

Epistemic Games of Religion

Rule-Bound Rationality, Insulated Victory Conditions,
and the Performance of Inquiry

Phil Stilwell

Independent Scholar

ORCID: 0009-0009-3725-8682

Revised August 2026

Central claim

A religious practice may rationally classify evidence, interpret texts, and award an internal victory without thereby establishing its central claims to outsiders. That victory earns public warrant only if some feasible challenge can be admitted, counted against the claim, and made to lower confidence under procedures the claim's own dependent authorities do not fully control. This independent-loss condition limits a public inference; it does not show that religious belief is false, meaningless, or personally irrational.

Preview: From an internal win to public warrant

Rule-bound practice	In-game success	Independent liability	Bounded result
A community sets admissible sources, moves, interpreters, and standards.	A claim fits those sources, survives approved objections, or sustains faithful participation.	A feasible challenge can enter, count negatively, and require revision without the claim controlling every stage.	Internal success can contribute to public warrant only when this liability is present; winning alone is insufficient.

“External” does not mean view-from-nowhere, secular, or laboratory-only. It means that the challenged position cannot guarantee its own success by controlling admission, scoring, and the consequence of a loss.

Visual preview:

Winning Is Not Yet Being Right



Read the image from left to right. Rules decide which moves are legal, evidence rules decide what may enter, and scorekeeping decides what counts as success; together they can produce an in-game win. The separate external check asks a different question: can a feasible challenge be admitted, scored against the central claim, and allowed to lower confidence under procedures the claim does not control? The image does not assume that rules are bad or that every religious inquiry is closed. It shows why winning under a system's own standards is not yet public evidence that the system accurately represents a shared world.

Abstract

Religious traditions can display serious reasoning. They distinguish trained from untrained interpretation, rank sources, answer objections, preserve doctrinal relationships, and revise secondary claims. The resulting activity may be coherent and rational by the standards of the practice. A public inference nevertheless remains: why should an internal victory count as evidence that the practice's central claims accurately describe a shared world? This paper develops the metaphor of an epistemic game to isolate that question without treating all rule-governed inquiry as defective. It distinguishes internal entitlement, practical value, and public warrant; then it identifies referee capture as a structural dependence in which the challenged claim or authorities whose standing depends on it control admission, negative scoring, or the consequence of apparent loss. The proposed independent-loss condition requires at least one feasible challenge that can be admitted, counted against the central claim, and made to reduce confidence under procedures not fully controlled by those dependent authorities. The condition does not demand neutrality, scientific experiment, or the abandonment of all prior commitments. It demands a route by which resistance can survive the framework's own interpretive machinery. Cross-tradition disagreement makes the requirement especially important because incompatible communities can each award sincere internal victories. A five-part accountability protocol shows how religious inquiry can improve public answerability through claim-specific exposure, independent admission, reciprocal scoring, revision consequences, and cross-community audit. The conclusion is conditional and jurisdictional. Internal authorities may settle what is orthodox or faithful within a tradition; they cannot, solely by ruling in its favor, settle whether its contested account of reality is publicly warranted. An annotated formal appendix states the inference and separates its valid logical form from the contestable premise that public warrant requires an independent loss condition.

Keywords: religious epistemology; social epistemology; public warrant; internal rationality; referee capture; epistemic games; loss conditions; religious disagreement

1 From Rational Performance to Public Warrant

Religious reasoning is often more structured than its critics allow. A tradition can train interpreters, preserve distinctions, reject invalid moves, compare proposed explanations, and maintain a long memory of objections. Its members may disagree vigorously while sharing standards for competent argument. None of this is mere theater. It can involve real intellectual discipline.

Yet discipline supports several different conclusions. It may show that a participant understands a tradition, that one doctrine coheres better than another with accepted sources, or that a practice gives its members a stable way to deliberate. A stronger conclusion says that the victorious claim accurately describes a reality shared by insiders and outsiders. That conclusion needs an additional bridge.

The game metaphor makes the gap visible. Wittgenstein uses *language-game* to emphasize that words acquire sense through patterned use within forms of life (Wittgenstein 1953). Huizinga treats play as culturally serious, and Suits emphasizes the constitutive role of rules and goals in making an activity the game it is (Huizinga 1949; Suits 1978). The metaphor therefore does not imply frivolity or deception. A game can demand intelligence, courage, and integrity. It can also have standards whose authority is limited to the activity they constitute.

Chess illustrates the jurisdictional point. Its rules settle whether a position is checkmate. They do not settle whether chess is the best preparation for governing a city. A court's rules may establish what is legally admissible without settling every historical fact. A church's rules may establish what counts as orthodox teaching without, by that fact alone, establishing divine authorship of the rule-defining sources. The rules have genuine authority, but their jurisdiction must be identified.

An *epistemic game* in this paper is a socially organized practice that allocates four things: admissible sources, permissible inferential moves, recognized adjudicators, and success conditions. The term is

descriptive before it is critical. Science, law, history, journalism, and ordinary conversation all contain such structures. The criticism begins only when an achievement defined within a practice is exported as public truth without showing that the practice remains answerable to resistance it cannot simply reclassify.

The governing question is thus not, “Does this religion reason?” It is, “What kind of success has its reasoning earned, and what could make the central claim lose?” The answer proposed here is an *independent-loss condition*. A public evidential victory requires at least one feasible challenge that can pass three stages—admission, negative scoring, and revision consequence—under procedures not fully controlled by authorities whose standing depends on the challenged claim.

The independent-loss condition

An in-game victory can support public warrant for a central claim only if some feasible challenge can be admitted, counted against that claim, and made to lower confidence under procedures not fully controlled by the claim or by authorities whose relevant standing depends on it.

This is not a falsification rule requiring one decisive experiment. A historical claim may face cumulative documentary defeat. A moral interpretation may lose credibility through counterexamples and incoherent implications. A personal explanation may be weakened by a stable pattern of behavior. The condition is deliberately broad. What matters is not one method but the presence of a genuine route from resistance to epistemic cost.

2 What an Epistemic Game Contains

Every inquiry begins with inherited concepts and standards. Kuhn’s account of paradigms made it difficult to imagine science as a stream of uninterpreted observations confronting isolated propositions (Kuhn 1962). Evidence is selected, described, and related to theory through learned practices. Religious inquiry is not disqualified merely because it also works through a tradition.

Four elements matter for the present analysis:

1. **Admission rules** decide which experiences, texts, testimony, arguments, and objections may enter consideration.
2. **Move rules** decide what participants may infer, which distinctions are legitimate, and which interpretive responses count as competent.
3. **Adjudicators** decide whether an objection succeeds, which sources take priority, and who possesses relevant authority.
4. **Victory conditions** decide when a claim has survived, an interpretation has prevailed, or inquiry has reached an acceptable stopping point.

These elements can produce knowledge. Testimonial networks, expert communities, and institutions are unavoidable in social epistemology. Goldman’s truth-oriented social epistemology evaluates such practices partly by whether they foster knowledge rather than error or ignorance (Goldman 1999). Longino likewise rejects the idea that objectivity requires solitary, value-free observers. On her account, social criticism can improve inquiry when there are recognized avenues for criticism, uptake of criticism, public standards, and sufficiently tempered equality of intellectual authority (Longino 1990; Longino 2002).

Those approaches block an easy contrast between rule-bound religion and rule-free science. The relevant difference is functional. Do the rules organize criticism, or do they guarantee that a protected conclusion survives it? Are adjudicators corrigible, or does dissent from their ruling automatically disqualify the dissenter? Can a challenge alter confidence, or only occasion another permitted interpretation?

The same formal structure can produce different epistemic effects. A canon may preserve hard-won knowledge or exclude inconvenient evidence. Expertise may filter noise or shield a guild. A shared vocabulary may improve precision or make contrary experience unintelligible. Stopping rules may prevent endless doubt or freeze a conclusion too early. The label *game* settles none of these questions; it directs attention to how the parts work together.

Berger and Luckmann’s sociology of knowledge helps explain why these structures feel natural to participants. Communities institutionalize roles, legitimate their shared world, and maintain its plausibility across generations (Berger and Luckmann 1966). This sociological fact neither proves nor disproves the world so maintained. It does show why internal stability and truth should not be treated as synonyms. False, true, and partly true outlooks can all acquire social supports.

3 Three Kinds of Success

The source of much confusion is that a single word—*rational*—is asked to cover several achievements. At least three should be separated.

Table 1: Three kinds of success in a religious practice.

Kind of success	What it can establish	What it does not establish by itself
Internal entitlement	A belief or interpretation follows competently from the practice’s accepted sources, rules, and authorities.	That outsiders should accept those sources or that the central claim is true.
Practical or existential value	A practice sustains identity, hope, moral formation, solidarity, or a meaningful life.	That the practice’s contested descriptions of supernatural reality are accurate.
Public warrant	Shared reasons favor a claim over live alternatives under procedures that can impose evidential loss on every side.	Certainty, complete neutrality, or immunity from future revision.

Internal entitlement is not fake justification. Suppose a scholar accepts a scriptural canon and well-established interpretive rules. One reading may fit grammar, genre, and the wider canon better than another. That is a real comparative achievement relative to the accepted materials. It may rationally guide members who already have reason to trust them.

Plantinga’s account of warranted Christian belief presents a stronger possibility. On the Aquinas/Calvin model, belief in a deity may arise from a properly functioning cognitive faculty, and distinctively Christian belief may receive warrant through the internal work of the Holy Spirit (Plantinga 2000). If the relevant God-concept is true and the proposed cognitive design operates as described, internal religious belief need not depend on an argument from public premises. This paper does not refute that conditional model.

It does, however, identify a limit on what follows in a dispute among people who do not already accept the model. A community’s recognition of a faculty as divinely designed cannot by itself establish to a rival that the faculty is reliable. Otherwise incompatible traditions could each settle the matter by

citing the spiritual source their own practice recognizes. The distinction is between possible personal warrant and a public argument capable of adjudicating the disputed source.

Practical value is also genuine. Ritual, narrative, and community can provide courage, identity, consolation, and disciplined attention. An insulated practice may even produce some reliable moral or psychological insights. Those goods deserve assessment on their own terms. They become epistemically misleading only when their presence is treated as sufficient evidence for every metaphysical or historical claim carried by the practice.

Public warrant is narrower than universal agreement but more demanding than internal permission. It requires reasons and correction procedures that rival participants can examine without first granting the contested conclusion. The standard is not that every person must actually be persuaded. It is that no party may secure victory simply by controlling every rule that determines whether its central claim has lost.

4 Referee Capture: A Dependence Problem

The phrase *referee capture* can sound like an accusation of corruption. Here it names a structure, not a motive. Capture occurs when the authority that controls a decisive stage of evaluation owes its relevant authority to the very claim under review, with no independent route for correcting its ruling.

Consider a dispute over whether a text is divinely authoritative. The text may be internally authoritative because a community has adopted it as canon. But if divine authority is established only by the text's self-description, interpreted by officials authorized only by that text, then the same dependence runs through source, interpreter, and verdict. A critic who rejects the conclusion is not merely outvoted; the critic may be classified as spiritually unable to judge. The result can be stable, coherent, and sincere while remaining circular as public adjudication.

The dependence can appear at three stages:

1. **Admission capture:** apparent counterevidence is excluded because only sources recognized by the target view may speak to the claim.
2. **Scoring capture:** the evidence is admitted, but the target view alone determines whether it counts negatively. Silence becomes hidden presence, failure becomes testing, and disagreement becomes rebellion without an independently constrained comparison.
3. **Consequence capture:** a negative score is acknowledged, but no result is permitted to lower confidence in the central claim. Only auxiliary beliefs, human understanding, or the critic's character may change.

Any one stage can block loss. If contrary testimony is inadmissible, it never reaches the scoreboard. If it is admitted but necessarily receives a preserving interpretation, it cannot count against the claim. If it counts against the claim but confidence is forbidden to fall, the score has no consequence. Genuine exposure requires all three connections.

Table 2: The independent-loss ledger.

Checkpoint	Open loss condition	Captured substitute
Claim specificity	A named claim is precise enough for contrary evidence to bear on it.	The target moves among broad theism, one tradition, one doctrine, metaphor, and mystery as pressure changes.
Admission	Relevant testimony, records, arguments, and experiences can enter even when they threaten the target.	The target’s sources or officials alone decide which challenges are spiritually or conceptually admissible.
Negative scoring	Rules stated independently of the focal outcome allow some results to lower the claim’s standing.	Each result receives a preserving interpretation chosen after it occurs.
Revision consequence	A recognized loss lowers confidence, narrows content, or changes practice in a recorded way.	Peripheral beliefs move while confidence in the central claim remains fixed by rule.
Reciprocal audit	Comparable claims from rival traditions face the same standards and access to criticism.	The home tradition receives favorable exceptions that would be rejected when used by a rival.

Independence here is not absolute causal separation. A Christian historian may remain Christian while evaluating a resurrection claim. A Muslim scholar may remain Muslim while assessing a hadith. A skeptic may possess background commitments that shape interpretation. The question is whether the procedure contains constraints that the favored conclusion cannot alter merely because it is threatened. Independent dating methods, publicly inspectable manuscripts, reciprocal source criticism, explicit confidence ranges, and recorded revision rules can all provide partial independence while participants retain their identities.

Boudry and Braeckman describe self-validating belief systems whose defensive mechanisms turn adverse information into support or explain it away (Boudry and Braeckman 2012). Referee capture explains how that resilience can become institutional rather than merely psychological. The issue is not only that one thinker rationalizes. A community may distribute the work: one authority admits evidence, another supplies approved interpretations, and a third assigns the spiritual meaning of dissent. No individual need intend insulation for the total system to lack an operative loss condition.

5 Cross-Game Comparison and Reciprocal Standards

Religious diversity raises the stakes. Christianity, Islam, Hindu traditions, Latter-day Saint communities, and many smaller movements contain intelligent interpreters, sacred histories, transformative experiences, and procedures for distinguishing faithful from unfaithful belief. Their central claims cannot all be true in their ordinary, incompatible forms. Internal victory therefore cannot be a sufficient public test unless some independent reason shows why one game’s victory conditions track truth while the others’ do not.

This point does not assume that all religions are equally plausible. Their historical records, metaphysical commitments, ethical consequences, and explanatory resources differ. The point is methodological: those differences must be shown through standards that do not automatically award the home tradition every disputed premise.

A reciprocal standard asks, “Would I accept this move if a rival used it?” If inner assurance confirms my canon, what status would comparable assurance have in another tradition? If apparent error in my text invites harmonization, what amount of harmonization would I allow a rival? If growth and moral fruit support my community, how do I score similar fruit elsewhere? If defeaters to my belief are explained by sin or spiritual blindness, may a rival explain my dissent in the same way?

Reciprocity is not mechanical equality. The evidence may justify different conclusions. Nor must every source be granted equal initial credibility. A forged document and a well-attested archive should not receive identical treatment. Reciprocity requires that differences in treatment be supported by differences the favored claim does not create solely for its own protection.

The comparison can be organized as a translation exercise. Insider terms should first be stated in claim-neutral form where possible. “The Spirit confirms this teaching” becomes a claim that a particular experience reliably indicates a divine source. “The tradition has preserved truth” becomes a historical and institutional reliability claim. “The doctrine transforms lives” becomes a comparative claim about psychological, behavioral, and communal outcomes. The translation does not dismiss the original meaning. It gives rival explanations a common object of assessment.

This is the paper’s distinct contribution relative to a general complaint about flexible explanation. Flexibility concerns how many outcomes a framework can accommodate. The independent-loss condition concerns jurisdiction over the evaluation process. A framework might make quite specific predictions yet remain captured if it alone decides whether failure occurred. Conversely, a tradition may permit broad interpretation yet preserve public accountability by clearly separating devotional meaning from evidential claims and allowing external evidence to revise the latter.

6 Why the Standard Is Not Scientism

A predictable objection says that the independent-loss condition forces religion to imitate experimental science. It does not. The condition is method-neutral. Historical inquiry, textual criticism, mathematics, law, and ordinary interpersonal knowledge expose claims to correction in different ways.

A historian cannot rerun an ancient event. Historical claims can nevertheless lose through anachronism, provenance failure, inconsistent records, better rival explanations, or the discovery that a source is dependent on another. A legal judgment may be overturned because evidence was excluded improperly, a precedent was misapplied, or a conflict of interest compromised adjudication. A mathematical argument loses when a counterexample or invalid step survives checking. A friend’s explanation loses credibility when it conflicts with a stable record of conduct. None of these practices requires a laboratory; each allows some resistance to impose cost.

Science itself is not a perfect contrast class. Kuhn shows how paradigms organize normal inquiry and how anomalies may be managed before a larger conceptual change occurs (Kuhn 1962). Scientific institutions can suffer hierarchy, publication bias, career incentives, and group loyalty. Longino’s work matters precisely because objectivity is an achievement of structured critical interaction, not an automatic property of scientists (Longino 1990; Longino 2002). The comparison should therefore be between more and less correction-capable practices, not between immaculate science and irrational religion.

Nor does public answerability exclude testimony or experience. Much knowledge depends on both. The questions are comparative reliability and correction. Who is likely to have the experience? How stable is its interpretation across cultures? What independent checks distinguish perception from expectation, memory reconstruction, social reinforcement, or deception? Can contrary experience count, or is its absence always attributed to the subject? Treating experience seriously requires such questions; exempting it from them does not honor it.

The standard also avoids the demand for a neutral foundation accepted by everyone. Public standards can themselves be criticized and revised. Logic, documentary methods, measurement practices, and norms of testimony are not self-justifying absolutes. They earn authority partly through broad applicability, correction of error, comparison of results, and vulnerability to challenge. The key

asymmetry is not sacred versus secular. It is whether a procedure can be corrected by participants who do not already owe allegiance to the conclusion it protects.

7 A Constructive Accountability Protocol

The argument is useful only if it identifies improvements. A religious community need not abandon its identity to increase public answerability. It can separate the kinds of claim it makes, state loss conditions prospectively, and open decisive stages to reciprocal criticism.

Table 3: A five-part public-accountability protocol.

Step	Question to answer	Public product
Name the claim	Which exact historical, metaphysical, experiential, or moral proposition is being supported?	A target that cannot move silently when challenged.
State admission rules	What information may enter, and who can introduce unwelcome evidence or argument?	A route for relevant resistance to reach deliberation.
State scoring rules	Which feasible observations or arguments would count against the claim, and how much?	A prospective loss condition rather than an after-the-fact verdict.
State revision consequences	What would lower confidence, narrow the claim, change practice, or suspend judgment?	A visible connection between losing and belief revision.
Run reciprocal audit	Would the same rules be accepted for a serious rival, and can critics contest the referees?	Evidence that the home position is not awarding itself special immunity.

The first step prevents target drift. A general claim that some divine reality exists faces different evidence from the claim that one named deity inspired one text, acted in one event, or answers one kind of petition. If evidence for a narrow claim is offered, a retreat to broad possibility cannot preserve the original evidential victory without acknowledging reduced content.

The second step concerns access. A community may screen out irrelevant or malicious challenges, just as every serious inquiry does. But criteria of relevance should be explainable without assuming that all threatening evidence is spiritually defective. Critics need not control the agenda; they need a recognized path to place a relevant challenge before adjudicators.

The third step is the core. A loss condition should be stated before the focal result is known whenever possible. It need not be numerical. A textual claim might specify what manuscript pattern would weaken it. A miracle claim might specify which medical documentation, timing, alternative causes, and base-rate comparison matter. A doctrine of guidance might specify the pattern of clarity or convergence expected. The more flexible the claim, the more important it is to distinguish possible compatibility from positive support.

The fourth step prevents symbolic concessions. A tradition can admit anomalies while protecting the central claim from every cumulative effect. Revision should therefore be linked to consequences: a lower confidence range, a narrower public claim, suspension, changed teaching, or a documented auxiliary whose future performance can be tested. Not every loss requires abandonment, but every genuine loss must cost something.

The fifth step checks jurisdiction. A standard that seems compelling only when it favors the home tradition is not yet public. Rival cases reveal hidden permissions. Reciprocal audit may also strengthen

the home claim: if its evidence remains superior under common standards, the argument becomes more rather than less credible.

This protocol is compatible with graded outcomes. A practice may be highly open in historical scholarship but closed in evaluating private revelation. It may permit external evidence to revise peripheral doctrines while protecting one central commitment. Public warrant should track that profile rather than receive an all-or-nothing label.

8 Strong Objections and Replies

8.1 Every framework has ultimate commitments

One reply says that all reasoning eventually reaches commitments it cannot prove without circularity. The independent-loss condition therefore appears selectively skeptical.

The reply confuses proof from nowhere with correction from somewhere. A standard can be fallible and historically situated while still exposing its applications to challenge. Measurement rules are revised when they generate unstable results. Historical methods are criticized through cases of forgery, dependence, and prediction. Logical systems can be compared for consistency, expressive power, and applicability. No absolute foundation is required for a real difference between a rule that can be revised under resistance and a rule that defines resistance as disobedience.

8.2 If the religion is true, its authorities may be reliable

This objection is correct as a conditional. If a perfectly informed divine agent authorized a text or faculty, following it may be reliable even when outsiders cannot verify the authorization. Plantinga's model makes room for this possibility (Plantinga 2000). The present argument does not derive falsity from public non-establishment.

The issue is what can be offered in the shared dispute. "This authority is reliable if my central theology is true" cannot independently establish that theology. It may explain how insiders possess warrant; it cannot make the authority a neutral referee between its own theology and a rival. Additional public evidence is needed to cross that gap.

8.3 Religious experience is not controlled by authorities

An unexpected experience may indeed resist official categories. Converts often describe events that disrupt prior commitments. Such cases show why the analysis must be applied, not assumed.

The later treatment of the experience remains crucial. Are disconfirming experiences admitted with the same seriousness as confirming ones? Can loss of a sensed divine presence lower confidence, or only confirm a test, hiddenness, or personal failure? Are comparable experiences in rival traditions credited symmetrically? An experience can be independent at its occurrence while its public score is captured during interpretation.

8.4 Coherence and moral fruit are evidence

They can be. Coherence is a virtue when rival accounts differ in contradiction, simplicity, or fit with independently supported beliefs. Moral and communal outcomes can also bear on claims that predict them. The criticism is not that these considerations count for nothing.

The difficulty is comparative underdetermination. Many incompatible systems can be coherent, and many communities produce mixed but meaningful fruit. Public use requires an independently specified link between the outcome and the central claim, a live comparison, and symmetric treatment of contrary cases. Coherence and fruit become captured only when the home system sets every relevant benchmark and excuses every failure.

8.5 The standard privileges outsiders

Outsiders can be uninformed, hostile, or careless. The independent-loss condition gives them no veto. Expertise remains necessary, and not every objection deserves uptake. The standard asks for a channel through which a competent challenge can matter even when it threatens central authority. Longino's emphasis on criticism and uptake does not require equal expertise; it resists arrangements in which social power makes criticism causally idle (Longino 2002).

The strongest design combines insider competence with outsider correction. Insiders understand the tradition's content and strongest reasons. Outsiders detect assumptions that have become invisible within it and supply rival explanations. Neither position guarantees accuracy. A practice becomes more publicly credible when both can correct it.

8.6 A religion may aim at faithfulness, not public proof

Then much of the criticism does not apply. A liturgy, community discipline, or devotional interpretation may openly aim at faithfulness, formation, or meaning. Its success can be assessed within that practice without being advertised as public evidence for contested supernatural claims.

The problem arises at export. If the same internal victory is used to persuade outsiders that a doctrine is true, the practice has crossed jurisdictions. It must then show why its success conditions track a shared reality rather than only faithful participation. Clear limitation is a legitimate response; it is not an epistemic defeat unless the narrower practice continues to claim the stronger warrant.

9 Conclusion

Religious inquiry can be learned, coherent, demanding, and sincere. It can also generate genuine internal entitlement and practical value. The error is to let those achievements silently settle a different question: whether a contested claim about a shared world has earned public warrant.

The game metaphor identifies where the silent transfer occurs. Rules admit evidence, authorize moves, designate adjudicators, and define victory. When the challenged claim or authorities dependent on it control admission, negative scoring, or revision consequence, the system may continue to win without being able to lose. This is referee capture in the structural sense.

The independent-loss condition supplies a bounded test. Some feasible challenge must be able to enter, count against the central claim, and impose epistemic cost under procedures the claim does not fully control. The standard is not a demand for neutrality, experimental science, or universal agreement. It is a demand that resistance survive interpretation long enough to matter.

The result concerns jurisdiction rather than truth. A religious authority may decide what counts as orthodox within its community. It may guide lives, preserve identity, and articulate profound meaning. But when its central claims are offered as public descriptions of reality, the decisive question remains: what can make them lose, who is permitted to say so, and what changes when they do?

A Formal Reconstruction of the Central Thesis

A.1 Ordinary-language argument

- P1.** A rule-governed practice can award an internal victory when a claim satisfies its accepted sources, moves, authorities, and success conditions.
- P2.** Satisfying those internal conditions does not by itself show that the claim accurately describes a reality shared with people who reasonably dispute the conditions.
- P3.** If an internal victory is to supply public warrant, the practice must contain at least one independently constrained route by which a feasible challenge can count against the claim and require some reduction or revision.
- P4.** Referee capture blocks that route when authorities dependent on the challenged claim control admission, negative scoring, or revision consequence so completely that no feasible challenge can impose epistemic cost.
- P5.** Therefore, where referee capture removes every independent loss condition, an internal victory does not supply public warrant for the central claim.
- P6.** This conclusion denies one route to public warrant; it does not entail that the central claim is false, meaningless, or incapable of receiving warrant from some other source.

The inference from P3 and P4 to P5 is valid once the terms are fixed: public warrant through the in-game victory requires an independent loss condition, while complete referee capture removes every such condition. P3 is the main philosophical premise. A critic may reject it by arguing that public warrant needs no independently constraining route to error correction. P4 is the main application premise. It must be established for a particular practice rather than inferred from the mere fact that the practice has authorities.

A.2 Symbols

Let the symbols mean the following:

- G : the epistemic game being evaluated.
- C : a central claim made within G .
- d : a possible challenge—an observation, document, experience, argument, or pattern of results.
- $F(d)$: d is feasible enough to be a genuine candidate challenge, rather than a contradiction or an impossible fantasy.
- $A_G(d)$: G admits d for consideration.
- $N_G(d, C)$: G allows d to count negatively against C .
- $R_G(d, C)$: if d is accepted, G requires some reduction of confidence in C or some material revision of it.
- $I_G(d, C)$: admission, negative scoring, and revision are constrained in a way not fully controlled by C or by authorities whose relevant standing depends on C .
- $L_G(C)$: G has an independent loss condition for C .
- $K_G(C)$: every otherwise feasible loss route for C is blocked; this is complete referee capture with respect to C .

- $W_G(C)$: an in-game victory in G supplies public warrant for C .

The subscripts matter. A_G , N_G , and R_G describe what the particular practice G permits. Another practice may admit the same challenge but score it differently. I_G then asks whether the decisive procedures are sufficiently independent of the target claim to prevent a guaranteed home victory.

A.3 The loss condition

Define an independent loss condition as follows:

$$\forall G \forall C [L_G(C) \leftrightarrow \exists d [F(d) \wedge A_G(d) \wedge N_G(d, C) \wedge R_G(d, C) \wedge I_G(d, C)]] \quad (1)$$

Plain-English annotation. The opening quantifiers in Equation (1) mean that the rule applies to any game G and central claim C . Such a practice can genuinely lose on C only if at least one realistic challenge d can complete the whole route. The challenge must be possible, admitted, allowed to count against C , connected to a real revision, and evaluated under procedures that C does not fully control. If even one link is missing, that particular challenge is not an independent loss condition.

Complete referee capture is the absence of any such completed route:

$$K_G(C) \leftrightarrow \neg \exists d [F(d) \wedge A_G(d) \wedge N_G(d, C) \wedge R_G(d, C) \wedge I_G(d, C)]. \quad (2)$$

Plain-English annotation. Equation (2) does not say that participants never criticize one another. It says that no feasible challenge can pass all five checkpoints and impose a loss on the central claim. The practice may allow debate while ensuring that the protected claim always survives.

Equations (1) and (2) yield the following relation:

$$K_G(C) \rightarrow \neg L_G(C). \quad (3)$$

Plain-English annotation. Equation (3) follows directly from the definitions. If complete capture means that no independent loss route exists, then the game lacks an independent loss condition. This step is logically secure; the difficult factual question is whether a real practice actually satisfies the definition of complete capture.

A.4 From loss exposure to public warrant

The governing normative premise is:

$$W_G(C) \rightarrow L_G(C). \quad (4)$$

Plain-English annotation. Equation (4) says that an in-game victory can provide public warrant only if the game can genuinely lose on the same claim. This is not a definition of all knowledge. It is a proposed requirement for exporting an internal victory as a reason that people outside the game should accept C . A defender may challenge this premise, but should then explain how a guaranteed internal win can discriminate truth from protected error.

Taking the logical opposite of both sides in the valid way called *contraposition* gives:

$$\neg L_G(C) \rightarrow \neg W_G(C). \quad (5)$$

Plain-English annotation. Contraposition preserves the meaning of a conditional: if public warrant requires a loss condition, then lacking that condition means the in-game victory does not supply public warrant. This line does not say that nobody may rationally believe C for other reasons.

Now combine Equations (3) and (5):

$$K_G(C) \rightarrow \neg W_G(C). \quad (6)$$

Plain-English annotation. Equation (6) is the paper’s bounded formal result. Complete referee capture prevents the internal victory from doing public evidential work. The conclusion concerns the evidential route from the game to public warrant, not the truth value of C .

To state that final limit explicitly:

$$\neg W_G(C) \not\Rightarrow \neg C. \quad (7)$$

Plain-English annotation. The crossed implication in Equation (7) means “does not logically entail.” A claim can be true even when one argument for it fails. The formal result therefore rejects an inference from captured victory to public warrant; it does not deduce atheism, religious error, or the falsity of any specific doctrine.

A.5 Validity, support, and possible defeat

The derivation from Equations (3)–(5) to Equation (6) is valid. Its sound application depends on two questions that symbolic notation cannot settle. First, is Equation (4) a justified norm for public warrant? The paper supports it through the need to distinguish truth-sensitive success from a success condition that guarantees itself. Second, does a particular practice satisfy $K_G(C)$? That requires evidence about its actual admission rules, scoring practices, authorities, and revision history.

A religious defender can answer constructively by identifying a specific d and showing each part of Equation (1): why d is feasible, how it can enter, how it would count against C , what revision would follow, and what prevents C -dependent authorities from neutralizing the loss. That response would not evade the test. It would show that the tested claim is not completely captured.

References

- Berger, Peter L. and Thomas Luckmann (1966). *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. Garden City, NY: Doubleday.
- Boudry, Maarten and Johan Braeckman (2012). “How Convenient! The Epistemic Rationale of Self-Validating Belief Systems.” In: *Philosophical Psychology* 25.3, pp. 341–364. DOI: [10.1080/09515089.2011.579420](https://doi.org/10.1080/09515089.2011.579420).
- Goldman, Alvin I. (1999). *Knowledge in a Social World*. Oxford: Oxford University Press. DOI: [10.1093/0198238207.001.0001](https://doi.org/10.1093/0198238207.001.0001).
- Huizinga, Johan (1949). *Homo Ludens: A Study of the Play-Element in Culture*. London: Routledge and Kegan Paul.
- Kuhn, Thomas S. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Longino, Helen E. (1990). *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press. DOI: [10.2307/j.ctvx5wbfz](https://doi.org/10.2307/j.ctvx5wbfz).
- (2002). *The Fate of Knowledge*. Princeton: Princeton University Press. DOI: [10.2307/j.ctv2tvzv0](https://doi.org/10.2307/j.ctv2tvzv0).
- Plantinga, Alvin (2000). *Warranted Christian Belief*. New York: Oxford University Press. DOI: [10.1093/0195131932.001.0001](https://doi.org/10.1093/0195131932.001.0001).
- Suits, Bernard (1978). *The Grasshopper: Games, Life and Utopia*. Toronto: University of Toronto Press. DOI: [10.3138/9781487574338](https://doi.org/10.3138/9781487574338).
- Wittgenstein, Ludwig (1953). *Philosophical Investigations*. Ed. by G. E. M. Anscombe and Rush Rhees. Trans. by G. E. M. Anscombe. Oxford: Basil Blackwell.