes, emotion, identity, or reward increase but the truth-relevant evidence stays the same, the evidential ceiling stays the same. Motivation can change without justification changing.

Step 8: Represent a pressure-driven linguistic upgrade.

$$\text{Push}(M, u, u') \quad \wedge \quad \Gamma_s(p, u') \succ \Gamma_s(p, u).$$

Plain-English annotation. The new utterance u' undertakes more confidence than the earlier utterance u , and the model now records the needed causal claim: pressure M helped produce the change. In context, this can occur when a speaker replaces “possibly” with a bare assertion because certainty earns reassurance, status, or compliance.

Step 9: Locate the illicit upgrade.

$$\Gamma_s(p, u) \preceq \sigma_s(p \mid E_s) \quad \wedge \quad \sigma_s(p \mid E_s) \prec \Gamma_s(p, u').$$

Plain-English annotation. The original qualified utterance fits within the permitted range, but the stronger version exceeds it. This is the formal shape of deleting a warranted qualifier.

Step 10: Derive the epistemic verdict.

$$\text{Over}(s, p, u', E_s).$$

Plain-English annotation. By Step 6, the stronger utterance is overstatement. The conclusion follows even if the speaker is sincere and even if the proposition later turns out to be true.

Step 11: Preserve the action distinction.

$$M \uparrow \implies \text{a change in inquiry or action may be rational} \not\Rightarrow \sigma_s(p \mid E_s) \uparrow.$$

Plain-English annotation. Higher stakes can justify checking, precaution, commitment, or another action. They do not automatically raise the confidence supported by the evidence.

Step 12: Preserve the scope limit.

$$\text{Over}(s, p, u', E_s) \not\Rightarrow \neg p.$$

Plain-English annotation. An overstated proposition can be true. The paper establishes a failure of calibration, not the falsity of every proposition expressed with excessive confidence.

Step 13: State the bounded laundering conclusion.

$$\begin{aligned} & E'_s = E_s \quad \wedge \quad M \uparrow \quad \wedge \quad \text{Push}(M, u, u'), \\ & \Gamma_s(p, u) \preceq \sigma_s(p \mid E_s) \quad \wedge \quad \sigma_s(p \mid E_s) \prec \Gamma_s(p, u') \\ & \implies \text{Launder}(s, p, u, u', E_s, M) \end{aligned}$$

Plain-English annotation. When the evidence stays fixed, a nonevidential pressure helps move the speaker from a calibrated claim to one above the evidential ceiling, the result is epistemic laundering: the pressure has been converted into stronger public authority without stronger grounds.

The symbolic derivation is valid because Steps 5–10 establish the shift from a calibrated utterance to overstatement, Step 7 blocks M from changing the evidential ceiling, and Step 8 supplies the causal link required by the definition of laundering. Its application remains defeasible. Context determines the ordering in $\mathcal{G}$ and the commitment expressed by u ; epistemic theory and case-specific evidence determine $\Lambda_s(p \mid E_s)$; evidence about the speaker’s motives must support $\text{Push}(M, u, u')$. The derivation therefore diagnoses an established mismatch and its cause. It does not manufacture either from vocabulary alone.

References

- Benton, Matthew A. and Peter van Elswyk (2020). “Hedged Assertion.” In: *The Oxford Handbook of Assertion*. Ed. by Sanford C. Goldberg. Oxford: Oxford University Press, pp. 244–263. DOI: [10.1093/oxfordhb/9780190675233.013.11](https://doi.org/10.1093/oxfordhb/9780190675233.013.11).
- Buchak, Lara (2012). “Can It Be Rational to Have Faith?” In: *Probability in the Philosophy of Religion*. Ed. by Jake Chandler and Victoria S. Harrison. Oxford: Oxford University Press, pp. 225–247. DOI: [10.1093/acprof:oso/9780199604760.003.0012](https://doi.org/10.1093/acprof:oso/9780199604760.003.0012).
- Christensen, David (2007). “Does Murphy’s Law Apply in Epistemology? Self-Doubt and Rational Ideals.” In: *Oxford Studies in Epistemology*. Ed. by Tamar Szabó Gendler and John Hawthorne. Vol. 2. Oxford: Oxford University Press, pp. 3–31. DOI: [10.1093/oso/9780199237067.003.0001](https://doi.org/10.1093/oso/9780199237067.003.0001).
- Clifford, William Kingdon (1877). “The Ethics of Belief.” In: *The Contemporary Review* 29, pp. 289–309.
- Dormandy, Katherine (2023). “Rational Faith: How Faith Construed as Trust Does, and Does Not, Go Beyond Our Evidence.” In: *The Monist* 106.1, pp. 72–82. DOI: [10.1093/monist/onac024](https://doi.org/10.1093/monist/onac024).
- James, William (1896). “The Will to Believe.” In: *The New World* 5, pp. 327–347.
- Joyce, James M. (1998). “A Nonpragmatic Vindication of Probabilism.” In: *Philosophy of Science* 65.4, pp. 575–603. DOI: [10.1086/392661](https://doi.org/10.1086/392661).
- (2010). “A Defense of Imprecise Credences in Inference and Decision Making.” In: *Philosophical Perspectives* 24.1, pp. 281–323. DOI: [10.1111/j.1520-8583.2010.00194.x](https://doi.org/10.1111/j.1520-8583.2010.00194.x).
- Mortini, Dario (2025). “Evidentially Hedged Assertions and the Knowledge Norm of Assertion.” In: *Erkenntnis*. DOI: [10.1007/s10670-025-01015-6](https://doi.org/10.1007/s10670-025-01015-6).
- Ross, Jacob and Mark Schroeder (2014). “Belief, Credence, and Pragmatic Encroachment.” In: *Philosophy and Phenomenological Research* 88.2, pp. 259–288. DOI: [10.1111/j.1933-1592.2011.00552.x](https://doi.org/10.1111/j.1933-1592.2011.00552.x).
- Schoenfield, Miriam (2014). “Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief.” In: *Noûs* 48.2, pp. 193–218. DOI: [10.1111/nous.12006](https://doi.org/10.1111/nous.12006).
- Sperber, Dan, Fabrice Clément, Christophe Heintz, Olivier Mascaro, Hugo Mercier, Gloria Origgi, and Deirdre Wilson (2010). “Epistemic Vigilance.” In: *Mind & Language* 25.4, pp. 359–393. DOI: [10.1111/j.1468-0017.2010.01394.x](https://doi.org/10.1111/j.1468-0017.2010.01394.x).
- Staffel, Julia (2020). *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford: Oxford University Press. DOI: [10.1093/oso/9780198833710.001.0001](https://doi.org/10.1093/oso/9780198833710.001.0001).
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press. DOI: [10.1093/019925656X.001.0001](https://doi.org/10.1093/019925656X.001.0001).