Sophisticated Ignorance and Our Relation to Reality

Ramin Saadat
Independent Researcher; Founder, British Institute for Ignorance Studies (BIIS) (biis.uk)

AI Disclosure and Translation Statement:
The original manuscript of this work was authored in Farsi (Persian) by the researcher. The English version was produced through AI-assisted translation and linguistic refinement, utilizing Google Translate and Large Language Models. The author has meticulously reviewed, edited, and verified the final English text to ensure that all theoretical nuances, philosophical arguments, and conceptual frameworks accurately reflect the original Persian content. The author remains fully responsible for the integrity and accuracy of the translation. The original Persian manuscript is available from the author upon reasonable request.

Abstract

This article is the third and final part of a trilogy that asks who sets the reference domain of ignorance, that is, the boundary that determines what counts as ignorance. Two earlier articles showed that, in the negative and positive faces of ignorance, this domain is neither external nor discoverable. This article pursues the same question for the third face, 'sophisticated ignorance'. It shows that here too the domain cannot be discovered, but that it is set for a different reason: the system sets it itself, and this shapes the system's relation to reality.

The article proposes an operational criterion for this claim, 'unabsorbed resistance of reality'. The criterion needs no judgement by the researcher about the truth of a belief. It shows only that a system received feedback from reality and did not change its central conclusion. This failure of absorption takes three states (non-registration, perpetual deferral and normalisation of deviance), which the system reaches by three modes (endogenisation of origin, endogenisation of competence and reinterpretation of mechanism). The article separates this mechanism from neighbouring concepts, from compound ignorance and the Dunning-Kruger effect to epistemic circularity and the echo chamber. It then traces the mechanism in six documented and independent cases, from the creationist RATE project and the flat-Earth movement to the Challenger and Columbia shuttle disasters and the Nimrod crash. The article openly accepts two limitations: the selective choice of cases, and the fact that the 'relation to reality' itself cannot be measured. It concludes that sophisticated ignorance begins where reality, despite repeated opportunities, is no longer a match for the system. Reality is not left unheard. It is heard, and then rewritten.

Keywords: sophisticated ignorance, resistance of reality, normalisation of deviance, ignorance studies, reference domain

---

Introduction and Background

The common expectation is that reality corrects ignorance by itself. Someone who ignores the signals of their environment will sooner or later face the consequences and correct their belief.

Sometimes, however, this does not happen, from folk beliefs in everyday life to institutional failures in science and engineering. In such cases, feedback from reality arrives again and again, yet the false belief persists.

This article asks why. It looks for the answer neither in a lack of information nor in the intention of a particular agent. It looks in a mechanism by which a belief system treats feedback from reality in one of three ways. The system may not register the feedback at all. It may recognise the feedback but defer its resolution indefinitely. Or it may register the feedback but, because repeated instances have no consequences, come to treat it as normal. In all three cases, despite repeated opportunities and despite the resistance of reality, the system's central conclusion does not change.

Before going further, one distinction is necessary. The author's earlier work uses three tripartite classifications. They are easily confused, although they have no relation to one another:

- the three faces of ignorance (negative, positive, sophisticated);
- the three research currents in ignorance studies (agnotology, epistemic injustice, critical race theory);
- and the three instances of the 'illiterate' in the 'sinister triangle of ignorance' (superstition, pseudo-science, nonsense).

This article concerns only the faces of ignorance. It does not enter the other two classifications.

This article is also the third part of a trilogy that has examined three faces of ignorance. The two earlier articles pursued one question: who sets the 'reference domain of ignorance', that is, the boundary that specifies 'what counts as ignorance'?

The first article concerned the negative face, which defines ignorance as the 'absence of knowledge'. It showed that, unlike other forms of absence, ignorance has no reference domain that can be discovered from its own nature. Darkness, blindness and missing data each have a domain. Either a causal mechanism produces that domain or an authoritative institution sets it. Ignorance has neither. The conclusion was that every general classification of ignorance must set its own domain, not discover it.

The second article pursued the same question for the positive face and reached a subtler answer. What reached an official boundary was always the 'expressible and clinical' aspect of the pattern, not its cognitive claim. As a result, the answer was negative for both faces, and the domain was not closed.

This third article concerns 'sophisticated ignorance' and, in the final analysis, our relation to reality. This face has had deadly consequences at times, as in the explosions of the Challenger and Columbia space shuttles. At other times it appears neither lethal nor dangerous. Examples include the belief in a 6,000-year-old Earth within an official research project (RATE), the belief that attributes healing properties to the urine of livestock, and a pseudo-scientific movement that holds the Earth to be flat, along with other cases that all of us have heard of and seen.

The natural expectation is that this article will repeat the same negative answer. Its answer, however, is of a different kind. In the sophisticated face, the reference domain is set, although the one who sets it is the 'system' itself. The article proposes that this endogenous setting explains the resistance to correction. When the system itself decides what counts as a 'gap' and what counts as a 'response to the gap', the criterion of assessment becomes part of what was meant to be assessed. Non-registration, deferral and normalisation are each a way of exercising this same authority. The trilogy reaches its conclusion at this point.

Heterogeneity in how ignorance is understood, both open and hidden, is not new in the literature of ignorance studies. For example, Le Morvan and Peels (2016) distinguish between the 'standard view' (ignorance as lack of knowledge) and the 'new view' (ignorance as lack of true belief). El Kassar (2018) shows that the contemporary literature, without saying so, draws on at least three heterogeneous conceptions: lack of knowledge or true belief, false attitudes that are actively guarded, and substantive epistemic practice. For this reason, she seeks a unified framework for this diversity.

These studies, however, from Le Morvan and Peels and El Kassar (2016 and 2018) to Nottelmann (2016) and Peels (2023), work at the level of 'what ignorance is': what components it consists of.

The question of this article lies at a different level: why do some forms of ignorance resist correction? This question concerns a mechanism, not the composition of components.

In this article, therefore, 'sophisticated ignorance' is the name of this mechanism: a behavioural pattern of the system in the face of feedback, not a neural, cognitive or evolutionary mechanism. A neural, cognitive or evolutionary explanation of this pattern is a separate question, and this article does not enter it. What the article offers in answer is a structural explanation (see the Conclusion).

A criterion formulated in Section 1 separates this mechanism from other forms of ignorance: 'unabsorbed resistance of reality'. This is feedback that reality gives despite repeated opportunities, but that the system does not count as a reason to revise its central conclusion. The feedback may go entirely unseen, or the system may see it and defer it, or the system may see it and treat it as normal. As far as the author has been able to establish, the traditional frameworks of ignorance studies have not discussed this criterion so explicitly.

The structure of this article is as follows:

- Section 1 formulates the definition of sophisticated ignorance and the criterion that distinguishes it.
- Section 2 examines the three modes by which a system rewrites feedback.
- Section 3 separates sophisticated ignorance from its nearest neighbouring concepts in the literature of ignorance studies.
- Section 4 presents six documented cases of this mechanism from different fields.
- Section 5 sets out, frankly, the limits of the evidence and of the claim.

---

1. Definitions and the Criterion of Distinction

In this article, 'sophisticated ignorance' is a condition in which a system (a mind, a network or an institution) does not change its central conclusion despite repeated feedback and warnings from reality, unless a disaster or external pressure forces it to.

To make the concept clear, we must also define several other terms:

- **'Central conclusion':** the conclusion on which the 'coherence of the system' depends, and whose abandonment would transform the whole structure (for example, the claim that 'the Earth is six thousand years old' in the project known as RATE).
- **'Gap':** the distance between this conclusion and the reality that the feedback has revealed.
- **'Unabsorbed resistance of reality':** feedback that reality gives to the system and that the system does not count as a reason to revise its central conclusion. The problem here is not the absence of feedback but its non-absorption. This is the criterion of this article.

Several points deserve attention.

First, the term 'sophisticated ignorance' entered the author's earlier work from several angles:

- from the angle of instruments, in the reconstruction of the 'conceptual model of the sources of ignorance', where intellectual tools and technical terms paradoxically consolidate a person's own ignorance;
- from the angle of historical lineage, in the re-reading of the path from *agnoia* to *amathia*;
- and from the angle of methodological consequence, in the discussion of unmeasurability, where the signals reverse.

This article claims that these angles are not rival definitions but 'cross-sections of a single core'. It seeks to formulate that core explicitly and to fix its boundaries.

Second, the concept of 'sophisticated ignorance' is not the same as 'compound ignorance' (in the Islamic tradition) or the 'Dunning-Kruger effect'. Section 3 discusses this in detail.

From these definitions, the definition and the range of instances of sophisticated ignorance rest on the 'behaviour' of the system in the face of feedback, not on the 'content' of a belief. The reason is that a belief is by nature 'unobservable', whereas behaviour is 'observable': an external observer can see a change in behaviour under the pressure of reality.

To say that feedback has 'arisen', however, we need a criterion that does not depend on the researcher's judgement about the truth. This criterion comes either from the system itself or from a source independent of it. From the system itself, it is a condition that the system had declared as a refutation or a requirement (for example, the '24-hour sun' among flat-Earthers, or the 'no debris' requirement at NASA). From an independent source, it is, for example, a warning from an external body. In this way the researcher need not judge the truth. It is enough to show that, by one of these criteria, feedback arose and the system did not change its 'central conclusion'.

From this perspective, the 'failure to absorb feedback' takes three forms. Documents separate the three, not the researcher's interpretation:

- **'Non-registration':** the feedback is never registered as a 'gap'. In other words, the system's documents contain no record of it as a gap.
- **'Perpetual deferral':** the gap is registered and named, and it continues to be called an 'unsolved problem' (for example, 'we know this is still dangerous'). Its resolution, however, is deferred continuously, and the central conclusion is never revised.
- **'Normalisation of deviance'** (Vaughan, 1996): the gap is registered, but each repetition without consequence withdraws the label 'gap', and what was a deviation comes to count as 'normal'. Its documentary sign is a 'change in the system's vocabulary', for instance from 'deviation' to 'acceptable risk'.

Three conditions are necessary for this 'normalisation of deviance', and all three must be documented:

1. an initial, dated record;
2. at least two later events with a similar or worse deviation and no obvious consequence;
3. and a documented change in the system's vocabulary to describe it.

The section on evidence gives examples.

The table below shows these three states of the 'failure to absorb feedback':

State (main example)	Is the gap registered?	What happens to the gap	Documentary sign
Non-registration (urine of livestock)	No	From the start, the feedback does not count as qualifying for the label 'gap'	The system's documents contain no record of the feedback as a gap
Perpetual deferral (RATE)	Yes	The gap is named and keeps the label 'unsolved problem', but its resolution is deferred continuously, and the central conclusion is never exposed to revision	The system still calls it an unsolved problem and only postpones action
Normalisation of deviance (Vaughan, 1996; Challenger, Columbia, Nimrod)	Yes, but only temporarily	With each repetition without consequence, the label 'gap' is withdrawn, and what was a deviation counts as 'normal' or 'acceptable'	The system's vocabulary changes over time

A concept must exclude something, or it is empty. A system belongs to this class if it has not changed its central conclusion in the face of many accumulated opportunities for revision. A procedure that has not once ruled against itself among a mass of such opportunities is not a procedure but a formality. A system that changes its central conclusion as feedback accumulates, before any disaster or external pressure, does not belong to this class.

By this criterion, the following cases fall outside the definition of sophisticated ignorance.

Normal science. Normal science does make errors. It falls outside the class because it has a procedure by which external findings can change the framework itself. History shows that it has used this procedure, although sometimes after decades of resistance.

Simple false beliefs. Someone who believes 'I use only ten per cent of my brain' does not belong to this class, because, once corrected, they usually change in the direction of the correction (Nyhan, 2021). This is not 'ignorance' but 'error' in the Cartesian sense (Descartes, 1641/1996, Fourth Meditation). In error, the will gives its judgement before the understanding has perceived clearly, and as soon as the understanding is corrected, the will follows. In this sense, error is neither a 'lack' (negative ignorance) nor an 'attribution to an external agent' (positive ignorance). The earlier article on the negative face already distinguished error from ignorance in its discussion of Descartes.

Doubt. A parent who is hesitant about a vaccine and changes their mind after a conversation is not an instance of this condition. The 'spectrum of doubt' runs from simple doubt that is open to correction to flat rejection that counts no feedback, not even official refutation, as a reason for revision. It includes both ends, but only part of it stays in this class: the part that does not change its central conclusion despite repeated and accumulated opportunities. The intensity of doubt does not fix the boundary. Behaviour in the face of feedback does.

Denial. In denial, the system identifies the 'contradicting evidence' but does not recognise it as evidence, and rejects it. None of the three states described above is the same as 'denial', because:

- in non-registration, the system does not identify the feedback as a gap;
- in perpetual deferral, the system identifies the feedback and does not reject it, but leaves it suspended;
- in normalisation of deviance, the system identifies the feedback and later withdraws the description 'gap' from it.

Normal science and simple false beliefs fall outside this class. This condition, by contrast, appears in religion and pseudo-science, in therapeutic beliefs and in engineering institutions, in different forms. This diversity follows from defining the core of sophisticated ignorance at the level of mechanism, not content. One mechanism takes different forms in different settings. Only the evidence of this article supports the claim, and the section on limitations sets out its limits.

The next section examines the three modes of rewriting by which this condition arises.

---

2. Three Modes of Rewriting

The previous section showed that the 'failure to absorb feedback' takes three states:

- the gap is never registered;
- the gap is registered, but its resolution is deferred continuously;
- or the gap is registered, and each repetition without consequence withdraws its label.

As we can see, these three states show where the gap between the central conclusion and the feedback from reality finally ends up. But how does the system arrive at this outcome? What path does it follow to reach one of the three states? This section addresses that question.

In the view of this article, feedback must pass through three gates before it can affect the central conclusion. The system checks three things:

- **It must count as new and external.** The system must see it as something new, and from elsewhere, not as something the system already knew or said.
- **Its speaker must be competent.** The source of the feedback must count, in the system's view, as qualified to judge the matter. That means the system regards the source as a member of an institution that the system itself recognises as authoritative for judging in that field.
- **It must connect to the central conclusion.** The system must establish a causal or logical link between the feedback and its central conclusion, and must not treat the feedback as unrelated to the conclusion.

Sophisticated ignorance closes none of these three gates by explicit rejection. It does not engage in denial. Instead, it rewrites the feedback. It changes the meaning that the feedback has for the system, so that the feedback no longer counts as a reason against the central conclusion. It does this in three modes, and each mode closes one of the gates named above.

First mode: endogenisation of origin. This mode closes the 'gate of novelty'. The system does not reject the external finding. It accepts the finding, but with the verdict that the finding is not new and 'was never absent'. The literature reports an example in the reading known as the 'scientific miracle' of the Qur'an (see Bigliardi, 2017; Guessoum, 2011; Edis, 2007). Here a scientific finding does not count as a gap. Readers re-read it as something already stated in advance: something old and repetitive that deserves no attention and can have no effect.

This mode, however, rejects exactly what must be accepted for a finding to count as a 'gap': that the finding was new and that the system did not have it before. By rejecting novelty, the mode leaves the resistance that reality has shown entirely unseen. The system does not see it and then justify it. A sign of this mode is that such readings always come after the finding, never before it. If the text really knew the finding in advance, one should have been able to predict a new finding from the text, and such a prediction never happens.

Second mode: endogenisation of competence. This mode closes the 'gate of the speaker's competence'. The feedback does not count, not because anyone has examined its content and found it false, but because its speaker 'has already lost standing'.

What loses standing here is usually the 'institution', not the 'method'. For example, flat-Earthers run experiments themselves and borrow the procedure of measurement, but they do not accept the judgement of the scientific institution. Similarly, the World Health Organization's warning about the urine of livestock is attributed to a political motive, not examined (Section 4).

Third mode: reinterpretation of mechanism. This mode closes the 'gate of connection'. The system accepts both the finding and the speaker. What it rewrites is the consequence that it draws from them: either the causal relation between the finding and the conclusion, or the threshold that determines what counts as a risk. In the RATE project, the project did not reject the radiometric dating data. It accepted them, and rewrote the causal relation between the data and the age. At NASA, engineers recorded the erosion of the seals and the shedding of foam, but the organisation rewrote the relation between these events and the danger: 'deviation' was downgraded to 'acceptable risk' (Section 4).

Note that 'reinterpretation' is not in itself a fault. Normal science sometimes does the same. For example, when the orbit of Uranus did not match the Newtonian prediction, astronomers did not abandon the theory. They posited an unknown planet. That hypothesis, however, yielded a new prediction, and the discovery of Neptune (1846) tested it independently.

What makes the third mode sophisticated ignorance, unlike the case of Uranus, is a reinterpretation that yields no new prediction and merely 'justifies' data already seen, while the central conclusion stays untouched each time. This is what Lakatos calls a 'degenerating research programme' (Section 3).

What, then, is the relation between these three modes and the three states listed in Section 1? The modes and the states are not two independent axes. They are two views of a single process. The states say what happens to the gap, and the modes say which gate the system closes to bring that about. The first and second modes act before the feedback counts as a gap at all, and so they lead to non-registration. The third mode acts after registration. Depending on whether the system keeps or withdraws the label 'gap', it leads to perpetual deferral or to normalisation of deviance.

Another point needs explaining here. The 'system' in this article, as Section 1 said, can be a mind, an institution or a network. What we have seen so far occurs at the level of the mind and the institution. It is worth saying a little about the level of the network too.

Although at the network level we face a combination of agents, at least two modes, disqualification and reinterpretation, operate at the same time at this level too, without needing the intention or coordination of any agent. Evidence shows (McPherson et al., 2001) that people tend to bond with others like themselves. At the scale of a network, this tendency leads to clusters in which the standing of the person who made an argument, not an assessment of its content, fixes the credibility of the argument. This is the structure of trust that Nguyen describes in his discussion of the echo chamber (Section 3). In such a cluster, two modes work at the same time:

- disqualification of the external source (second mode);
- and reinterpretation, to fit the feedback into the existing model (third mode).

As a result, one member's registering of a gap does not, by itself, change the system (here, the network). Either the network repairs its model, or it expels that member. Section 4 gives a documented example of this pattern.

Thus the three modes make clear how the system rewrites feedback. The next question is what separates 'sophisticated ignorance' from neighbouring concepts, from compound ignorance to the echo chamber and epistemic circularity. The next section addresses these boundaries.

---

3. Sophisticated Ignorance and Neighbouring Concepts

To clarify the boundaries, this section compares sophisticated ignorance with the nearest related concepts in the field of ignorance studies.

3.1. Compound ignorance and the Dunning-Kruger effect (Kruger and Dunning, 1999). Compound ignorance is a lack of awareness of one's lack of knowledge. It operates at the individual level and involves no institutional or network mechanism. Sophisticated ignorance is limited neither to the individual nor to 'not knowing what one does not know'. A system may see the gap and name it (perpetual deferral), or register the gap and count it as normal (normalisation of deviance).

The Dunning-Kruger effect is likewise a correctable defect in an individual's power of judgement. As real skill increases, the gap closes, and the mechanism corrects itself. In sophisticated ignorance, more intellectual tools do not close the distance, and sometimes they widen it. The two move in opposite directions, and this reversed prediction is the test that separates them.

3.2. Motivated reasoning. Kunda (1990) showed that a person's goals affect how they assess evidence, although the person must still be able to justify the conclusion. Kahan et al. (2012) showed that, on the risk of climate change, higher scientific literacy and numeracy went with more polarisation, not less. These studies explain a psychological mechanism at the individual level: why a person reads evidence in a biased way. Sophisticated ignorance describes a condition defined by the system's behaviour in the face of feedback, and it does not require a motive. Motivated reasoning can be one way in which the condition is realised at the level of the mind, but the definition does not depend on it. Kahan's finding agrees with the claim of the previous paragraph, that more intellectual tools can widen the gap, although only within that one domain.

3.3. Epistemic circularity (Alston, 1993). This theory shows that the justification of the reliability of a faculty rests on premises that are themselves products of that same faculty. The subject there, however, is the 'structure' of justification, and the difference does not lie in distance from the conclusion. It lies in the kind of exclusion. In epistemic circularity, the proposition is formulated and present in the field, and the subject can refuse it. In what this article describes, the proposition never takes a form about which one could ask a question. An external proposal is also retranslated

into the categories of the system itself. The difference is between refusing to accept something and failing to grasp something as a thing that one could accept or refuse.

The RATE project (Section 4) shows this distinction in documents. The project did not reject the data. It accepted and re-read them, and the system's own official document is the evidence of this re-reading, not an external report on the behaviour of followers. This case moves the claim of retranslation from purely theoretical analysis closer to documented evidence, although it remains a single case, not a systematic sample (Section 5).

3.4. Coherentism and the isolation objection. The 'isolation objection' is one of the classic objections to epistemic coherentism. The literature has also examined it in connection with 'Lehrer's theory' (Lehrer, 2000), but its best-known formulation is BonJour's (1985, p. 108).

The difference here is that the coherentist at least knows that their system is a set of beliefs, and can ask whether a new belief exists that is not yet in it. In sophisticated ignorance, this very question loses its object, because the system itself fixes the boundary of what could have been a new belief.

3.5. Immunising strategies (Albert, 1968; Boudry and Braeckman, 2011). These are defensive mechanisms. They ensure that contradicting evidence cannot do harm. They presuppose that the system has identified the contradicting evidence and then neutralised it. In non-registration, sophisticated ignorance acts one step earlier and makes the work of defence unnecessary. In perpetual deferral and normalisation of deviance, the system does identify the evidence, so the two concepts overlap to that extent. The difference lies in scale (see the next paragraph).

3.6. Degenerating research programme (Lakatos, 1970). An immunising strategy describes a single move: contradicting evidence arrives, and an auxiliary hypothesis neutralises it. What happens in the state of 'perpetual deferral' is a longer path. The system recognises the gap, names it and even writes about it. Then, generation after generation, it builds new auxiliary hypotheses without once abandoning its central conclusion. This is exactly what Lakatos calls a 'degenerating research programme': a programme that produces no new predictions and only patches, one after another, so that its hard core stays untouched. The difference from an immunising strategy lies in scale. One analyses a single move. The other analyses a history of moves, all in one direction. The RATE project is a clear example of this pattern (Section 4).

3.7. Kuhn (1962). Kuhn sees two rival frameworks as symmetrical in that each brings its own criteria of assessment. In his view, however, accumulated anomalies eventually lead to a crisis and a change of framework. In other words, the system finally absorbs the feedback. In sophisticated ignorance, accumulated feedback does not reach a crisis. In this sense, the claim of this article is asymmetrical: a framework builds its own criterion and remains immune to external assessment.

3.8. The echo chamber and the epistemic bubble (Nguyen, 2020). Nguyen separates two phenomena. In an epistemic bubble, relevant voices are left out through inattention. In an echo chamber, outsiders are systematically discredited, and exposure to opposing evidence can strengthen belief. The bubble concerns what reaches the person. The echo chamber, in Nguyen's sense, is very close to the second mode and to the network level. What this article adds is that the mechanism does not come only from discrediting the source. In the cases of RATE and NASA, the

systems accept the source and rewrite the causal relation. The definition rests on the system's behaviour in the face of feedback, not on a structure of trust.

Here we should also mention the concept of 'emerged ignorance', discussed elsewhere by the author. 'Emerged ignorance' is 'structural' and 'agentless', in the sense that no single agent designs or coordinates it, and it resists correction. This concept is not a rival to 'sophisticated ignorance'. It is the same mechanism appearing at the network level.

3.9. Nottelmann and Peels (Nottelmann, 2016; Peels, 2023). Sophisticated ignorance does not fit into Nottelmann's (2016) threefold map of ignorance, which Peels (2023) takes up in his book: propositional, objectual and technical. Propositional ignorance requires a formulated proposition, which does not exist here. Technical ignorance concerns 'how to do something', not this condition. The nearest neighbour is objectual ignorance, that is, ignorance of a domain, not of a proposition. The difference is that here the system itself makes it impossible to formulate that object.

One can also trace the phenomenon and mechanism of 'sophisticated ignorance' in history, and find evidence of its conceptual precedent.

Take the thought of al-Ghazali, for example. The mechanism has a history of roughly nine hundred years, at least as a conceptual formulation, though not as documented evidence of a system's behaviour. In *Fada'ih al-Batiniyya* (487–488 AH / 1094–1095 CE), al-Ghazali describes the Ismaili doctrine of *ta'lim* as follows. Valid knowledge comes only from the infallible imam. The system has already declared any external rational argument that challenges this claim to lack standing, because only the imam has the right to judge. In al-Ghazali's account, the followers could not prove the imam's infallibility independently of the imam's own word.

If this account is right, it is the same asymmetry that Section 2 described in the second mode. But the description is not direct evidence. *Fada'ih al-Batiniyya* is a polemic commissioned by the Abbasid caliphate at the height of its military conflict with the Nizari Ismailis. It is not an official document of the Ismaili system itself, and it does not meet the criterion that we valued in the case of RATE (a document from within the system).

Contemporary scholarship (see Mitha, 2001) shows that al-Ghazali probably presented the doctrine of *ta'lim* in a more caricatured form than the Ismailis themselves believed. This case therefore belongs under conceptual precedent, not evidence. It shows that a sharp observer could have formulated the mechanism about nine hundred years ago. It does not show that the mechanism historically occurred in the Ismaili system.

---

4. Evidence

The previous section separated sophisticated ignorance from its nearest neighbouring concepts. We now show that this mechanism is more than a conceptual possibility. It has occurred in six documented and independent cases. In each of the six, a 'system-level' failure (not necessarily that of every individual member) to absorb feedback from reality appears:

4.1. The RATE project (Radioisotopes and the Age of the Earth). The RATE project was an official, funded research project (2000 to 2005) run by the Institute for Creation Research (ICR) and the Creation Research Society. It did not reject the radiometric dating data, which give the Earth an age of billions of years. It accepted the data. To reconcile them with the view of a 'six-thousand-year-old Earth', it proposed a hypothesis called 'accelerated decay', operating in short periods of history.

Vardiman and colleagues, in the project's official report (2005), turned to supernatural intervention to resolve the 'heat problem', that is, the enormous heat that such decay would release (Vardiman et al., 2005: 736).

Loechelt (2009) independently showed that the RATE model of helium diffusion in zircon contains technical errors: an incorrect effective radius, neglect of the high permeability of biotite, and an error in Gentry's estimate of the amount of helium produced. He also offered an alternative (multi-domain) model that made the same data consistent with an age of about 1.5 billion years.

As we can see, the appeal of the project's official report to 'supernatural intervention' is exactly what Lakatos calls a 'sign of a degenerating programme': an auxiliary hypothesis that does not 'predict' new facts and merely 'justifies' existing data after seeing it.

Loechelt's finding, independently of this, shows that even at the more technical level of the model, the same pattern of post-hoc justification recurs. The gap was registered and named, but its resolution was deferred continuously: perpetual deferral. The importance of this case is that the system's own official document, not an external interpretation of followers' behaviour, is the evidence of this reinterpretation.

4.2. The history of flat-Earthers. Flat-Earth thinking in our time follows a path from 'isolated denial' to 'networked deferral'. In 2018, Bob Knodel of the Globebusters group used a ring laser gyroscope, costing about twenty thousand dollars, and recorded a drift of nearly fifteen degrees per hour, exactly what the rotation of the Earth predicts. The documentary *Behind the Curve* recorded his reaction: 'That was a problem. Obviously we were not willing to accept that, and we started looking for ways to disprove that the [result] was the motion of the Earth' (*Behind the Curve*, 2018). By the definition in Section 1, this isolated reaction is denial, not non-registration. On its own, it is not an instance of sophisticated ignorance.

Six years later, a far larger opportunity arose. In the expedition 'The Final Experiment' (December 2024), a group of flat-Earthers and round-Earth defenders travelled together to Antarctica to observe the 24-hour sun directly. The flat-Earthers themselves had earlier treated this phenomenon as a 'decisive criterion of refutation'. Jeran Campanella, a leading figure of the movement, openly admitted after observing it: 'Sometimes you are wrong in life… I thought there was no 24-hour sun' (Al-Sibai, 2024). He later formally withdrew from belief in a flat Earth (ChurchLeaders, 2024). This is a negative case. It shows that sufficient and accumulated opportunity can change the central conclusion, at least at the level of one member, and in this way it frees the criterion of Section 1 from a logical circle.

At the level of the network, however, the story went differently. Austin Whitsitt, another member on the same trip, tried to fit the new data into a rearranged model of the flat Earth. He called it 'only a single data point' that neither refutes the flat Earth nor proves a globe (Hogue, 2025; ScienceAlert, 2024). This is reinterpretation of mechanism, of the same kind as we saw in the case of RATE. A larger part of the flat-Earth community called the trip itself fake and part of a conspiracy (Wikipedia contributors, 2025). This is disqualification of the source, this time even against one of their own. The result was that the system registered the gap once at the level of the individual but, at the level of the network, through two parallel mechanisms, stayed immune to any change in its central conclusion. This is exactly what Section 2 (network level) says about a credibility that depends on the position of a node, not on the content of the proposition: Campanella's testimony was either accepted as a forgivable error, with the model repaired, or rejected as betrayal.

4.3. The healing properties of domestic livestock urine. In several independent and geographically scattered cultural traditions, the folk belief in the healing properties of animal urine has coexisted with official institutional research. In one case, after identifying zoonotic transmission of a lethal respiratory virus, the World Health Organization formally issued a health warning. The public belief continued, and the warning was either ignored or attributed to a political motive (WHO, 2015; 2019).

In another case, a national government body for animal husbandry launched an official programme of research and public examination to establish scientifically the same properties. The programme included a claim of a clinical trial with several hundred participants and a high rate of effectiveness, which was never published (ThePrint, 2021; Rashtriya Kamdhenu Aayog, 2021).

In both cases the pattern is the same. The source of the feedback (the World Health Organization in one case, the absence of published evidence in the other) loses standing without anyone examining its content, either by attribution to a political motive or by simple neglect.

4.4. NASA and the normalisation of deviance: Challenger. The Space Shuttle Challenger disaster is the best-documented secular and institutional case in this article. In her foundational study (Vaughan, 1996), Diane Vaughan showed that engineers at Morton Thiokol, the maker of the solid rocket boosters, observed erosion of the O-ring seals at low temperatures on several flights before the disaster and formally recorded it in engineering reports. In Vaughan's own words, 'each time, evidence initially interpreted as a deviation from expected performance was reinterpreted as within the bounds of acceptable risk' (Vaughan, 1996). This pattern repeated over several flights. These were exactly the repeated and accumulated opportunities that the definition in Section 1 requires. It continued until the night before the Challenger launch, when engineers formally issued a warning and NASA managers, under production pressure, ignored it. This is an example of normalisation of deviance: the gap was registered, but with each repetition without consequence, the label 'deviation' was downgraded to 'acceptable risk'.

4.5. NASA and the Columbia shuttle. Seventeen years later, the same organisation had declared an explicit requirement in its original design: 'No debris may be shed from the critical zone of the external tank, on the launch pad or during ascent.' Yet after dozens of flights without accident, the shedding of insulating foam, which at the start breached exactly this requirement, was gradually

downgraded from a ‘deviation’ to a ‘routine maintenance concern’ (Columbia Accident Investigation Board, 2003). On 16 January 2003, a piece of foam broke off Columbia’s external tank and struck the leading edge of the left wing. On re-entry, the damage broke the shuttle apart. The Columbia Accident Investigation Board (CAIB), to which Diane Vaughan contributed as an expert, used in Chapter 8 of its official report, titled ‘History as Cause: Columbia and Challenger’, the very concept of ‘normalisation of deviance’ that she had coined seventeen years earlier for Challenger (Columbia Accident Investigation Board, 2003). The institutional reforms after Challenger, namely the Rogers Commission, the redesign of the joints and the new safety procedures, had not destroyed the mechanism. They had only changed its object, from the O-ring to the foam.

4.6. The Royal Air Force, Nimrod XV230 and the normalisation of deviance. On 2 September 2006, a Nimrod MR2 maritime patrol aircraft of the Royal Air Force, number XV230, caught fire after air-to-air refuelling over Afghanistan, and all fourteen people on board died. The source of the fire was leaking fuel that came into contact with a hot air duct. The independent review by Haddon-Cave (Haddon-Cave, 2009) judged the accident avoidable and, among the twelve similarities it drew with the Columbia disaster, named ‘normalisation of deviance’. Fuel leaks in the Nimrod fleet had been recorded for years. Because none had led to a disaster, they came to seem routine and inevitable, and the very record of no disaster was taken as evidence of the system’s soundness, not as a warning.

The Nimrod safety case also rested on the assumption that an aircraft that had flown for thirty years was ‘safe anyway’. The value of this case is that its institution (the Royal Air Force and the Ministry of Defence, together with contractors) and its country are distinct from NASA. The review itself, not an outside researcher, explicitly linked the similarity to Columbia and Challenger.

These six cases come from different fields: religion (creationism) and science (RATE), a pseudo-scientific movement (flat-Earthism), a therapeutic belief and a government institution (livestock urine), and engineering and safety (NASA and the Royal Air Force). This diversity of fields separates them from examples drawn from a single setting, although the section on limitations discusses their geographical concentration. The second case (the flat-Earthers) is the only place in this article where the criterion of Section 1 is tested against a documented negative case (Campanella’s change of mind), and so is freed from the charge of unfalsifiability. It also shows why an individual’s change does not lead to a change in the network. Cases four to six show that the mechanism is not limited to religion and pseudo-science, and that even serious institutional reforms do not necessarily destroy it. In secular engineering organisations, with trained engineers and formal documentation procedures, it has acted in the same way three times, in two institutions and two countries.

Classification of the six cases of evidence

Case	Mode (Section 2)	State (Section 1)	Documentary evidence in the text
RATE	Third mode	Perpetual deferral	'An auxiliary hypothesis that does not predict new facts…'
Flat-Earthers (network level)	Third mode (Whitsitt) and second mode (disqualification of the source), at the same time	Perpetual deferral and non-registration, at the same time	'Through two parallel mechanisms, stayed immune to any change in its central conclusion'
Livestock urine	Second mode	Non-registration	'The warning was either ignored or attributed to a political motive'
Challenger	Third mode	Normalisation of deviance	'Each time… was reinterpreted as within the bounds of acceptable risk'
Columbia	Third mode	Normalisation of deviance	'Was downgraded from "deviation" to "routine maintenance concern"'
Nimrod XV230	Third mode	Normalisation of deviance	'The record of no disaster was taken as evidence of the system's soundness, not as a warning'

What the table shows: each of the three states has at least one clear representative. RATE represents perpetual deferral, livestock urine represents non-registration, and the three engineering cases represent normalisation of deviance. The distribution, however, is unequal. Non-registration has only one clear representative. The flat-Earthers, at the network level, show two states at the same time. This is not a fault in the classification. Two parallel mechanisms operate together at the network level. The third mode appears in five of the six rows, so the table shows the diversity of states more than the diversity of modes.

These six cases, for all their diversity, are not enough on their own and without qualification. The next section takes up the limits of their selection and of their generalisation.

---

5. Limitations

The six cases in the previous section showed the mechanism. Before we draw any conclusion beyond the cases themselves, however, we must draw the limits of this evidence frankly.

First, the evidence of this article is six hand-picked cases, and the author alone classified them, with no independent coder. Three cases come from the range of religion, pseudo-science and therapeutic belief, and three are engineering and secular. But the 'institutional diversity' is limited.

Two cases come from a single institution (NASA, seventeen years apart), and only one (Nimrod) comes from another country and institution. Four of the six cases (RATE, the flat-Earthers, Challenger, Columbia) are American, and Nimrod is British. Only the livestock urine case comes from other cultural contexts. Among these cases, none comes from politics or from a computer algorithm (for example, a social media feed).

Second, the distribution of evidence across states and modes is unequal. Non-registration has only one clear representative (livestock urine). Normalisation of deviance has three, but all three are technical disasters with human casualties. Five of the six cases fit the third mode. The first mode, whose documented example is the literature on the scientific miracle, has not been studied as a case in this article.

Third, we always identify normalisation of deviance after a disaster, and this carries a risk of hindsight bias. Organisations that tolerated a deviation and suffered no disaster are not studied, and tolerating a deviation can sometimes be attributed to 'reasonable risk-taking'. The test of Section 1 (an initial dated record, later events, and a documented change in vocabulary) relies on the system's own criterion, not on the outcome. This keeps it fairly far from hindsight judgement, though not entirely. The Nimrod case, too, has been mapped through the lens of Columbia, so we must regard the similarity as both evidence and a product of the analyst's framework.

Fourth, the state of 'perpetual deferral' has two documented cases (RATE and the history of flat-Earthers), but this opens a new problem. Lakatos distinguishes between degenerating and progressive research programmes, and both can carry an unresolved anomaly for years. In standard physics, for example, the problem of dark matter has stayed unresolved for decades, and no one counts it a sign of sophisticated ignorance. Lakatos's criterion rests on whether a programme produces new predictions, not on the mere length of time without resolution. For RATE, this criterion has been applied and documented (Section 4: the reliance on supernatural intervention, and post-hoc fitting instead of prior prediction). For the flat-Earthers, this test has not yet been explicitly carried out.

Fifth, the claim that increasing intellectual tools intensifies ignorance (and so separates sophisticated ignorance from the Dunning-Kruger effect) has support in only two areas: the polarisation that Kahan and colleagues (2012) found, and the literature on the scientific miracle. For the central cases of Section 4, it is an open hypothesis, not a finding.

Sixth, and most fundamental, the definition of this article rests on 'relation to reality'. Automatic measurement of this relation requires a judgement about the 'truth of a belief'. Data-driven systems, however, hide this judgement behind a claim of objectivity and automation. Such a measure, at scale, must rely only on surface signals (fluency of expression, confidence, sophistication of argument), and these are the very signals that reverse in sophisticated ignorance: the more sophisticated the argument, the stronger the defence of the false belief. The gap between what the measure assesses and what it claims to assess therefore stays open, even with complete honesty on the part of the person who designs the measure. What this article claims is conceptual analysis, not an automatic measure. Conceptual analysis makes no claim to final access to truth, and it stays open to criticism.

Even so, what these six cases show, with all their limitations, remains unchanged. ‘Sophisticated ignorance’ begins where reality, however hard it presses, is no longer a match. Reality is not silent. It is heard, and then rewritten in an ‘artful’ manner: an art peculiar to human beings and to the institutions they have built.

Conclusion

The trilogy on ignorance, of which this article is the third and final part, began from a single question: who sets the reference domain of ignorance, that is, the boundary that determines what counts as ignorance?

The first article showed that, in the negative face, this domain, unlike darkness or blindness, cannot be discovered from its own nature. The second article showed that, in the positive face, what reached an official boundary was only the clinical and expressible aspect of the phenomenon, not its cognitive claim. In both faces, the answer was negative: the domain was not closed.

This article pursued the question for the third face and reached an answer of a different kind. In the sophisticated face too, the domain is not ‘discovered’, but for a different reason it is also ‘set’: not because no external institution or mechanism has fixed it, but because the system itself sets it. This ‘endogenous setting’, which was absent in the negative and positive faces, explains the resistance to correction in the sophisticated face. When the system decides what counts as a ‘gap’, the criterion of assessment becomes part of what was meant to be assessed. The trilogy reaches its conclusion at this point. In all three faces, the reference domain was not given in advance. What distinguished the three articles was whether the setting of that domain remained external or became internal.

What this article added in particular is a criterion that turns this endogenous setting from a merely philosophical claim into a documented proposition: unabsorbed resistance of reality.

This criterion needs no judgement by the researcher about the truth of a belief. It is enough to show that the system, by its own criterion or by a source independent of it, received feedback and did not change its central conclusion. The failure of absorption leads to three states (non-registration, perpetual deferral and normalisation of deviance), and the system reaches them by three modes (endogenisation of origin, endogenisation of competence and reinterpretation of mechanism). Together they show that this resistance is neither accidental nor exceptional. It is a recurring pattern, which has taken a single form in six entirely different systems, from official creationism to secular space engineering.

This finding, however, is not a claim about the content of beliefs. It is a claim about the behaviour of systems in the face of feedback. For precisely this reason, as Section 3 showed, sophisticated ignorance is not the same as compound ignorance, which is an individual and self-correcting phenomenon. Nor is it the same as Alston’s epistemic circularity or Nguyen’s echo chamber, each of which rests on a different mechanism (the structure of justification, or the structure of trust). And for precisely this reason, as Section 5 frankly acknowledged, the claim of the article remains limited to six hand-picked cases with limited institutional and geographical diversity. Its most fundamental limitation is that ‘relation to reality’ itself cannot be handed over to an automatic

measure without reproducing, in the instrument of measurement, the very resistance that is the subject of the analysis.

This article described the mechanism at the behavioural level, not at the neural, cognitive or evolutionary level. Why biological and social systems adopt such a mechanism is a separate question, and the article deliberately left it out of the discussion.

Also, although Section 2 showed that this mechanism remains active at the network level too, through at least two modes, the article did not examine as a case its systematic extension to the scale of larger networks, including recommendation algorithms or public policy. This remains a topic for separate research.

Within these limits, the claim of the article is modest but specific. Sophisticated ignorance begins where reality, however hard it presses, is no longer a match. Reality is not silent. It is heard, and then rewritten. If the trilogy on ignorance began from a question about a boundary, in this final part it arrives at the conclusion that, in the sophisticated face, nothing outside the system draws the boundary any longer. The system itself draws it anew, each time it meets reality.

References

Albert, H. (1968). *Traktat über kritische Vernunft*. Mohr Siebeck.

Al-Ghazali, Abu Hamid. *Fada'ih al-Batiniyya (al-Mustazhiri)*. 487–488 AH / 1094–1095 CE.

Al-Sibai, N. (2024, December 18). Flat Earthers Befuddled as They Visit Antarctica and Earth Appears to Be Round. *Futurism*.

Alston, W. P. (1993). *The Reliability of Sense Perception*. Cornell University Press.

Behind the Curve [Documentary]. (2018). Directed by D. J. Clark. Netflix.

Bigliardi, S. (2017). Islam and the Quest for Science and Religion Dialogue: Analogies, Overlaps, and Alternatives. *Zygon*, 52(1), 146–171.

BonJour, L. (1985). *The Structure of Empirical Knowledge*. Harvard University Press.

Boudry, M., & Braeckman, J. (2011). Immunizing Strategies and Epistemic Defense Mechanisms. *Philosophia*, 39(1), 145–161.

ChurchLeaders. (2024, December). Pastor Takes Flat-Earthers to Antarctica, Hoping To Convince Them the Earth Is Round.

Columbia Accident Investigation Board. (2003). *Report of the Columbia Accident Investigation Board* (Vol. I). NASA and Government Printing Office.

Descartes, R. (1996). *Meditations on First Philosophy* (J. Cottingham, Trans.). Cambridge University Press. (Original work published 1641)

Edis, T. (2007). *An Illusion of Harmony: Science and Religion in Islam*. Prometheus Books.

El Kassar, N. (2018). What Ignorance Really Is: Examining the Foundations of Epistemology of Ignorance. *Social Epistemology*, 32(5), 300–310.

Guessoum, N. (2011). *Islam's Quantum Question: Reconciling Muslim Tradition and Modern Science*. I.B. Tauris.

Haddon-Cave, C. (2009). *The Nimrod Review: An Independent Review into the Broader Issues Surrounding the Loss of the RAF Nimrod MR2 Aircraft XV230 in Afghanistan in 2006* (HC 1025). The Stationery Office.

Hogue, L. (2025). Mixed Results for The Final Experiment. *Substack*. https://larryhogue.substack.com/p/mixed-results-for-the-final-experiment

Kahan, D. M., Peters, E., Wittlin, M., Slovic, P., Ouellette, L. L., Braman, D., & Mandel, G. (2012). The Polarizing Impact of Science Literacy and Numeracy on Perceived Climate Change Risks. *Nature Climate Change*, 2, 732–735.

Kruger, J., & Dunning, D. (1999). Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments. *Journal of Personality and Social Psychology*, 77(6), 1121–1134.

Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.

Kunda, Z. (1990). The Case for Motivated Reasoning. *Psychological Bulletin*, 108(3), 480–498.

Lakatos, I. (1970). Falsification and the Methodology of Scientific Research Programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the Growth of Knowledge*. Cambridge University Press.

Lehrer, K. (2000). *Theory of Knowledge* (2nd ed.). Westview Press.

Le Morvan, P., & Peels, R. (2016). The Nature of Ignorance: Two Views. In R. Peels & M. Blaauw (Eds.), *The Epistemic Dimensions of Ignorance*. Cambridge University Press.

Loechelt, G. H. (2009). *A Response to the RATE Team Regarding Helium Diffusion in Zircons*. American Scientific Affiliation.

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, 27, 415–444.

Mitha, F. (2001). *Al-Ghazali and the Ismailis: A Debate on Reason and Authority in Medieval Islam*. I.B. Tauris.

Nguyen, C. T. (2020). Echo Chambers and Epistemic Bubbles. *Episteme*, 17(2), 141–161.

Nottelmann, N. (2016). The Varieties of Ignorance. In R. Peels & M. Blaauw (Eds.), *The Epistemic Dimensions of Ignorance*. Cambridge University Press.

Nyhan, B. (2021). Why the Backfire Effect Does Not Explain the Durability of Political Misperceptions. *Proceedings of the National Academy of Sciences*, 118(15), e1912440117.

Peels, R. (2023). *Ignorance: A Philosophical Study*. Oxford University Press.

Rashtriya Kamdhenu Aayog, Ministry of Fisheries, Animal Husbandry and Dairying, Government of India. (2021). *Kamdhenu Gau-Vigyan Prachar-Prasar Exam: Study Guide*.

ScienceAlert. (2024, December). Flat-Earthers Went to Antarctica to Look at the Sun. Here's What Happened. https://www.sciencealert.com/flat-earthers-went-to-antarctica-to-look-at-the-sun-heres-what-happened

ThePrint. (2021, February). Cow Urine, Milk Can Treat Skin Diseases, Psoriasis: Modi Govt's Agency in Document for Exam.

Vardiman, L., Snelling, A. A., & Chaffin, E. F. (Eds.). (2005). *Radioisotopes and the Age of the Earth, Vol. II*. Institute for Creation Research / Creation Research Society.

Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press.

Wikipedia contributors. (2025). The Final Experiment (expedition). *Wikipedia*. Retrieved 29 September 2026, from https://en.wikipedia.org/wiki/The_Final_Experiment_(expedition)

World Health Organization. (2015). *WHO Statement on the Tenth Meeting of the IHR Emergency Committee Regarding MERS*.

World Health Organization. (2019). *Middle East Respiratory Syndrome Coronavirus (MERS-CoV): Fact Sheet*.
