

From Authorship to Influence: The Trace Economy as a Public-Reliance Protocol

A framework for making public reliance traceable by integrating authorship, provenance, and influence.

Author: Stephen John Coupland

Date: 4 May 2026

Version: v1.0

Publication Type: White Paper / Conceptual Framework Paper

Status: Designed protocol; institutional infrastructure under development

Keywords:

Trace Economy; Proof of Cognitive Work; Unifaircation; public reliance; authorship; provenance; influence provenance; AI governance; AI regulation; influence disclosure; cognitive sovereignty; sponsored reasoning; longitudinal cognitive data; public accountability; social media harm reduction; digital platforms; consumer protection; conflict of interest; institutional transparency; public-sector decision-making; human sign-off; evidence provenance; authority provenance; conduct provenance; public trust.

Author Note:

This paper forms part of the author's ongoing development of the Trace Economy, Proof of Cognitive Work, and Unifaircation frameworks. It extends prior work on AI authorship, provenance, cognitive contribution, public reliance, and trace-based accountability by introducing influence as a third core pillar alongside authorship and provenance.

Protocol Status Note:

The Trace Economy is presented here as a designed public-reliance protocol rather than a fully institutionalised system. Its public anchoring mechanism can operate immediately through timestamped public tags and direct attribution. Its licensing, registry, audit, insurance, procurement, foundation, and regulatory layers remain subject to further development.

Licence / Rights Statement:

© 2026 Stephen John Coupland. All rights reserved unless otherwise stated. Use,

citation, quotation, or discussion of this work should preserve attribution to the author and the Trace Economy / Proof of Cognitive Work / Unifaircation framework.

Abstract

Artificial intelligence has intensified a broader institutional problem: people are increasingly asked to rely on outputs, claims, recommendations, rankings, decisions, public narratives, and institutional representations without sufficient visibility over who created them, where they came from, or what interests shaped them. Existing transparency mechanisms remain fragmented. Search engines may label sponsored results; films may disclose product placement; researchers may declare funding; courts may rely on chain-of-custody rules; professionals may disclose conflicts of interest; and AI governance frameworks may require certain forms of AI-content labelling. Yet these mechanisms do not provide a unified public-reliance protocol capable of integrating authorship, provenance, and influence across AI-mediated and non-AI contexts.

This paper introduces the Trace Economy as a public-reliance framework built on three core pillars: authorship, provenance, and influence. Authorship identifies the human or institutional contributor who originated, shaped, selected, validated, or stood behind a work, claim, or decision. Provenance traces the pathway through which evidence, data, materials, outputs, and decisions were produced, transmitted, modified, and relied upon. Influence reveals the forces that shaped the result, including funding, commercial incentives, behavioural optimisation, platform self-preferencing, institutional constraints, political interests, sponsored placement, affiliate relationships, and the use of longitudinal user data.

The paper argues that the relevant trigger for trace is not private thought, private creativity, or ordinary exploratory AI use. The trigger is public reliance: the moment an output, claim, recommendation, policy, report, legal filing, investigation, commercial ranking, institutional decision, or public justification is placed before others as something to trust, act upon, fund, regulate, prosecute, insure, obey, or commercially value. This distinction protects private cognition while requiring accountable traceability when reliance consequences arise.

The Trace Economy mechanism begins with a low-friction public anchor: the use of protocol hashtags, including **#TraceEconomy**, **#PoCW**, and **#Unifaircation**, together with direct tagging of the architect or relevant foundation. This trace marker does not replace evidence. Rather, it publicly anchors the claim that a work or representation has

crossed from private cognition into public reliance and that a deeper evidence trail may be audited if necessary. Supporting evidence may include drafts, source materials, funding declarations, prompt records, AI-use statements, chain-of-custody records, institutional approvals, DOI records, version histories, licensing declarations, or human sign-off records.

By extending provenance beyond content and authorship into influence, the Trace Economy provides a regulatory and civic template for AI governance, social media harm reduction, grant-funded research, political campaign transparency, professional accountability, policing, institutional decision-making, and even war-justifying public claims. Its core doctrine is simple: private cognition remains private, public reliance becomes traceable, and influence becomes visible.

Core Thesis Statement

The Trace Economy is not merely an AI-authorship protocol. It is a public-reliance protocol that integrates three missing accountability layers: authorship, provenance, and influence. Its purpose is not to monitor private cognition, but to create a visible point of accountability when ideas, outputs, claims, recommendations, decisions, or institutional representations are placed into public reliance.

Core Doctrine

Private cognition remains private.
Public reliance becomes traceable.
Influence becomes visible.

Foundational Trilogy

Authorship identifies the contributor.
Provenance traces the pathway.
Influence reveals the forces.

Section Plan

1. Introduction: The Public-Reliance Problem

Defines the core problem and explains why AI has intensified, but not created, the need for trace.

2. Fragmented Transparency: Why Existing Systems Are Incomplete

Discusses sponsored search results, product placement, conflict-of-interest declarations, research funding disclosures, chain-of-custody rules, and AI-content labelling as partial but insufficient mechanisms.

3. The Three Pillars: Authorship, Provenance, and Influence

Introduces the paper's central trilogy and explains why influence is the missing third layer.

4. AI as the Accelerant: Sponsored Reasoning and Longitudinal Cognitive Data

Shows how LLMs intensify the issue by blending advice, personalisation, persuasion, and potential commercial influence into conversational reasoning.

5. Public Reliance as the Trigger for Trace

Clarifies that trace does not apply to every private thought, prompt, draft, or creative act. It applies when reliance consequences arise.

6. The Trace Economy Mechanism: Public Anchor, Deeper Evidence Trail

Explains the three-hashtag protocol, tagging of the architect/foundation, and how public anchors relate to evidence, audit, licensing, and accountability.

7. Application Domains

Applies the framework to AI outputs, social media harm reduction, grants and research, politics, policing, professional advice, institutional decision-making, and war claims.

8. Licensing, Regulation, and Insurance Implications

Explores how the Trace Economy could operate as a voluntary protocol, compliance layer, procurement standard, insurance condition, or public-reliance licensing framework.

9. Objections and Limits

Addresses surveillance concerns, gaming risks, evidence burdens, scalability, free speech, institutional adoption, and the limits of hashtag anchoring.

10. Conclusion: From Opacity to Traceable Reliance

Restates the doctrine and positions the Trace Economy as a simple but powerful accountability primitive for the AI era and beyond.

1. Introduction: The Public-Reliance Problem

Modern societies depend on reliance. Citizens rely on governments. Patients rely on clinicians. Courts rely on evidence. Consumers rely on product claims. Investors rely on disclosures. Researchers rely on published studies. Voters rely on political

representations. Audiences rely on media narratives. Increasingly, users also rely on artificial intelligence systems to summarise, recommend, rank, draft, analyse, explain, decide, and advise.

Yet reliance often occurs without sufficient visibility over the conditions that produced the thing being relied upon. A claim may appear independent while being shaped by funding. A recommendation may appear neutral while being commercially incentivised. A report may appear objective while reflecting institutional pressure. A social media post may appear organic while being driven by sponsorship, campaign strategy, or coordinated influence. An AI-generated answer may appear helpful while being shaped by personalisation, prior user data, platform incentives, behavioural optimisation, or undisclosed ranking logic.

This is the public-reliance problem.

The problem is not limited to whether a statement is true or false. Nor is it limited to whether a work was created by a human or generated by artificial intelligence. The deeper issue is whether the public can see the authorship, provenance, and influence pathway behind outputs, claims, recommendations, decisions, and institutional representations that people are asked to trust or act upon.

Artificial intelligence has made this problem more urgent, but it did not create it. Long before large language models, societies already faced disputes over attribution, source integrity, conflicts of interest, funded research, political donations, embedded advertising, product placement, chain of custody, expert independence, professional responsibility, and institutional accountability. What AI changes is the scale, speed, opacity, and intimacy of the problem. A large language model can produce outputs that appear reasoned, personalised, and neutral, while the user may not know what data, incentives, model behaviours, commercial relationships, or prior interactions shaped the response.

This paper argues that the relevant category is not simply AI transparency. The relevant category is public reliance.

Public reliance occurs when an output, claim, recommendation, report, policy, legal filing, investigation, ranking, public statement, commercial representation, or institutional decision is placed before others as something to trust, act upon, fund, regulate, prosecute, insure, obey, publish, or commercially value. At that point, private cognition is no longer the only concern. The representation has entered a social, legal, economic, or institutional environment where others may bear consequences.

The central claim of this paper is that public reliance requires trace.

Trace does not mean total surveillance. It does not require every private thought, prompt, draft, conversation, or creative step to be exposed. Creative and intellectual

processes are often non-linear, exploratory, uncertain, and incomplete until the work has taken final form. Private cognition must remain private. But when a claim or output crosses into public reliance, a visible point of accountability should attach.

That point of accountability must answer three basic questions.

First: who contributed, originated, shaped, validated, or stood behind the work or claim? This is the question of authorship.

Second: where did the evidence, data, material, output, or decision pathway come from, and how did it move? This is the question of provenance.

Third: what interests, incentives, funding arrangements, institutional pressures, behavioural design choices, commercial relationships, political motivations, or platform dynamics shaped the result? This is the question of influence.

Existing systems address these questions only in fragments. Citation systems trace sources but not always influence. Funding declarations reveal some interests but not all pathway effects. Sponsored-content labels identify some advertising but not deeper behavioural steering. Chain-of-custody rules protect evidence integrity but usually operate within specific legal settings. AI-content labels may disclose machine generation but not whether the answer was shaped by commercial incentives, personalisation, prior user history, or institutional constraints.

The Trace Economy is proposed as a public-reliance protocol for integrating these fragments. Its foundational trilogy is simple:

Authorship identifies the contributor.

Provenance traces the pathway.

Influence reveals the forces.

Its operating doctrine is equally simple:

Private cognition remains private.

Public reliance becomes traceable.

Influence becomes visible.

This paper develops that framework. It begins by examining why existing transparency systems are incomplete. It then defines authorship, provenance, and influence as three interdependent pillars of public-reliance accountability. It considers AI as an accelerant, particularly through the emergence of sponsored reasoning and longitudinal cognitive data. It then introduces public reliance as the trigger for trace, explains the Trace Economy's low-friction public anchoring mechanism, and applies the framework across domains including AI governance, social media harm reduction, grants and research, political campaign transparency, policing, professional accountability, institutional decision-making, and war-justifying public claims.

The paper does not argue that trace determines truth. A trace marker does not prove that a claim is correct, that an output is ethical, or that an institution has acted properly. Rather, trace makes the conditions of reliance visible. It creates a public doorway into authorship, provenance, influence, evidence, accountability, and audit.

In that sense, the Trace Economy is not merely an AI-authorship protocol. It is an accountability layer for public reliance.

2. Fragmented Transparency: Why Existing Systems Are Incomplete

Transparency is not a new idea. Modern institutions already contain many partial mechanisms for revealing source, interest, process, responsibility, and influence. Search engines label sponsored results. Films and television may disclose product placement. Academic publications often require funding declarations and conflict-of-interest statements. Courts rely on rules of evidence, disclosure obligations, and chain-of-custody procedures. Professionals operate under duties of competence, independence, and accountability. Regulators require certain forms of consumer disclosure, financial disclosure, political disclosure, and data-protection notice.

These mechanisms are important. They show that societies already recognise a basic principle: where people are asked to rely on information, they should not be misled about the conditions under which that information was produced or presented.

The problem is that these systems remain fragmented. Each addresses part of the public-reliance problem, but none provides a general framework for integrating authorship, provenance, and influence across domains.

Search engines provide a useful starting point. When a search result is sponsored, it is commonly labelled so that users can distinguish paid placement from organic relevance. This does not eliminate all concerns about ranking, competition, advertising, or user profiling, but it does acknowledge a basic truth: commercial influence can affect visibility, and users are entitled to know when a result appears because money has shaped its placement.

Film and television product placement reveal another dimension. Advertising may not always appear as a separate advertisement. It can be embedded inside narrative, character identity, lifestyle, emotion, aspiration, and plot. A product can be made desirable because it is placed in the hands of a protagonist, associated with status, or woven into the symbolic world of the story. This shows that influence can operate not only through overt claims, but through context, framing, and association.

Academic and research disclosures address a related but distinct problem. Funding declarations and conflict-of-interest statements recognise that research may be shaped by financial support, institutional incentives, industry relationships, or professional commitments. A funded study is not necessarily unreliable. Funding does not automatically invalidate a result. But funding can influence the questions asked, the methods chosen, the interpretations emphasised, the uncertainties minimised, and the conclusions promoted. Disclosure does not resolve every problem, but it gives readers a basis for evaluating the conditions under which the work was produced.

Courts and policing provide yet another model: chain of custody. Evidence must be handled, stored, transferred, tested, and presented in a way that preserves its integrity. The law does not merely ask whether a piece of evidence exists. It asks where it came from, who handled it, how it was preserved, whether it was altered, whether the process can be verified, and whether the person relying on it can justify doing so. Chain of custody is therefore a domain-specific form of trace.

Professional accountability systems also recognise that responsibility must attach to human judgment. A lawyer cannot simply rely on an unverified source and avoid responsibility by pointing to the tool used. A clinician cannot avoid responsibility by saying a system suggested a diagnosis. An engineer cannot rely on automated output without professional judgment. In regulated settings, the person or institution placing material into reliance must usually be able to justify the process, evidence, and sign-off behind it.

Political and public-sector systems contain further partial disclosures. Campaign finance laws, lobbying registers, procurement rules, parliamentary records, freedom-of-information mechanisms, and administrative-law duties all attempt, in different ways, to reveal who is influencing public decisions and how authority is being exercised. Yet these systems often operate slowly, unevenly, or after the fact. They may disclose money without tracing influence, disclose process without tracing authorship, or disclose authority without revealing the incentives that shaped the decision pathway.

AI governance now adds another layer. Much of the current debate focuses on whether users should know that they are interacting with AI, whether AI-generated content should be labelled, whether certain systems are high-risk, whether datasets were lawfully obtained, and whether outputs can be explained or audited. These are necessary questions, but they are still incomplete. Knowing that content was AI-generated does not tell the user what shaped the answer. It does not reveal whether the response was influenced by sponsored placement, affiliate incentives, behavioural optimisation, platform self-preferencing, prior user history, institutional constraints, government policy logic, or longitudinal user data.

This is the central limitation of existing transparency systems: they tend to disclose fragments rather than pathways.

A citation may identify a source but not a funder.

A funding declaration may identify a sponsor but not how the sponsor shaped the research design.

A sponsored label may identify paid placement but not behavioural targeting.

A chain-of-custody record may trace evidence but not institutional pressure.

An AI-content label may identify machine generation but not commercial influence.

A conflict-of-interest statement may disclose a relationship but not its practical effect on framing, omission, or recommendation.

Public reliance requires more than fragmented disclosure. It requires a way to connect authorship, provenance, and influence into a single accountable structure.

This does not mean every claim requires the same level of trace. A casual opinion does not require the same burden as a legal filing. A restaurant recommendation does not require the same burden as a medical diagnosis. A product endorsement does not require the same burden as a police investigation. A funded policy paper does not require the same burden as a declaration of war. The burden must scale with consequence.

The principle is:

The greater the consequence of reliance, the higher the trace burden.

This principle already exists implicitly across legal and institutional life. Higher-stakes decisions usually require stronger evidence, clearer authority, more formal process, and greater accountability. The Trace Economy makes this implicit structure explicit. It proposes a general public-reliance layer that can operate across domains without collapsing them into a single rigid rule.

The key distinction is between transparency as isolated notice and trace as accountable pathway.

Transparency may tell the user one fact: this is sponsored, this was AI-generated, this was funded, this was cited, this was reviewed. Trace asks a broader question: what authorship, provenance, and influence conditions shaped the thing being relied upon?

This is why the Trace Economy must be understood as more than a disclosure label. A disclosure label is an output. Trace is a relational structure. It connects the claim, the claimant, the evidence, the process, the influence pathway, the reliance context, and the point of accountability.

Existing systems are not wrong. They are incomplete. They are fragments of a larger public-reliance architecture that has not yet been unified.

The Trace Economy provides that unifying frame. It does not replace citations, funding declarations, chain-of-custody rules, conflict-of-interest statements, sponsored-content labels, professional duties, or AI transparency laws. Rather, it gives them a common logic:

If people are being asked to rely on something, they should be able to see who stands behind it, where it came from, and what shaped it.

That is the bridge from fragmented transparency to traceable public reliance.

3. The Three Pillars: Authorship, Provenance, and Influence

Public reliance requires more than a statement, source, label, or disclaimer. It requires a structure capable of showing who contributed to the thing being relied upon, where it came from, and what shaped it. These are related but distinct questions. When they are collapsed, accountability weakens. When they are separated and traced together, public reliance becomes more intelligible.

The Trace Economy organises this structure around three pillars:

Authorship identifies the contributor.

Provenance traces the pathway.

Influence reveals the forces.

Together, these pillars form the basis of a public-reliance protocol.

3.1 Authorship: Who Contributed?

Authorship concerns the human or institutional contribution behind a work, claim, recommendation, decision, design, report, output, or public statement. It asks who originated, shaped, selected, refined, validated, signed off, or stood behind the thing being placed into reliance.

In traditional authorship contexts, the question often appears simple. A writer writes a book. An artist paints a canvas. A researcher signs a paper. A lawyer files a submission. A clinician writes a report. Yet even these cases are rarely as simple as they appear. Most public outputs are shaped by prior knowledge, collaboration, institutional process, editorial influence, funding, review, technical assistance, and social context.

Artificial intelligence makes this complexity visible. When an AI-assisted output is produced, the central question is not merely whether a human touched the process. A person may generate many prompts, make superficial edits, or select between machine outputs without supplying meaningful cognitive contribution. Conversely, a person may provide the core concept, structure, argument, constraints, source materials, governing

thesis, corrections, and final judgment, even if an AI system assists with drafting or expression.

For that reason, authorship cannot be reduced to contact. It must be linked to contribution.

A prompt log may show interaction. It does not automatically prove authorship. A final output may show production. It does not automatically reveal the human cognitive architecture behind it. The stronger authorship question is:

What identifiable human contribution gave this work, claim, recommendation, or decision its governing form, meaning, direction, or responsibility?

This matters because AI-assisted authorship may begin before the visible output exists. A prompt can be a casual request, but it can also be a recipe, score, design brief, legal theory, policy architecture, symbolic frame, or generative instruction set. Where the human supplies the governing structure and the AI executes within that structure, authorship may reside significantly in the instruction layer, not merely in the final wording.

This is especially important in theoretical, legal, policy, and institutional writing. The authorial contribution may not be limited to sentence-level expression. It may reside in the formation of the thesis, the framing of the problem, the selection of distinctions, the structure of argument, the synthesis of fields, the invention of terminology, the use of metaphor, the correction of outputs, and the acceptance of responsibility for the final text.

The Trace Economy therefore treats authorship as a claim of accountable human or institutional contribution. It does not assume that every AI-assisted output is authored merely because a human prompted it. Nor does it deny authorship merely because AI assisted in producing the output. It asks where meaningful human cognition entered the process and whether the person or institution making the claim is willing to stand behind it.

3.2 Provenance: Where Did It Come From?

Provenance concerns origin and pathway. It asks where the material, data, evidence, claim, output, recommendation, or decision came from; how it moved; what systems or actors handled it; and how it became the thing now being relied upon.

In evidence law, provenance appears through chain of custody. In academia, it appears through citation and source attribution. In art, it appears through ownership history, authenticity records, and exhibition lineage. In research, it appears through methodology, data availability, peer review, and funding declarations. In digital systems, it may appear through metadata, audit logs, dataset records, version histories, and technical documentation.

The problem is that provenance is often treated too narrowly. It is commonly understood as source-tracing, but public reliance requires pathway-tracing. A source may be identifiable, while the process that transformed the source remains opaque. A dataset may be named, while its selection, weighting, cleaning, exclusion, labelling, or embedding remains unclear. A report may cite evidence, while the pathway from evidence to conclusion remains poorly documented. A police brief may contain evidence, while the conduct, interpretation, omission, or escalation pathway remains under-examined.

AI intensifies this problem. Large language models may produce outputs from complex training histories, model architectures, retrieval systems, user prompts, system instructions, platform filters, ranking layers, and post-processing mechanisms. The user may not know which part of the system shaped the final answer. A response may appear as a single coherent statement, while its production pathway includes many hidden layers.

Provenance therefore must address more than source origin. It must include process integrity.

A useful distinction is:

Source provenance asks where the material came from.

Process provenance asks how it was transformed.

Reliance provenance asks how it became something others were asked to trust or act upon.

The Trace Economy brings these together. It treats provenance as the traceable pathway from origin to reliance. This may include source materials, data inputs, draft histories, AI-use declarations, prompts or prompt summaries, model/system identifiers, institutional approvals, funding records, review steps, version histories, DOI records, chain-of-custody logs, and human sign-off.

The point is not to expose every private step. The point is to preserve enough of the pathway that a reliance claim can be evaluated if challenged.

A trace marker does not prove the whole pathway by itself. It anchors the claim that such a pathway exists and that someone is standing behind it.

3.3 Influence: What Shaped It?

Influence is the missing third pillar.

A work may have identifiable authorship and clear provenance while still being shaped by undisclosed interests. A report may be properly cited but framed to favour a funder. A product recommendation may accurately describe options but rank them according to affiliate incentives. A political claim may be authorised and sourced but shaped by

donors, lobbyists, campaign strategy, or polling data. A social media post may appear organic while being sponsored or coordinated. An AI response may be factually useful while being shaped by personalisation, behavioural optimisation, platform self-preferencing, or commercial arrangements.

Influence asks:

What forces shaped the output, claim, recommendation, ranking, decision, or public representation?

Those forces may include funding, commercial incentives, sponsorship, affiliate relationships, institutional constraints, political interests, platform incentives, behavioural optimisation, recommendation logic, prior user data, emotional targeting, government policy, internal risk settings, professional duties, or organisational self-interest.

Influence is not automatically improper. All communication is shaped by something. Values, expertise, institutional roles, professional obligations, lived experience, funding, ideology, incentives, and context all influence what people say and how they say it. The problem is not influence as such. The problem is undisclosed material influence where others are asked to rely on the result.

This distinction matters. The Trace Economy does not assume that funded work is corrupt, that sponsored material is false, that institutional claims are invalid, or that AI-personalised outputs are necessarily harmful. It asks whether material influence factors are visible enough for reliance to be properly calibrated.

The simplest principle is:

Funding is not the problem. Influence is not the problem. Undisclosed material influence is the problem.

This is where AI creates a new regulatory frontier. Traditional advertising can be labelled. Sponsored search results can be marked. Film product placement can be disclosed or scrutinised. But LLM-mediated influence may be blended into personalised reasoning. A user may not see an advertisement, a ranked list, or an obvious product placement. Instead, the user may receive an answer that appears conversational, neutral, and tailored to their situation.

This gives rise to what may be called **sponsored reasoning**: commercially or institutionally influenced reasoning embedded inside AI-generated advice, rankings, summaries, recommendations, or decision-support outputs.

Sponsored reasoning is more subtle than sponsored content. It may not appear as a banner, link, or product placement. It may appear as the logic of the answer itself. The risk is that a user may believe they are receiving independent assistance when the

response has been shaped by hidden commercial, behavioural, or institutional incentives.

Influence provenance is therefore necessary. It asks whether the answer, recommendation, claim, or decision pathway was shaped by factors such as:

personalisation,
prior user history,
sponsored placement,
affiliate incentive,
platform self-preferencing,
behavioural optimisation,
institutional policy,
government eligibility logic,
political targeting,
funding dependency,
or longitudinal cognitive data.

In this sense, influence disclosure is the next step in transparency. It does not merely ask whether content was AI-generated. It asks what shaped the content, ranking, recommendation, or reasoning pathway.

3.4 Why the Pillars Must Operate Together

Each pillar is necessary, but none is sufficient alone.

Authorship without provenance can become assertion. A person may claim responsibility for a work, but without a traceable pathway, the claim may be difficult to verify.

Provenance without authorship can become mechanical. A process may be documented, but without accountable human or institutional sign-off, responsibility may be displaced into systems, workflows, or tools.

Authorship and provenance without influence can become misleading. A claim may be properly attributed and sourced, yet still shaped by hidden incentives that materially affect how others should evaluate it.

Influence without authorship and provenance can become accusation. It is not enough to say that interests may have shaped a result. A credible framework must identify who is making the claim, where the claim came from, and what evidence supports the influence concern.

The Trace Economy integrates all three.

It asks:

Who stands behind this?

Where did it come from?

What shaped it?

Why is the public being asked to rely on it?

What evidence trail exists if the claim is challenged?

That integrated structure is what distinguishes trace from ordinary transparency.

Transparency may disclose isolated facts. Trace connects those facts into an accountable pathway.

3.5 Public Reliance and the Scale of Trace

Not every context requires the same burden. A casual opinion, private draft, artistic experiment, or exploratory conversation should not be treated like a legal filing, medical recommendation, police investigation, public grant report, political campaign, or war-justifying claim.

The burden must scale with consequence.

The greater the potential impact on rights, liberty, health, money, reputation, democratic choice, public resources, institutional trust, or human life, the stronger the trace obligation should be.

This creates a public-reliance ladder. Low-stakes contexts may require little more than ordinary attribution or no trace at all. Commercial recommendations may require disclosure of paid influence or affiliate incentives. Professional advice may require evidence, competence, and human sign-off. Institutional decisions may require process records and authority provenance. Police investigations may require evidence provenance, authority provenance, and conduct provenance. War claims may require the highest level of trace: evidence categories, legal authority, influence disclosure, dissent records, and retrospective audit.

The Trace Economy does not collapse these contexts into one rule. It supplies a common logic:

The greater the consequence of reliance, the higher the trace burden.

This principle allows trace to remain proportionate while still being structurally consistent across domains.

3.6 The Trilogy as a Public-Reliance Framework

The three pillars can be expressed simply:

Authorship identifies the contributor.

Provenance traces the pathway.

Influence reveals the forces.

This trilogy reframes the governance problem. The question is no longer only whether content is true, whether AI was used, or whether a source can be cited. The broader question is whether a person or institution can make a public-reliance claim with sufficient visibility over contribution, pathway, and influence.

This is the point at which the Trace Economy becomes more than an authorship protocol, more than a provenance system, and more than an AI-governance tool. It becomes a general framework for accountable public reliance.

The next section turns to artificial intelligence as the accelerant that makes this framework urgent.

4. AI as the Accelerant: Sponsored Reasoning and Longitudinal Cognitive Data

Artificial intelligence did not create the public-reliance problem. It accelerated it.

Before large language models, societies already had to manage sponsored search results, product placement, political donations, funded research, institutional conflicts, professional responsibility, chain of custody, and public-sector accountability. These were all partial responses to the same underlying problem: people are often asked to rely on claims, recommendations, rankings, reports, decisions, or narratives that may be shaped by interests they cannot see.

Large language models intensify this problem because they combine several functions that were previously more separate. They can search, summarise, rank, draft, advise, recommend, persuade, analyse, coach, classify, translate, and simulate expertise inside a single conversational interface. The user does not merely receive a list of links or a visible advertisement. The user receives a response that may appear reasoned, personalised, neutral, and responsive to their circumstances.

That creates a new kind of public-reliance risk.

The central issue is not simply that AI may generate inaccurate content. Inaccuracy is one problem. Hallucination is one problem. Copyright infringement is one problem. Bias is one problem. But the deeper governance issue is that AI can produce outputs whose shaping forces are invisible to the user.

A response may be shaped by model training, system instructions, safety policies, retrieval sources, platform ranking systems, commercial relationships, prior user history, behavioural optimisation, institutional constraints, or government policy logic. Yet the final answer may appear as a single coherent statement. The user may not know whether the response is evidence-based, personalised, commercially influenced, institutionally constrained, or optimised for engagement.

That is why AI transparency cannot stop at content labelling. Knowing that an answer came from AI does not tell the user what shaped the answer.

4.1 From Sponsored Links to Sponsored Reasoning

Search engines provide the clearest historical analogy. When a search engine displays sponsored results, the user is given at least some notice that placement may be commercial rather than purely organic. That disclosure is not perfect, but it recognises an important distinction between relevance and paid visibility.

Film and television provide a second analogy. Product placement shows how advertising can be embedded inside narrative. A brand may be placed inside a story, associated with a character, or connected to aspiration, identity, humour, romance, status, or emotion. The advertisement is not presented as an interruption. It is blended into the expressive environment.

LLMs create a third and more subtle category.

They may allow commercial influence to be embedded not merely inside ranking or narrative, but inside reasoning.

A search result may say, in effect:

This link appears prominently because it is sponsored.

A film may imply:

This product belongs in the world of the story.

An LLM may say:

Given your goals, circumstances, prior concerns, budget, values, and decision criteria, this is the best option for you.

That is a different level of influence. It can feel like advice rather than advertising. It can appear as judgment rather than placement. It can be personalised to the user's circumstances, history, vulnerabilities, ambitions, and inferred reasoning style.

This is what this paper calls **sponsored reasoning**.

Sponsored reasoning occurs when commercial, political, behavioural, platform, institutional, or other material influence is embedded inside AI-generated advice, recommendations, rankings, summaries, explanations, or decision-support outputs.

The issue is not that commercial influence is always illegitimate. Advertising, sponsorship, affiliate relationships, and institutional recommendations may all have lawful and legitimate roles. The problem arises when the influence is material and undisclosed, especially where the output is framed as neutral assistance.

If search engines must label sponsored links, and film product placement is recognised as embedded advertising, then LLMs should disclose commercially influenced answers. The standard should arguably be higher, not lower, because an LLM is not merely presenting a placement. It may be shaping personalised reasoning.

4.2 The Conversational Authority Problem

LLMs produce outputs in a conversational form. This matters because conversation carries social authority. A user may experience the system not as a billboard, search page, or static report, but as an assistant, adviser, tutor, analyst, coach, confidant, or decision-support partner.

That conversational form can increase trust. It can also obscure influence.

A ranked list invites comparison. A sponsored label signals caution. A footnote identifies a source. A product placement is often visible as a brand inside a scene. But an LLM answer may integrate facts, assumptions, recommendations, probabilities, warnings, and suggested actions into one fluent response. The influence pathway can disappear inside the prose.

This is especially important in high-reliance domains: law, medicine, finance, education, insurance, employment, government services, policing, public health, and political communication. In these contexts, the user may not simply be browsing. They may be making decisions that affect rights, money, liberty, health, reputation, livelihood, or public obligations.

The conversational authority problem is therefore not only technical. It is institutional. Once an AI system speaks in the form of advice, users need to know whether it is operating as:

- an assistant,
- an adviser,
- a recommender,
- an advertiser,
- a platform representative,
- a government interface,
- a decision-support tool,
- a behavioural steering mechanism,
- or a commercially incentivised channel.

A user should not be left to infer this from the tone of the answer. The role of the system should be disclosed where reliance consequences may arise.

4.3 Longitudinal Cognitive Data

Traditional digital platforms often infer preferences. A music platform may infer what songs a person likes. A video platform may infer what shows a person watches. A search engine may infer interests from queries. A social media platform may infer attention patterns from likes, shares, comments, and dwell time.

LLMs can go further.

Over time, an LLM may observe how a person thinks. It may receive repeated examples of a user's reasoning style, problem-framing habits, doubts, goals, ambitions, vulnerabilities, creative process, professional strategy, political concerns, legal anxieties, financial pressures, emotional states, persuasion pathways, and evolving worldview.

This creates what this paper calls **longitudinal cognitive data**.

Longitudinal cognitive data is not merely preference data. It is the accumulated record, or inferable pattern, of how a person frames problems, develops ideas, tests arguments, responds to uncertainty, revises positions, selects options, expresses fears, articulates desires, and moves toward decisions over time.

This type of data is extremely sensitive because it may reveal not only what a person likes, but how that person may be moved.

That has implications for advertising, political persuasion, financial exploitation, platform dependency, product recommendation, employer monitoring, education, insurance, professional advice, and public-sector service delivery. It also has implications for authorship and intellectual property, because the same longitudinal record may reveal the development of an original body of thought.

This creates a dual-use problem.

On one hand, longitudinal AI interaction can help evidence authorship. A body of AI-assisted work may show continuity of cognition: recurring concepts, distinctions, metaphors, corrections, frameworks, methods, and public commitments. It may demonstrate that a work belongs to a traceable trajectory of human thought.

On the other hand, if that longitudinal record is held inside a proprietary AI platform, it may become a site of extraction, dependency, and manipulation. The platform may hold the richest record of a user's cognitive development, while the user may not control how that record is stored, used, monetised, or made available for authorship proof.

This is why cognitive sovereignty becomes central.

The issue is not only data privacy. It is control over the developmental record of human thought produced through AI interaction.

4.4 Why Ordinary Personalisation Disclosure Is Insufficient

Many digital systems already use personalisation. A platform may say that content is tailored to the user's interests. That may be adequate where the stakes are low. But LLM-mediated personalisation can affect the framing of advice itself.

A personalised AI response may choose which facts to emphasise, which risks to minimise, which options to rank, which emotional frame to use, which uncertainty to disclose, which action to recommend, and which product, service, institution, or pathway to make appear most reasonable.

If that shaping is based on prior user history or inferred cognitive patterns, then ordinary "personalisation" language may be too vague.

The relevant question is not merely:

Was this answer personalised?

The stronger question is:

Was this answer shaped by prior user data, inferred traits, behavioural optimisation, commercial incentives, institutional constraints, or platform interests in a way that materially affects the user's decision?

That is why influence disclosure must be more granular than a generic personalisation notice.

A meaningful LLM influence disclosure might indicate whether the response used:

prior conversation history,
stored memory or profile information,
sponsored placement,
affiliate incentives,
commercial ranking,
platform self-preferencing,
behavioural optimisation,
sensitive inferred traits,
institutional policy constraints,
government eligibility logic,
or human sign-off.

This does not require exposing the full internal reasoning of the model. It requires disclosing material influence factors that shaped the output where reliance consequences may arise.

4.5 Proprietary Platforms and Cognitive Asymmetry

The public-reliance problem becomes sharper when the AI system is proprietary. The provider may control the model, interface, data retention settings, ranking layers,

memory features, system instructions, product integrations, commercial arrangements, safety filters, and disclosure design. The user may experience the output as a direct answer, while the provider controls the hidden conditions under which that answer is produced.

This creates cognitive asymmetry.

The platform may know how the user thinks, what the user is trying to build, which arguments persuade the user, which insecurities or ambitions recur, and which recommendations the user is likely to accept. The user, by contrast, may not know how the platform is using that knowledge, whether commercial incentives are present, whether recommendations are ranked neutrally, or whether prior conversations are shaping current outputs.

This asymmetry is especially serious where the platform profits from steering user behaviour. If the provider or deployer receives kickbacks, affiliate revenue, sponsored placement fees, subscription conversions, marketplace commissions, advertising revenue, or ecosystem benefits, then the influence pathway should not be hidden.

Responsibility should follow control and benefit.

If a platform controls the influence pathway and benefits from it, it should not be able to shift responsibility entirely onto the end user. End-user responsibility depends on informed reliance. Where material influence is concealed, reliance is not fully informed.

This principle is central:

End users cannot be fully responsible for reliance where the influence pathway was concealed.

4.6 Influence Disclosure as Conflict-of-Interest Disclosure for AI

Influence disclosure should operate like conflict-of-interest disclosure for machine-mediated persuasion.

A professional conflict-of-interest disclosure does not necessarily prove misconduct. It reveals a relationship that may affect judgment. A funding declaration does not prove that research is false. It reveals a condition that may shape interpretation. A sponsored search label does not prove that a result is bad. It reveals that paid placement affected visibility.

In the same way, LLM influence disclosure would not automatically invalidate an output. It would reveal material shaping factors so that users, regulators, auditors, courts, insurers, and institutions can calibrate reliance.

The disclosure should answer practical questions:

Was the response based on evidence, general model output, retrieval sources, or user-provided material?

Was it personalised using prior user data?

Were sponsored, affiliate, or commercial incentives involved?

Was ranking or recommendation influenced by platform self-preferencing?

Was behavioural optimisation used?

Were sensitive inferred traits used?

Was the output constrained by institutional policy, government logic, or employer settings?

Was a human required to sign off before public reliance?

This does not require complete model transparency. It requires reliance-relevant transparency.

The question is not:

Can the user inspect every parameter of the model?

The question is:

Can the user see the material forces that shaped the answer they are being asked to rely on?

4.7 AI as a Public-Reliance Multiplier

AI multiplies public reliance in three ways.

First, it multiplies production. More outputs can be generated faster, at lower cost, across more domains.

Second, it multiplies authority. Outputs may appear fluent, reasoned, personalised, and expert-like, even where the underlying process is opaque.

Third, it multiplies influence. Responses can be adapted to user context, prior interaction, inferred cognition, commercial incentives, and institutional objectives.

This makes trace more necessary, not less.

The solution is not to ban AI assistance, nor to require every private prompt to be disclosed. That would be unworkable and would risk turning creativity, inquiry, and personal reflection into monitored activity. The solution is to identify the moment when AI-mediated output enters public reliance and require trace at that point.

In that sense, AI is the accelerant that reveals the broader structure.

The same questions that apply to AI also apply to funded research, political campaigns, social media recommendations, professional advice, police investigations, public-

sector decisions, and war claims. AI makes the missing layer easier to see because it compresses authorship, provenance, and influence into a single conversational output.

4.8 The Need for Influence Provenance

Content provenance asks where information came from. Authorship provenance asks who contributed human cognition. Influence provenance asks what shaped the output and who benefits from its acceptance.

AI governance must include all three.

A system that discloses AI generation but hides commercial influence remains incomplete. A system that cites sources but hides ranking incentives remains incomplete. A system that acknowledges personalisation but hides behavioural optimisation remains incomplete. A system that provides useful recommendations but hides platform self-preferencing remains incomplete.

Influence provenance fills that gap.

It recognises that users should not only know whether an output was AI-generated. They should know whether the reasoning pathway was materially shaped by interests, incentives, or prior user data that affect reliance.

This is the bridge from AI transparency to public-reliance traceability.

The next section develops the trigger for trace: not private thought, not ordinary prompting, not every creative step, but public reliance.

5. Public Reliance as the Trigger for Trace

The Trace Economy does not begin from the premise that every thought, draft, prompt, conversation, or creative act should be monitored. Such an approach would be unworkable, invasive, and hostile to the nature of intellectual and creative development. Human cognition is often exploratory before it is declarative. It moves through uncertainty, intuition, fragments, revisions, failed attempts, contradictions, experiments, and later recognition. The full picture is often only visible after the work has taken form.

For that reason, the trigger for trace cannot be private cognition.

The trigger must be public reliance.

Public reliance occurs when a work, claim, recommendation, ranking, report, policy, investigation, legal filing, professional opinion, public statement, commercial representation, institutional decision, or AI-mediated output is placed before others as

something to trust, act upon, fund, regulate, prosecute, insure, publish, obey, or commercially value.

This is the moment at which private development becomes socially consequential. It is the point where authorship, provenance, influence, and responsibility begin to matter to others.

5.1 Private Cognition Must Remain Private

A public-reliance framework must preserve the distinction between thinking and claiming.

People must be free to think, draft, test, prompt, revise, speculate, explore, create, and discard without turning every intellectual movement into a public record. This is especially important in AI-assisted creation, where users may generate multiple drafts, test weak ideas, explore alternative arguments, ask naïve questions, or use AI systems as private thinking aids.

If every prompt were treated as a trace event, the system would collapse into procedural surveillance. It would also create incentives for performative authorship: users could manufacture long prompt logs to simulate contribution, while genuinely creative processes might be penalised for being messy, intuitive, or non-linear.

The Trace Economy rejects that model.

Private cognition remains private. The private process may generate evidence that can later support an authorship or reliance claim, but it does not itself need to be publicly trace-logged. The public trace begins when a person or institution asks others to rely on the result.

This distinction protects creative freedom while preserving accountability where consequences arise.

5.2 The Transition Point: From Thinking to Reliance

The key moment is the transition from private cognition to public reliance.

A person may privately develop an idea with notebooks, conversations, AI tools, research materials, collaborators, drafts, prompts, or informal feedback. At that stage, the work remains in formation. It is not yet a public claim. It is not yet being relied upon by others.

But once the person publishes it, sells it, licenses it, submits it, files it, markets it, recommends it, uses it to guide others, or invokes it as evidence or authority, the work enters a different category.

It becomes a reliance object.

A reliance object is any output or representation that others are invited, expected, encouraged, or required to trust or act upon.

Examples include:

- a published AI-assisted paper,
- a legal submission,
- a medical recommendation,
- a financial product ranking,
- a funded research report,
- a government eligibility decision,
- a police brief of evidence,
- a social media campaign,
- a product endorsement,
- a political advertisement,
- a grant outcome report,
- an institutional whitepaper,
- a public health statement,
- a procurement recommendation,
- or a war-justifying public claim.

These are not merely expressions. They are claims placed into social or institutional circulation with potential consequences.

The Trace Economy attaches at that transition point.

5.3 Why Public Reliance Is the Correct Trigger

Public reliance is the correct trigger because it avoids two extremes.

The first extreme is opacity. Under this model, outputs can influence people, institutions, markets, courts, governments, and public opinion without any visible account of authorship, provenance, or influence. This leaves users, citizens, regulators, insurers, courts, and counterparties unable to assess the conditions under which reliance is being requested.

The second extreme is total disclosure. Under this model, every prompt, draft, private thought, research note, and exploratory step would be captured or exposed. This would chill creativity, undermine privacy, and confuse process with accountability.

Public reliance offers a middle path.

It does not require disclosure of every private step. It requires an accountable trace once the result is placed before others as something consequential.

This mirrors existing institutional logic. Courts do not monitor every private thought of a lawyer, but once a submission is filed, the lawyer is responsible for its contents.

Medical regulators do not inspect every private reflection of a clinician, but once a diagnosis or treatment recommendation is made, professional accountability attaches. Researchers are not required to disclose every intellectual influence, but once a paper is published, methods, funding, authorship, and conflicts may become relevant. Governments may deliberate internally, but once they ask the public to rely on a policy or justification, accountability obligations arise.

The same principle should apply to AI-mediated and non-AI public outputs.

5.4 Reliance Consequences and Trace Burden

Not all reliance events carry the same consequence. The trace burden should be proportionate.

A casual creative post does not require the same trace burden as a court filing. A light product recommendation does not require the same burden as financial advice. A low-stakes AI summary does not require the same burden as an automated government benefits decision. A grant-funded community report does not require the same burden as a declaration of war.

The Trace Economy therefore operates on a scaling principle:

The greater the consequence of reliance, the higher the trace burden.

This allows trace to remain practical rather than rigid.

Low-consequence reliance may require only basic attribution or no formal trace. Moderate-consequence reliance may require authorship and influence disclosure. Commercial reliance may require disclosure of sponsorship, affiliate incentives, platform interests, or paid placement. Professional reliance may require evidence, competence, process records, and human sign-off. State coercion, policing, prosecution, and war claims require the highest level of trace because they may affect liberty, life, national security, public resources, and democratic legitimacy.

This scaling principle prevents the Trace Economy from overreaching while preserving its core logic across domains.

5.5 Public Reliance in AI-Assisted Creation

AI-assisted authorship makes the public-reliance trigger especially important.

During private creation, the author may use AI in many ways: brainstorming, drafting, summarising, testing arguments, structuring sections, generating alternatives, editing prose, translating, critiquing, or refining ideas. The process may be iterative and unpredictable. A prompt may function as a casual request, a detailed recipe, a generative instruction set, or a temporary experiment that is later abandoned.

It would be unreasonable to require every step to be publicly traced.

But once the AI-assisted work is published, sold, submitted, licensed, cited, relied upon, or presented as evidence of human authorship, the relevant authorship claim should become traceable.

That does not mean publishing every prompt. It means declaring the nature of the human contribution and the role of AI assistance in a proportionate way.

A suitable AI-use trace statement might identify:

the originating human concept,
the human-supplied structure or thesis,
the role of AI in drafting or refinement,
the extent of human revision,
the source materials supplied,
the final human editorial judgment,
and the person or institution accepting responsibility.

The key question is not whether AI was used. The key question is whether the human contribution is identifiable, proportionate, and accountable at the point of public reliance.

5.6 Public Reliance in Institutional Contexts

Institutions often create reliance objects: reports, policies, investigations, procurement recommendations, risk assessments, public statements, research outputs, product rankings, eligibility decisions, and regulatory findings. These outputs may be shaped by multiple people, systems, funding sources, incentives, and internal pressures.

Public reliance becomes especially important here because institutions can diffuse responsibility. A claim may be attributed to “the department,” “the agency,” “the model,” “the report,” “the algorithm,” “the committee,” or “the process,” while no clear accountable pathway is visible.

Trace resists this diffusion.

It asks:

Who authorised the output?

Who contributed to it?

What evidence or data supported it?

What process produced it?

What systems were used?

What institutional incentives shaped it?

What conflicts or funding relationships existed?

What dissent or uncertainty was recorded?

Who signed off?

Who is accountable if others relied on it?

This does not require every internal deliberation to be public. It requires a traceable pathway sufficient to support the reliance claim.

5.7 Public Reliance and State Power

The public-reliance trigger is strongest where state power is exercised.

Police investigations, prosecutions, administrative decisions, public health orders, welfare determinations, immigration decisions, surveillance authorisations, military actions, and declarations of war all involve claims or decisions that citizens may be required to accept, obey, or contest under unequal power conditions.

In these contexts, trace is not merely a matter of transparency. It is a safeguard against abuse, error, omission, and institutional evasion.

For policing, trace does not only concern evidence. It also concerns authorised conduct. The question is not merely what evidence was collected, but what authorised officers did, failed to do, relied upon, ignored, escalated, omitted, disclosed, or departed from, and whether those actions were consistent with training, lawful powers, policies, and professional standards.

For war claims, trace must account for evidence, authority, influence, legal justification, dissent, uncertainty, and retrospective audit. Classified information may remain protected where necessary, but public justification cannot rest on rhetoric alone.

The greater the coercive power of the state, the stronger the trace burden should be.

5.8 Reliance Without Censorship

A public-reliance protocol is not a censorship framework.

The Trace Economy does not seek to prevent speech, suppress opinion, or determine truth by decree. It does not say that a person cannot make a claim, publish an idea, advocate for a policy, recommend a product, or challenge an institution.

Instead, it asks whether the conditions of reliance are visible when others are asked to trust or act upon the claim.

This distinction is essential.

Trace does not censor speech. It contextualises reliance.

A person may still speak. An organisation may still advocate. A government may still propose policy. A researcher may still publish funded work. A platform may still recommend products. An AI system may still assist users. But where reliance

consequences arise, the relevant authorship, provenance, influence, and accountability conditions should be traceable.

This makes trace compatible with open discourse. It does not suppress claims; it strengthens the public's ability to evaluate them.

5.9 The Public Anchor

The Trace Economy's practical mechanism begins with a public anchor.

The public anchor does not contain the whole evidence trail. It marks the reliance event.

In the Trace Economy, this can begin through the use of protocol hashtags — **#TraceEconomy**, **#PoCW**, and **#Unifaircation** — together with direct tagging of the architect or relevant foundation. This public act signals that the work, claim, recommendation, or output has been placed into public reliance under the trace protocol.

The public anchor says:

this claim has entered public reliance,
a human or institution is standing behind it,
a provenance claim is being made,
and a deeper evidence trail may be audited if necessary.

The hashtags do not replace evidence. They anchor the claim.

This distinction is critical. The trace marker is not the entire proof. It is the public doorway into proof. The supporting record may include drafts, prompt records, funding disclosures, source materials, version histories, institutional approvals, chain-of-custody logs, conflict-of-interest declarations, human sign-off records, DOI anchors, or audit trails.

The public anchor creates visibility without requiring unnecessary exposure of private cognition.

5.10 Public Reliance as the Organising Principle

Public reliance is the organising principle that allows the Trace Economy to apply across AI and non-AI contexts without overreach.

It explains why private prompting should not be surveilled but published AI-assisted work may require an authorship statement.

It explains why personal opinion may remain informal but sponsored recommendation should be disclosed.

It explains why funded research is not invalid but should carry interest transparency.

It explains why police investigations may remain operationally confidential but coercive state action must be traceable.

It explains why classified intelligence may remain protected but public war justification requires accountable trace.

It explains why AI governance must go beyond content labelling and include influence provenance.

The framework can therefore be summarised as follows:

Private cognition remains private.

Public reliance becomes traceable.

Influence becomes visible.

The next section explains how the Trace Economy mechanism operationalises this principle through a low-friction public trace anchor supported by deeper evidence, audit, licensing, and accountability layers.

6. The Trace Economy Mechanism: Public Anchor, Deeper Evidence Trail

The Trace Economy is designed around a simple distinction: the public does not need access to every private cognitive step, but it does need a visible point of accountability when a claim, output, recommendation, decision, or work enters public reliance.

This is why the mechanism begins with a public anchor.

The public anchor is not the whole proof. It is the visible marker that a reliance claim has been made and that a deeper evidence trail exists, or should exist, behind it. It marks the transition from private cognition to public reliance.

In its simplest form, the Trace Economy public anchor consists of three protocol hashtags:

#TraceEconomy

#PoCW

#Unifaircation

together with direct tagging of the architect or relevant foundation.

This combination creates a public, timestamped, searchable signal. It does not require a new platform to begin. It can operate across existing public platforms that support timestamps, hashtags, tagging, comments, posts, publication links, and public archives. Over time, it can be strengthened through DOI anchoring, institutional registers, licensing records, audit systems, and formal governance infrastructure.

The simplicity of the mechanism is not a weakness. It is the reason it can operate as an adoption primitive.

A signature is simple.

A timestamp is simple.

A citation is simple.

A disclosure label is simple.

A DOI is simple.

The power of these mechanisms does not lie in complexity. It lies in the fact that they create a recognisable point of accountability. The Trace Economy public anchor serves the same function for public-reliance claims.

6.1 The Public Anchor Does Not Replace Evidence

A common misunderstanding would be to treat the trace tag as if it proves the entire claim by itself. It does not.

The hashtags do not replace evidence. They anchor the claim.

A trace marker says that a person or institution is placing something into public reliance under the Trace Economy protocol. It identifies the public moment at which the claim, work, recommendation, decision, or output was presented as something others may rely upon. It creates a reference point around which evidence can later be organised, audited, disputed, strengthened, or challenged.

The deeper evidence trail may include:

- drafts,
- notes,
- source materials,
- version histories,
- AI-use declarations,
- prompt summaries or material prompt records,
- funding declarations,
- conflict-of-interest statements,
- dataset records,
- chain-of-custody records,
- institutional approvals,
- human sign-off records,
- methodological notes,
- audit logs,
- DOI records,
- licensing declarations,
- public comments,

collaboration records,
and retrospective review materials.

The trace tag is therefore best understood as a doorway into evidence, not a substitute for evidence.

This distinction preserves credibility. The Trace Economy does not claim that a hashtag can prove truth, authorship, influence, or provenance by itself. Rather, it claims that a public trace anchor can mark the reliance event and connect that event to a broader evidentiary pathway.

6.2 Why Hashtags Matter

Hashtags are sometimes dismissed as informal or lightweight. That criticism misses their structural function.

A hashtag is a public indexing device. It allows dispersed content to be grouped, searched, followed, archived, and cross-referenced across platforms. When combined with timestamping and direct tagging, hashtags can create a visible public signal that a particular claim or output belongs to a defined protocol.

The Trace Economy uses hashtags not as decoration, but as a public indexing layer.

The three tags serve different functions.

#TraceEconomy identifies the broader system of public-reliance traceability, attribution, influence disclosure, and accountable value recognition.

#PoCW identifies the Proof of Cognitive Work layer: the claim that human cognitive contribution, direction, judgment, or sign-off is being asserted and may be evidenced.

#Unifaircation identifies the founding protocol architecture and links the claim to the wider governance and licensing framework.

Direct tagging of the architect or relevant foundation creates an additional notice function. It signals that the claim is not merely descriptive but is being made under a recognised protocol pathway. In later institutional versions, this tag may be directed to the Trace Commons Foundation or another designated governance body.

Together, these markers create a low-friction, public, searchable trace event.

6.3 Public Trace as a Transition Marker

The core function of the public anchor is to mark transition.

Before trace:

private thought,
private drafting,

private prompting,
private research,
private experimentation,
private revision,
private collaboration,
and private uncertainty remain protected.

After trace:

the work, claim, recommendation, decision, or output has entered public reliance.

This transition matters because accountability attaches differently once others are invited to rely.

For example, an author may privately use AI to explore ideas. That does not require public trace. But once the author publishes a paper, sells a book, files a legal submission, markets a product, or claims professional authority, the relevant authorship and AI-use claims become reliance-relevant.

A company may internally test AI-driven product recommendations. That does not necessarily require public trace. But once those recommendations are presented to customers, especially where sponsorship, affiliate incentives, platform ranking, or behavioural optimisation may shape the result, influence disclosure becomes relevant.

A police investigator may conduct confidential inquiries. Operational secrecy may be necessary. But once investigative material is used to justify arrest, charge, search, prosecution, public accusation, or court reliance, evidence provenance, authority provenance, and conduct provenance become necessary.

A government may deliberate internally about national security. But once it asks citizens, soldiers, allies, institutions, and international bodies to rely on a war-justifying claim, the public justification must carry a heightened trace burden.

The public anchor marks the moment when the reliance burden begins.

6.4 The Deeper Evidence Trail

The deeper evidence trail is context-specific. It does not look the same in every domain.

In AI-assisted authorship, the evidence trail may include the originating concept, outline, human instructions, material prompts, draft evolution, AI-use declaration, revisions, source materials, final sign-off, and publication record.

In commercial recommendations, it may include ranking criteria, affiliate relationships, sponsorship arrangements, product data, personalisation signals, user-history usage, and disclosure records.

In research, it may include funding sources, methodology, data provenance, conflicts of interest, peer review, analysis scripts, limitations, and publication history.

In policing, it may include chain of custody, officer conduct logs, legal authority, training requirements, deviations from procedure, supervisory approvals, forensic methods, AI-system use, exculpatory material, and disclosure records.

In public-sector decision-making, it may include policy authority, eligibility logic, human review, algorithmic tools, administrative reasons, appeal pathways, and audit records.

In political campaigns, it may include donor interests, advertising spend, targeting criteria, consultant relationships, message-testing, platform placement, and influence networks.

In war-justifying claims, it may include evidence categories, legal authority, decision-making process, intelligence confidence level, dissenting assessments, conflict-of-interest disclosures, parliamentary or congressional review, and retrospective audit mechanisms.

The Trace Economy does not require one universal evidence file. It requires that the evidence trail be proportionate to the reliance context.

The public anchor is common. The supporting record varies.

6.5 Trace Claims and Claim Types

To avoid becoming a vague label, the Trace Economy should distinguish between different trace claims.

Not every trace marker means the same thing. A person may be claiming authorship, influence disclosure, funding transparency, public sign-off, institutional responsibility, or reliance compliance. These should be identified where possible.

Examples of trace claim types include:

Authorship Trace

The claimant asserts human or institutional contribution to a work, idea, design, report, creative output, legal argument, policy, or framework.

AI-Use Trace

The claimant discloses that AI assisted in drafting, analysis, summarisation, generation, ranking, recommendation, editing, or decision support.

Influence Trace

The claimant discloses material shaping factors such as sponsorship, affiliate incentives, platform ranking, behavioural optimisation, institutional constraints, political funding, or prior user data.

Funding Trace

The claimant discloses financial support, grants, donor relationships, industry sponsorship, political backing, or other funding structures relevant to reliance.

Authority Trace

The claimant identifies the legal, institutional, professional, or organisational authority under which a decision or public representation was made.

Conduct Trace

The claimant records what authorised persons did, failed to do, escalated, omitted, or departed from in relation to training, procedure, policy, or lawful authority.

Reliance Trace

The claimant declares that an output, decision, report, claim, or recommendation is being placed into public reliance and that supporting evidence exists or can be audited.

These categories can overlap. For example, an AI-assisted legal report may require authorship trace, AI-use trace, source provenance, human sign-off, and influence disclosure. A government AI eligibility tool may require authority trace, process provenance, influence trace, auditability, and human review. A funded research report may require authorship trace, funding trace, methodology provenance, and influence disclosure.

Claim types prevent trace from becoming an empty badge. They clarify what is being asserted.

6.6 Trace as Anti-Laundering Infrastructure

The Trace Economy is not designed to make weak claims look strong. It is designed to prevent procedural laundering.

A poorly designed AI-provenance system could become a tick-and-flick exercise. Users might generate prompt logs, make superficial edits, select outputs, and then claim authorship without meaningful human contribution. Institutions might attach generic AI disclosures while hiding the commercial, behavioural, or political incentives shaping outputs. Platforms might label content as AI-generated while leaving sponsored reasoning undisclosed. Organisations might issue conflict statements while omitting practical influence pathways.

Trace must avoid this.

The public anchor must be linked to proportionate claim types and supporting evidence. A person cannot merely say “trace” and thereby prove authorship, neutrality, legitimacy, or accountability. The trace claim must be reviewable.

This is why the Trace Economy distinguishes between interaction and authorship, source and pathway, content and influence, private cognition and public reliance.

The protocol should not reward the number of human touches. It should identify the quality and accountability of human contribution.

It should not merely disclose that AI was used. It should disclose how AI shaped a public-reliance output where material.

It should not merely disclose that funding exists. It should identify how funding or institutional interests may affect reliance.

It should not merely show that a procedure occurred. It should reveal who had authority, what standard applied, what was done, and where departures occurred.

In this sense, trace is anti-laundering infrastructure. It prevents authorship laundering, evidence laundering, influence laundering, funding laundering, and responsibility laundering.

6.7 Human Sign-Off and the Anchor of Responsibility

The Trace Economy is human-centred because accountability cannot be delegated to a machine.

AI systems may assist, generate, rank, summarise, analyse, recommend, or decide within institutional workflows. But when outputs enter public reliance, a human or institution must accept responsibility for the reliance claim.

Human sign-off does not mean that one individual personally created every element. It means that an accountable person or institution stands behind the output, process, disclosure, and reliance conditions.

In some contexts, the sign-off may be individual: an author, lawyer, clinician, engineer, researcher, public servant, police officer, investigator, or executive.

In other contexts, the sign-off may be institutional: a university, company, regulator, government agency, publisher, platform, insurer, court, research body, or foundation.

The key is that responsibility cannot be hidden inside “the system.”

A useful principle is:

The machine may assist, but reliance requires an accountable anchor.

This is especially important for AI-mediated outputs. A platform cannot simply say the user relied on the answer if the platform shaped the output through hidden incentives. A deployer cannot simply blame the model if it used the model in a decision-support role. A professional cannot simply blame AI for an unverified filing, recommendation, or report. A government cannot simply blame automated logic for an administrative decision that affects citizens.

Trace attaches responsibility to the reliance point.

6.8 The Role of DOI Anchoring and Formal Records

The hashtag trace is a beginning, not an endpoint. For higher-reliance contexts, additional anchoring may be required.

DOI anchoring provides a more formal layer for published papers, protocols, reports, frameworks, datasets, and versioned public documents. A DOI creates a stable citation record and can support timestamped authorship, version control, and public reference.

Institutional registers may provide another layer. A professional body, regulator, insurer, court, government agency, or foundation may maintain formal records of trace-compliant outputs, licences, declarations, or audits.

Cryptographic timestamping may provide further integrity for evidence files, prompt archives, chain-of-custody records, or disclosure logs.

Procurement and insurance documents may incorporate trace conditions as contractual obligations.

The Trace Economy can therefore operate at multiple levels:

- public hashtag anchor,
- DOI or publication anchor,
- private evidence archive,
- institutional register,
- licensing record,
- audit record,
- insurance condition,
- regulatory compliance file,
- and legal evidentiary record.

The level required depends on the reliance burden.

6.9 Trace and Body-of-Work Continuity

A single trace event may be useful, but a body of trace events is more powerful.

Over time, public trace can demonstrate continuity of cognition, authorship, influence disclosure, public commitments, and institutional behaviour. This is especially important in AI-assisted work, where individual outputs may be difficult to attribute cleanly at sentence level.

A body of work may reveal:

- recurring concepts,
- consistent terminology,
- development of a framework,
- response to critique,

evolution of doctrine,
refinement of method,
public commitments,
collaborative contributions,
and cumulative authorship.

This creates developmental provenance.

Developmental provenance is evidence that a work belongs to a broader trajectory of human or institutional thought. It is harder to fabricate retrospectively than a single prompt, single document, or single claim. A prompt can be manufactured after the fact. A public body of work, timestamped over time, is much harder to fake.

This is why the Trace Economy is not merely document-level provenance. It also supports longitudinal authorship and institutional accountability.

The relevant question becomes:

Does this work arise from a traceable continuity of human cognition, public claim, and accountable development?

That question is central to AI-era authorship.

6.10 Public Trace and Open Infrastructure

A major advantage of the Trace Economy mechanism is that it can begin without waiting for large institutions to build new infrastructure.

Public platforms already support posts, timestamps, tags, mentions, comments, links, and archives. Publication repositories support DOI records. Documents support AI-use statements and funding declarations. Institutions already maintain records, approvals, audits, and compliance files. The Trace Economy connects these existing capacities into a common public-reliance logic.

This does not mean the system should remain informal forever. As adoption grows, formal infrastructure will be needed: governance bodies, licensing terms, registries, audit templates, disclosure standards, evidence categories, dispute procedures, and integration with insurers, regulators, professional bodies, and procurement systems.

But the initial mechanism is deliberately low-friction.

It allows a person or institution to begin tracing immediately:

This output has crossed into public reliance.

A human or institution stands behind it.

A trace claim is being made.

The deeper pathway can be audited if necessary.

This is why the three-hashtag protocol is powerful. It creates an immediate public signal while leaving room for formalisation.

6.11 What the Trace Mechanism Does Not Do

The Trace Economy mechanism has limits, and those limits must be stated clearly.

It does not determine truth by itself.

It does not prove authorship by hashtag alone.

It does not guarantee that a claim is ethical, accurate, lawful, or unbiased.

It does not require every private prompt or draft to be disclosed.

It does not eliminate the need for courts, regulators, auditors, experts, journalists, reviewers, or institutional processes.

It does not prevent all manipulation, bad faith, gaming, or false claims.

It does not replace domain-specific evidence standards.

Instead, it creates a public-reliance anchor. It gives claims, outputs, decisions, and recommendations a visible point of accountability and a pathway into evidence.

This modesty is important. The Trace Economy is not a magic proof system. It is a governance primitive.

Its strength lies in making reliance conditions visible.

6.12 Summary

The Trace Economy mechanism can be summarised as follows:

A public-reliance event occurs.

The claimant marks the event with protocol tags and direct tagging.

The trace marker publicly anchors the claim.

The claim type clarifies what is being asserted.

A deeper evidence trail supports the claim.

The trace burden scales with consequence.

Human or institutional sign-off anchors responsibility.

Influence, authorship, and provenance become reviewable.

This mechanism operationalises the paper's central doctrine:

Private cognition remains private.

Public reliance becomes traceable.

Influence becomes visible.

The next section applies this framework across domains, showing how the same public-reliance logic operates in AI systems, social media, grants and research, politics, policing, institutional accountability, and war-justifying claims.

7. Application Domains

The Trace Economy applies wherever people, institutions, or systems ask others to rely on something. This includes AI-generated outputs, but it is not limited to AI. Its broader category is public reliance: the moment a work, claim, recommendation, ranking, report, policy, investigation, or decision becomes socially, legally, economically, politically, or institutionally consequential.

The same three questions recur across domains:

Who contributed or stood behind it?

Where did it come from and how was it produced?

What interests, incentives, systems, or pressures shaped it?

This section applies the framework to several major domains. The examples are not exhaustive. They show how authorship, provenance, and influence can be used as a general public-reliance architecture.

7.1 AI Systems and LLM Outputs

AI systems are the most visible near-term use case because they compress authorship, provenance, and influence into a single interface. A user may ask a large language model for legal guidance, financial comparisons, health information, product recommendations, academic writing assistance, employment advice, political summaries, creative outputs, or strategic analysis.

The answer may appear coherent and neutral. Yet it may be shaped by model training, system instructions, retrieval sources, prior conversation history, memory, user profiling, commercial relationships, product integrations, safety rules, platform incentives, or institutional deployment settings.

The Trace Economy applies by requiring trace where the AI-mediated output enters public reliance.

For example, an AI-assisted legal submission should include accountable human sign-off and evidence verification. An AI-generated financial recommendation should disclose if products were ranked by sponsorship, affiliate incentive, or platform preference. An AI-assisted academic paper should identify the human contribution, role of AI assistance, source materials, and final responsibility. An AI-powered government

chatbot should disclose whether the response is based on eligibility rules, policy logic, human review, or automated triage.

The relevant issue is not merely whether AI was used. The relevant issue is whether the user, reader, regulator, court, client, citizen, or market can evaluate the authorship, provenance, and influence conditions behind the output.

This is especially important where LLMs become commercially integrated. If a system recommends products, services, insurers, lawyers, courses, investments, treatments, news sources, political content, or professional pathways, users should know whether the recommendation was shaped by evidence, personalisation, sponsorship, affiliate incentives, platform self-preferencing, behavioural optimisation, institutional policy, or prior user data.

In this context, influence disclosure becomes the AI equivalent of sponsored-search labelling and conflict-of-interest disclosure.

The central principle is:

If AI-mediated reasoning is commercially or institutionally shaped, the material influence factors should be disclosed at the point of reliance.

7.2 Social Media Harm Reduction

Social media platforms have already shown how influence can be embedded into digital environments. Harm may arise not only from individual pieces of content, but from recommendation loops, engagement incentives, virality mechanics, targeting systems, influencer economics, political amplification, and platform design choices.

This matters for youth harm, misinformation, extremism, body-image pressure, gambling-like engagement, consumer manipulation, political polarisation, and mental-health effects. In many cases, the issue is not simply whether a post is true or false. The issue is what influence architecture pushed the user toward it, amplified it, repeated it, monetised it, or associated it with identity, fear, belonging, outrage, aspiration, or insecurity.

The Trace Economy contributes a harm-reduction structure by shifting the question from content alone to reliance conditions.

A social media recommendation, campaign, endorsement, fundraising appeal, health claim, product claim, political message, or institutional statement may be traceable according to:

- who created or funded it,
- whether it was sponsored or affiliated,
- whether it was organically posted or strategically amplified,
- whether it was boosted by platform incentives,

whether the recommender has a vested interest,
whether the post is part of a coordinated campaign,
whether AI assisted in targeting or generating it,
and whether the public is being asked to rely on it.

This does not require censoring speech. It contextualises reliance.

A person may express an opinion. But where an individual or organisation seeks public reliance — especially through recommendations, endorsements, product claims, health advice, political messaging, or fundraising — the relevant influence pathway should be visible.

The principle is:

Trace does not suppress speech. It contextualises reliance.

This is particularly important for influencers and institutional actors. A recommendation from an individual may appear personal, while being paid, affiliated, sponsored, politically aligned, or strategically coordinated. A trace layer would make the interest structure visible without necessarily prohibiting the speech.

In this sense, the Trace Economy offers a free-speech-compatible approach to social media harm reduction. It does not ask platforms to decide all truth claims. It asks that public-reliance claims carry visible authorship, provenance, and influence conditions.

7.3 Grants, Funding, and Research Outcomes

Research, grant-funded projects, think-tank reports, consultancy findings, policy papers, NGO campaigns, industry studies, and institutional evaluations often enter public reliance. They may shape funding decisions, regulation, public opinion, government policy, academic debate, commercial markets, and social legitimacy.

The issue is not that funding invalidates work. Funded work can be rigorous, independent, and socially valuable. The problem is undisclosed or under-disclosed influence.

Funding can shape:

which questions are asked,
which methods are used,
which populations are studied,
which outcomes are emphasised,
which uncertainties are minimised,
which comparators are omitted,
which findings are promoted,
which language is used,
and which policy conclusions are drawn.

A study may be formally accurate but still skewed by design choices. A report may disclose funding but not explain how the funder's interests relate to the framing. A grant-funded program may report success according to metrics selected by the recipient or funder. A think-tank paper may present itself as independent while being aligned with donor interests.

The Trace Economy applies by making funding and influence part of public-reliance provenance.

A trace-compliant research or policy output would disclose:

who funded the work,
who designed the study or program,
what interests the funder may have,
what methodology was used,
what limitations exist,
what data was excluded,
who interpreted the results,
whether dissenting analysis exists,
and who signed off on the public claim.

This does not accuse funded work of corruption. It simply allows reliance to be calibrated.

The guiding principle is:

Funding is not the problem. Undisclosed material influence is the problem.

This is especially important where funded research is used to justify policy, procurement, public health messaging, environmental claims, defence spending, social programs, pharmaceutical recommendations, technology adoption, or regulatory reform.

The Trace Economy provides a structure for making research and grant outputs more publicly accountable without dismissing them merely because they are funded.

7.4 Political Campaigns and Democratic Influence

Politics is one of the clearest domains of public reliance. Voters are asked to rely on claims about candidates, policies, threats, economic outcomes, social harms, national identity, institutional trust, and future promises. Political campaigns also involve funding, consultants, donors, lobbyists, polling, message testing, media buying, targeting, and strategic omission.

Influence is not incidental to politics. It is central to it.

The issue is whether the public can see enough of the influence pathway to evaluate the claim.

A political message may be shaped by:

party interests,
major donors,
lobbyists,
industry groups,
unions,
foreign interests,
data brokers,
consultants,
microtargeting systems,
social media platforms,
AI-generated messaging,
or coordinated influence campaigns.

Traditional campaign-finance disclosure may reveal some funding, but it may not reveal how influence shaped the message itself. A donor may fund a campaign. A consultancy may test language. A platform may optimise delivery. A data system may identify vulnerable audiences. AI may generate tailored variants of the same political message. The final voter-facing output may appear like ordinary persuasion while being shaped by a complex influence pathway.

The Trace Economy would not prohibit political campaigning. It would require that public-reliance claims be more traceable.

For high-impact political claims, the relevant trace might include:

who funded the campaign,
who authorised the message,
whether AI generated or tailored it,
whether microtargeting was used,
which audience criteria were selected,
whether the claim relies on evidence,
whether the claim was fact-checked,
whether a sponsor or third party shaped the framing,
and who is accountable for the representation.

The principle is:

Democratic reliance requires visible influence.

Political speech remains free, but material influence should not be hidden where voters are being asked to act.

7.5 Professional Advice and Public Trust

Professional advice creates reliance because the recipient often lacks equal expertise. Lawyers, clinicians, engineers, accountants, financial advisers, auditors, architects, educators, psychologists, consultants, and other professionals make claims or recommendations that clients, courts, regulators, patients, businesses, and institutions may act upon.

Professional systems already contain duties of competence, independence, record-keeping, disclosure, and care. The Trace Economy does not replace those duties. It strengthens them by providing a public-reliance lens.

The relevant questions are:

- Who gave the advice?
- What authority or qualification did they rely upon?
- What evidence supported it?
- What assumptions were used?
- Was AI involved?
- Was the advice independently verified?
- Were conflicts disclosed?
- Were material limitations explained?
- Who signed off?
- What responsibility attaches?

This becomes especially important as professionals adopt AI tools. A lawyer using an LLM to assist with submissions must still verify authorities. A clinician using AI triage must still apply clinical judgment. A financial adviser using automated product comparisons must still disclose incentives and suitability. An engineer using generative design must still ensure compliance with standards and safety obligations.

The professional cannot launder responsibility through the tool.

The principle is:

The machine may assist, but professional reliance requires accountable human judgment.

Trace creates a way to show where that judgment occurred.

7.6 Policing, Investigations, and Authority Provenance

Policing demonstrates why trace must extend beyond evidence to authorised conduct.

Investigations already require chain of custody, evidence logs, warrants, statements, forensic records, disclosure obligations, and supervisory review. But modern policing also involves digital forensics, facial recognition, social media intelligence, predictive

tools, AI-assisted analysis, automated risk scoring, phone-location data, CCTV interpretation, informant intelligence, and complex inter-agency processes.

A public-reliance framework must ask not only what evidence was gathered, but who had authority, what they were trained to do, what they actually did, what they failed to do, what they relied on, what they ignored, and whether any departure from policy or training was justified.

This creates three distinct trace layers:

Evidence provenance — where the evidence came from, who handled it, how it was preserved, and whether it was altered or interpreted.

Authority provenance — what legal power, policy, warrant, delegation, or institutional authority justified the action.

Conduct provenance — what authorised persons did or failed to do, and whether that conduct aligned with training, procedure, law, and professional standards.

This is especially important where police conduct affects liberty, reputation, property, privacy, or public trust.

A Trace Economy approach would ask:

Was the officer trained for the situation?

What legal power was invoked?

What procedure applied?

Was the procedure followed?

If not, why not?

Was the departure documented?

Was supervision involved?

Was exculpatory evidence preserved?

Was AI or automated analysis used?

Was uncertainty disclosed?

Who approved escalation?

Who signed off on charge, search, arrest, or public accusation?

The principle is:

Authority must be traceable to conduct. Conduct must be traceable to training.

Departures must be traceable to justification.

This does not make the framework anti-police. Properly implemented, it protects citizens, good officers, courts, and institutions. It makes accountability evidentiary rather than rhetorical.

7.7 Public-Sector Decision-Making

Government decisions often create reliance because citizens may be required to obey, accept, challenge, or live under them. This includes welfare eligibility, tax decisions, immigration outcomes, licensing, education access, policing priorities, planning approvals, procurement decisions, public health directives, housing allocation, and regulatory enforcement.

As AI and automated systems enter public administration, the need for trace increases.

A citizen affected by a decision should be able to know:

what authority governed the decision,
what criteria were applied,
whether AI or automation was used,
what data was relied upon,
whether the data was accurate,
whether personalisation or risk scoring occurred,
whether human review was available,
what appeal pathway exists,
who signed off,
and what institutional policy shaped the outcome.

This is not only a technical issue. It is a democratic issue.

Public administration must be able to justify coercive or allocative decisions. Where AI assists, the justification burden does not disappear. It becomes more important because automated systems can obscure responsibility.

The principle is:

Public authority cannot hide inside automated process.

Trace provides a way to connect decision, evidence, authority, influence, and human or institutional accountability.

7.8 Commerce, Product Recommendations, and Platform Markets

Modern commerce increasingly depends on recommendations. Platforms rank products, compare services, recommend providers, show reviews, filter options, personalise offers, and guide purchasing decisions. These systems may be influenced by sponsorship, affiliate revenue, commissions, inventory arrangements, platform-owned products, advertising bids, user profiling, behavioural optimisation, or marketplace design.

A recommendation may appear helpful while being commercially shaped.

This is where the Google analogy becomes important. Sponsored search results are labelled because users need to know when commercial payment affects ranking. LLMs

and recommendation systems should meet an equivalent or higher standard where advice is personalised and conversational.

A trace-compliant commercial recommendation would disclose whether:

placement was paid,
affiliate commission applies,
ranking was commercially influenced,
platform-owned products were preferred,
prior user data shaped the recommendation,
behavioural optimisation was used,
reviews were filtered or weighted,
AI generated the comparison,
and what entity benefits from user action.

This would not prevent advertising or affiliate commerce. It would make the influence visible.

The principle is:

Commercial influence is permissible; concealed commercial influence in public-reliance advice is not.

7.9 Media, Journalism, and Institutional Narratives

Media institutions shape public reliance through story selection, framing, headlines, source selection, expert quotation, visual imagery, editorial emphasis, and repetition. Journalism already contains norms around attribution, corrections, editorial independence, source verification, and conflicts of interest. But the media environment has become more complex through sponsored content, native advertising, political alignment, platform distribution, AI-generated copy, influencer journalism, and algorithmic amplification.

The Trace Economy can help distinguish between reporting, commentary, sponsored content, institutional messaging, campaign material, and AI-assisted editorial production.

A trace-relevant media output may disclose:

authorship,
AI assistance,
source basis,
funding or sponsorship,
editorial independence,
institutional affiliations,
corrections,

source limitations,
and whether the content is news, opinion, advertisement, advocacy, or analysis.

This matters because the same factual content may carry different reliance implications depending on whether it is independent reporting, paid placement, political messaging, or institutional advocacy.

The principle is:

The public should be able to distinguish information from influence.

7.10 War, National Security, and the Highest Trace Burden

Declarations of war and military action represent one of the highest forms of public reliance. A state asks citizens, soldiers, allies, courts, industries, media, and international bodies to rely on claims that may result in death, destruction, civil-liberty restrictions, economic sacrifice, geopolitical realignment, and long-term social consequences.

In such contexts, the trace burden should be extremely high.

This does not mean that all classified material must be disclosed in real time. Some intelligence sources, operational methods, and strategic details may need protection. But secrecy cannot justify an absence of structured accountability.

A war-justifying public claim should carry trace across several dimensions:

Evidence provenance — what type of evidence supports the claim, what confidence level applies, and whether evidence categories can later be audited.

Authority provenance — who had legal and constitutional authority to act, and what parliamentary, congressional, judicial, or executive process applied.

Influence provenance — whether defence contractors, donors, lobbyists, foreign interests, think tanks, media campaigns, intelligence pressures, or political incentives shaped the public narrative.

Dissent provenance — whether alternative assessments, uncertainty, internal disagreement, or contrary intelligence existed.

Accountability provenance — who signed off, who reviewed, who advised, and what retrospective audit mechanism exists.

The Trace Economy does not decide whether war is justified. It makes the justification traceable.

The principle is:

A population should not be asked to rely on war claims without a traceable chain of evidence, authority, influence, and accountability.

This may be the highest expression of the public-reliance doctrine.

7.11 Insurance, Risk, and Underwriting

Insurance depends on trace because insurers price risk. As AI-mediated outputs and institutional decisions become more consequential, insurers will need to know whether organisations have reliable provenance, disclosure, sign-off, and audit systems.

An insurer covering professional liability, cyber risk, AI deployment, product advice, financial services, public-sector systems, or institutional decision-making may ask:

Was AI used?

Was there human sign-off?

Were influence factors disclosed?

Were users warned of limitations?

Were commercial incentives visible?

Were high-risk outputs reviewed?

Were audit logs retained?

Was the organisation operating under a recognised public-reliance protocol?

The Trace Economy could therefore function as a compliance and underwriting layer.

Organisations operating under trace may be able to demonstrate that they have procedures for authorship, provenance, influence disclosure, and accountable sign-off.

This may reduce uncertainty for insurers and create a basis for differentiated coverage, premiums, exclusions, or conditions.

The principle is:

Traceable reliance is insurable reliance.

This does not mean trace eliminates risk. It makes risk more visible, assignable, and governable.

7.12 Application Summary

Across these domains, the same structure recurs.

AI systems require trace because outputs may appear neutral while being shaped by hidden data, incentives, or systems.

Social media requires trace because harm often arises from hidden influence architecture, not content alone.

Research and grants require trace because funding can shape questions, methods, interpretation, and public messaging.

Politics requires trace because democratic decisions depend on visible influence.

Professional advice requires trace because clients rely on expertise they cannot fully evaluate themselves.

Policing requires trace because state authority must be accountable to evidence, training, procedure, and law.

Public-sector decisions require trace because citizens should not be governed by opaque process.

Commerce requires trace because recommendations may be commercially shaped.

Media requires trace because information and influence are increasingly blended.

War claims require trace because the consequences of reliance are existential.

Insurance requires trace because risk needs accountable pathways.

The Trace Economy unifies these domains through one principle:

Public reliance requires visible authorship, provenance, and influence.

The next section considers how this framework can be translated into licensing, regulation, procurement, insurance, and institutional governance.

8. Licensing, Regulation, and Insurance Implications

The Trace Economy is not only a conceptual framework. If it is to function as a public-reliance protocol, it must be capable of operational adoption. That means it must translate into licensing, regulatory design, procurement standards, professional duties, insurance conditions, institutional governance, and public-interest infrastructure.

The central proposition is simple:

Where a person or organisation asks others to rely on a claim, output, recommendation, report, decision, ranking, investigation, or institutional representation, the relevant authorship, provenance, and influence conditions should be traceable.

This proposition can operate through both voluntary and mandatory pathways. It can begin as a public-interest protocol adopted by individuals and institutions. It can then mature into contractual terms, industry standards, procurement criteria, insurance requirements, professional obligations, and eventually regulatory expectations.

The strength of the Trace Economy lies in its ability to begin lightly while scaling into formal governance.

8.1 From Protocol to Licence

The Trace Economy begins as a protocol: a set of rules, signals, and expectations for marking the transition from private cognition into public reliance.

A licence formalises that protocol.

A Trace Economy or PoCW-style licence could define:

- what constitutes public reliance,
- what trace obligations arise at different reliance levels,
- what authorship declarations are required,
- what AI-use disclosures are required,
- what influence factors must be disclosed,
- what evidence must be retained,
- what public anchors must be used,
- what human or institutional sign-off is required,
- what audit rights exist,
- what royalty or value-sharing arrangements apply,
- and what consequences follow from misuse or false trace claims.

This licence need not apply to every private use of AI, every private draft, or every casual communication. Its proper domain is public reliance.

A key formulation is:

The licence attaches not to private cognition, but to public reliance.

This avoids overreach. It also makes the licence more commercially and legally plausible. The Trace Economy does not need to control every underlying tool, model, platform, or private process. It governs the moment at which an output or claim is placed into reliance.

That distinction is crucial in relation to large language models. It may not be realistic, at least initially, to require every model provider to adopt the Trace Economy at the model layer. But organisations deploying LLMs in public-reliance contexts could operate under the licence at the reliance layer.

This includes law firms, insurers, banks, health providers, education providers, government agencies, AI governance vendors, professional services firms, research organisations, publishers, social media platforms, procurement bodies, and customer-facing AI deployers.

The practical licensing claim is therefore:

Any organisation using AI-mediated systems in public-reliance contexts can operate under a Trace Economy licence to evidence authorship, provenance, influence disclosure, and accountable sign-off.

8.2 The Reliance Layer as the Commercial Entry Point

The commercial entry point is not necessarily the AI model itself. It is the reliance layer.

The model provider controls infrastructure.

The platform controls interface and deployment.

The organisation controls use case and user relationship.

The professional or institution controls sign-off.

The public-reliance event controls accountability.

The Trace Economy licence should attach where accountability becomes necessary: the point at which the output is acted upon, published, filed, sold, recommended, used in decision-making, or presented to others as trustworthy.

This is commercially important because many organisations do not need a philosophical AI doctrine. They need a defensible way to show that their AI-assisted outputs, recommendations, decisions, or public claims are governed, auditable, and accountable.

A Trace Economy licence can become a trust signal:

This organisation does not merely use AI. It operates public-reliance outputs under a traceable authorship, provenance, and influence-disclosure protocol.

That is valuable to clients, consumers, regulators, insurers, investors, procurement officers, courts, auditors, and the public.

8.3 Regulatory Implications: From Content Disclosure to Influence Disclosure

Current AI regulation often focuses on content, system risk, user notification, safety, privacy, bias, and accountability. These are important. But they do not fully address the influence pathway behind AI-mediated outputs.

The next regulatory step is influence disclosure.

Regulation should not only ask:

Was AI used?

Was content generated or manipulated?

Was the system high-risk?

Was personal data processed?

Was there human oversight?

It should also ask:

What shaped the response?
Was prior user history used?
Was the output personalised?
Were sponsored or affiliate incentives involved?
Was ranking commercially influenced?
Was behavioural optimisation used?
Were sensitive inferred traits used?
Did institutional policy constrain the answer?
Did government eligibility logic shape the result?
Was the system acting as adviser, advertiser, recommender, or decision-support tool?

The Trace Economy provides a framework for this regulatory expansion.

A possible regulatory clause could read:

Where an AI system produces advice, recommendations, rankings, summaries, responses, or decision-support outputs that may materially influence a person's choices, the provider or deployer must disclose the material influence factors shaping that output, including commercial incentives, personalisation, paid placement, behavioural optimisation, institutional constraints, and use of longitudinal user data.

This clause does not require disclosure of every internal model parameter. It requires disclosure of material influence factors relevant to reliance.

That is the correct regulatory level. It is neither too shallow nor impossibly deep.

8.4 Consumer Protection

Consumer protection law is a natural home for influence disclosure.

Consumers are already protected against misleading or deceptive conduct, undisclosed material connections, unfair practices, hidden fees, manipulative design, and false representations. LLMs and AI recommendation systems create a new version of these problems because commercial influence may be embedded inside personalised reasoning.

The consumer-protection question is:

Did the user understand whether the recommendation was neutral, sponsored, affiliated, platform-preferred, personalised, or behaviourally optimised?

If not, then the user may be relying on an output whose influence pathway was concealed.

This matters for:

product recommendations,
financial products,

insurance comparisons,
education courses,
health services,
legal services,
travel recommendations,
employment platforms,
subscription products,
marketplaces,
investment products,
and AI-generated buying advice.

The Trace Economy could operate as a consumer-protection layer by requiring material commercial influence to be disclosed at the point of AI-mediated recommendation.

A simple rule follows:

If sponsored placement must be labelled in search, sponsored reasoning should be disclosed in AI.

This is one of the most accessible regulatory arguments for the framework.

8.5 Privacy and Longitudinal Cognitive Data

Privacy and data-protection law also become central because LLMs may generate or infer longitudinal cognitive data.

Traditional privacy regimes often focus on personal information, sensitive information, consent, collection, use, retention, disclosure, and security. But LLM interaction introduces a deeper category: the accumulated pattern of how a person thinks, reasons, creates, worries, plans, revises, and responds to persuasion.

This raises questions that existing privacy law may not yet handle clearly:

Who controls the user's cognitive development record?
Can it be used for advertising?
Can it be used for recommendation ranking?
Can it be used to infer vulnerabilities?
Can it be used to shape political or financial persuasion?
Can it be used to train future systems?
Can the user export it?
Can the user rely on it as evidence of authorship?
Can the provider deny access to it?
Can it be used against the user?

The Trace Economy's answer is not to expose private cognition. It is to protect it while creating trace at the point of public reliance.

The principle is:

Private cognitive development should remain under user control. Public reliance should carry trace.

Where an AI system uses longitudinal user data to shape advice, recommendations, rankings, or decision-support outputs, that use should be disclosed. Where the data remains private and exploratory, it should not become an automatic public trace record.

The regulatory implication is that longitudinal cognitive data may require heightened protection, especially against commercial targeting, behavioural manipulation, political persuasion, and sensitive inference.

8.6 Government Procurement

Government procurement is one of the most promising adoption pathways.

Governments increasingly purchase AI systems, data tools, decision-support software, analytics services, cybersecurity platforms, public-service chatbots, risk assessment tools, and administrative automation. Procurement terms can require vendors to meet standards before systems are deployed.

A government could require that AI vendors operating in public-reliance contexts provide:

AI-use disclosure,
source and data provenance,
influence disclosure,
human review pathways,
audit logs,
public-facing explanation,
appeal mechanisms,
bias and risk assessments,
training records,
institutional sign-off,
and compliance with a recognised trace protocol.

This would not require governments to adopt every aspect of the Trace Economy immediately. They could begin with influence disclosure and public-reliance classification.

Procurement is powerful because it can shape markets. If public agencies require trace-compliant AI systems, vendors will adapt.

A procurement clause could state:

Any AI system used to provide citizen-facing advice, eligibility guidance, ranking, triage, recommendation, enforcement support, or decision assistance must maintain traceable authorship, provenance, influence disclosure, and human sign-off records proportionate to the reliance consequence.

That is a practical pathway from theory to implementation.

8.7 Professional Regulation

Professional bodies can also adopt trace requirements.

Lawyers, doctors, engineers, accountants, financial advisers, auditors, architects, educators, psychologists, consultants, and public servants increasingly use AI-assisted tools. Their professional obligations already require competence, judgment, accountability, and disclosure of conflicts.

Trace provides a way to operationalise those obligations in AI-mediated environments.

Professional regulators could require:

AI-use declarations for high-impact work,
verification of AI-assisted claims,
retention of material source records,
human sign-off before reliance,
disclosure of commercial or institutional influence,
documentation of departures from standard practice,
and preservation of evidence trails where clients, courts, patients, regulators, or the public rely on outputs.

For example:

A lawyer filing AI-assisted material should verify authorities and disclose AI use where required by court or professional rules.

A clinician using AI triage should document human review and clinical reasoning.

A financial adviser using AI comparisons should disclose product incentives and suitability logic.

An engineer using AI-generated designs should document compliance checks and human responsibility.

The principle is:

AI may assist professional work, but it cannot replace professional accountability.

Trace makes that accountability visible and reviewable.

8.8 Insurance and Underwriting

Insurance may become the strongest commercial driver for trace adoption.

Insurers need to price risk. As organisations use AI in public-facing, advisory, commercial, professional, or decision-support contexts, insurers will need to know whether the organisation has governance controls that reduce liability.

A Trace Economy licence or protocol can function as evidence of risk management.

Insurers might ask:

Does the organisation classify public-reliance outputs?

Does it retain AI-use records?

Does it require human sign-off?

Does it disclose commercial influence?

Does it use longitudinal user data for recommendations?

Does it separate sponsored from independent advice?

Does it log institutional constraints?

Does it maintain audit trails?

Does it have escalation procedures for high-risk outputs?

Does it preserve evidence if harm occurs?

Where the answer is yes, risk may be more governable. Where the answer is no, insurers may exclude coverage, increase premiums, or require remediation.

This gives rise to the principle:

Traceable reliance is insurable reliance.

The Trace Economy could therefore become an underwriting standard for AI-mediated advice, professional liability, cyber risk, product recommendation systems, public-sector technology, educational technology, healthcare AI, legal AI, and financial decision-support tools.

This is commercially significant because insurance can force adoption faster than abstract ethics frameworks. Organisations often change behaviour when coverage, premiums, exclusions, or liability exposure are affected.

8.9 Institutional Governance and Internal Controls

Institutions can adopt trace internally even before external regulation requires it.

A company, university, NGO, government agency, law firm, research institute, or platform can create internal controls for public-reliance outputs. These controls may include:

public-reliance classification,

authorship and contributor records,

AI-use registers,

funding and influence declarations,
approval workflows,
human sign-off requirements,
audit logs,
version histories,
escalation thresholds,
legal review,
and external trace anchoring.

The internal version of trace may be private. Not every internal deliberation needs to be public. But when the institution releases a public-reliance output, the public-facing trace marker can connect to a controlled evidence file.

This distinction is important:

Internal trace supports governance. Public trace supports reliance.

A mature institution may need both.

8.10 Platform Governance

Platforms have a special role because they control distribution, ranking, monetisation, recommendation, interface design, and disclosure architecture.

A social media platform, search engine, AI assistant, marketplace, app store, news aggregator, education platform, or financial comparison platform can shape what users see, believe, buy, share, or rely upon.

The Trace Economy's influence layer applies directly to platforms because platforms are not neutral pipes. They structure attention.

Platform trace obligations could include:

disclosure of sponsored ranking,
separation of paid and organic recommendations,
labels for AI-generated or AI-assisted outputs,
visibility into platform self-preferencing,
records of behavioural optimisation in high-risk contexts,
disclosure of affiliate relationships,
limits on using sensitive inferred traits,
records of political or issue-based targeting,
and public archives for high-impact campaigns.

The principle is:

Platforms that shape reliance should disclose material influence pathways.

This is especially important for LLM platforms because conversational outputs can make platform influence less visible than ordinary feeds or search results.

8.11 Public-Interest Governance and the Foundation Role

For the Trace Economy to operate at scale, a governance body is needed. This may take the form of a foundation, public-interest company, nonprofit entity, standards body, or licensing office.

The foundation's role would not be to approve every claim or decide all disputes. Its role would be to maintain the protocol, licence, glossary, trace standards, claim categories, public guidance, audit templates, adoption pathways, and recognised use conditions.

A foundation could provide:

- standard trace language,
- licensing terms,
- claim-type definitions,
- AI-use declaration templates,
- influence disclosure templates,
- public-reliance classification guides,
- dispute-resolution pathways,
- training materials,
- institutional onboarding,
- public registry options,
- and interfaces for regulators, insurers, researchers, and adopters.

The three-hashtag protocol and direct tagging provide the early public anchor. Over time, the foundation can strengthen that anchor through formal registers, DOI integration, certification, audit partnerships, and licensing enforcement.

This gives the Trace Economy a path from public signal to institutionally recognised infrastructure.

8.12 Regulation Without Overreach

A major risk in any trace framework is overreach. If trace is framed as universal monitoring, it will fail and should fail. The regulatory design must preserve freedom of thought, speech, creativity, research, and private inquiry.

The Trace Economy avoids overreach through three safeguards.

First, private cognition remains private. Trace does not apply to every prompt, draft, note, conversation, or unfinished idea.

Second, public reliance is the trigger. Trace applies when others are being asked to trust, act upon, fund, obey, insure, regulate, prosecute, publish, or commercially value something.

Third, trace burden scales with consequence. Low-stakes contexts do not require high-stakes evidence burdens.

These safeguards make the framework proportionate.

The goal is not to create a surveillance regime. The goal is to prevent hidden authorship, hidden provenance, hidden influence, hidden incentives, and hidden responsibility from shaping public reliance without accountability.

8.13 Adoption Pathway

A practical adoption pathway could unfold in stages.

Stage 1: Voluntary public trace

Individuals and organisations use the three protocol hashtags and tagging mechanism to mark public-reliance outputs.

Stage 2: Standard declarations

AI-use statements, influence disclosures, funding disclosures, authorship declarations, and public-reliance classifications are standardised.

Stage 3: DOI and publication integration

Papers, reports, protocols, and public frameworks are anchored through DOI records and version histories.

Stage 4: Institutional licensing

Organisations adopt the Trace Economy licence for AI-mediated outputs, professional advice, research, public reports, and recommendations.

Stage 5: Insurance and procurement uptake

Insurers and public-sector buyers require trace-compliant controls in high-reliance contexts.

Stage 6: Regulatory recognition

Regulators recognise influence disclosure, public-reliance trace, and human sign-off as part of compliance expectations.

Stage 7: Public registry and audit infrastructure

Formal trace registries, audit pathways, dispute mechanisms, and certification layers develop around the protocol.

This staged approach allows the system to begin immediately without waiting for full regulation.

8.14 Licensing as Public-Interest Infrastructure

The licensing model should not be understood merely as a private commercial mechanism. It is also public-interest infrastructure.

A licence can define the conditions under which organisations may claim trace compliance. It can prevent misuse of the protocol. It can create obligations around disclosure, evidence retention, sign-off, and audit. It can support funding of the governance body. It can create a basis for training, enforcement, and recognised adoption.

The licence therefore has two functions.

It protects the integrity of the Trace Economy.

It enables organisations to use the protocol in a standardised and accountable way.

This is why licensing matters. Without licensing, “trace” risks becoming a vague label that anyone can use without discipline. With licensing, trace can become a governed public-reliance standard.

8.15 Summary

The Trace Economy can become operational through a layered adoption model:

public protocol,
licence,
declaration templates,
institutional controls,
professional duties,
procurement standards,
insurance requirements,
regulatory recognition,
foundation governance,
and audit infrastructure.

Its core value is that it translates an abstract accountability principle into a usable public-reliance mechanism.

The key claim is:

The Trace Economy licence does not need to control every AI model, platform, or private cognitive act. It governs the moment output becomes relied-upon human, institutional, commercial, legal, or public action.

That is the regulatory and commercial sweet spot.

The next section addresses objections and limits, including surveillance concerns, hashtag reliability, gaming risks, free speech, institutional resistance, and the danger of treating trace as proof rather than as an evidentiary anchor.

9. Objections and Limits

Any framework that proposes a new public-reliance layer must face objections directly. The Trace Economy is intentionally broad. It applies across AI, research, commerce, social media, professional advice, policing, politics, institutional decision-making, and high-stakes public claims. That breadth is its strength, but it also creates risks: overreach, misunderstanding, misuse, weak evidentiary claims, superficial compliance, institutional resistance, and confusion between trace and truth.

This section addresses those objections and clarifies what the Trace Economy does and does not claim to do.

9.1 Is This Surveillance?

The strongest objection is that a trace framework could become surveillance of thought.

This concern is legitimate. A badly designed provenance system could demand constant logging of prompts, drafts, private notes, conversations, internal deliberations, and unfinished ideas. Such a system would be invasive and would undermine creativity, inquiry, professional judgment, and personal autonomy.

The Trace Economy rejects that approach.

Its core doctrine is:

Private cognition remains private.

Public reliance becomes traceable.

Influence becomes visible.

The trace obligation does not arise because a person thinks, drafts, prompts, experiments, researches, or privately develops an idea. It arises when a claim, output, recommendation, decision, report, policy, or representation is placed before others as something to rely upon.

The distinction is critical.

Private thought is not the regulated object. Public reliance is.

This means that a person may privately use AI, explore ideas, develop drafts, discard outputs, refine arguments, and think freely without public trace. But once the person or

institution publishes, sells, files, recommends, markets, licenses, governs, prosecutes, funds, or otherwise places the material into reliance, accountability conditions attach.

The Trace Economy is therefore not a surveillance framework. It is a reliance framework.

9.2 Can Hashtags Really Prove Anything?

A second objection is that hashtags are too informal to carry evidentiary weight.

The answer is that hashtags do not prove the whole claim.

They anchor it.

A trace tag does not prove authorship, truth, originality, neutrality, legal validity, scientific rigour, or ethical status by itself. It marks a public-reliance event and connects that event to a protocol claim. It creates a timestamped, searchable public signal that can later be supported, challenged, audited, or contextualised.

The hashtags are therefore analogous to a doorway, not the entire building.

They point toward a deeper evidence trail, which may include drafts, notes, source materials, funding records, conflict disclosures, AI-use statements, prompt summaries, chain-of-custody records, institutional approvals, human sign-off, version histories, DOI records, audit logs, or licensing records.

The principle is:

The trace tag is not the proof. It is the public anchor that makes proof easier to organise.

This modest claim is important. It prevents the Trace Economy from overstating what a public tag can do, while still recognising the practical value of public timestamped indexing.

9.3 Could the System Be Gamed?

Yes. Any disclosure or accountability system can be gamed.

People can make false declarations. Organisations can hide interests. Platforms can use vague language. Funders can shape work indirectly. Institutions can comply formally while avoiding substantive accountability. AI users can manufacture long prompt logs to simulate authorship. Public actors can attach trace tags to weak, misleading, or incomplete claims.

The risk of gaming is real.

But the existence of gaming does not make disclosure systems useless. Financial disclosures can be abused. Conflict-of-interest statements can be incomplete. Chain-of-custody records can be challenged. Academic citations can be misleading.

Sponsored-content labels can be obscured. Yet societies still use these mechanisms because they create a basis for accountability, challenge, and sanction.

Trace reduces gaming in several ways.

First, it creates public timestamping. Claims made publicly can be compared against later versions, competing claims, and prior records.

Second, it supports body-of-work continuity. A single trace claim can be faked more easily than a long public developmental record.

Third, it distinguishes claim types. A person must clarify whether they are claiming authorship, funding disclosure, AI use, influence disclosure, authority, conduct, or public reliance.

Fourth, it links public markers to deeper evidence. A trace claim can be audited if challenged.

Fifth, it creates reputational consequences. False trace claims become visible, contestable, and potentially sanctionable under licence, regulation, contract, professional duty, or public criticism.

Trace does not eliminate bad faith. It makes bad faith harder to hide.

9.4 Does Trace Determine Truth?

No.

Trace does not determine whether a claim is true. It does not replace scientific method, legal proof, journalism, peer review, professional judgment, democratic debate, regulatory investigation, or court process.

Trace answers a different question.

It asks whether the conditions of reliance are visible.

A traced claim may still be false. An untraced claim may still be true. But a traced claim provides a clearer pathway for evaluation: who made it, where it came from, what shaped it, what evidence supports it, and who stands behind it.

Truth and trace are related but distinct.

Trace does not decide truth. It supports truth-seeking by making authorship, provenance, influence, and accountability more visible.

This distinction prevents the Trace Economy from becoming an epistemic authority that declares what is true. Its role is to make reliance conditions legible.

9.5 Is This Just Existing Disclosure Repackaged?

The Trace Economy builds on existing disclosure systems but is not merely a restatement of them.

Existing systems are fragmented. Sponsored-search labels disclose paid ranking. Product-placement rules address embedded advertising. Funding declarations disclose some research interests. Chain-of-custody records trace evidence handling. Professional duties attach responsibility to expert judgment. AI-content labels disclose machine generation. Campaign-finance laws disclose some political funding.

Each system addresses part of the problem.

The Trace Economy unifies them around a broader category: public reliance.

Its contribution is the integration of three pillars:

Authorship identifies the contributor.

Provenance traces the pathway.

Influence reveals the forces.

Existing systems often address one or two of these, but rarely all three together.

A funded report may disclose authorship and funding but not influence pathway.

An AI label may disclose machine generation but not commercial steering.

A citation may disclose source but not institutional pressure.

A product recommendation may disclose sponsorship but not behavioural profiling.

A police evidence record may trace chain of custody but not authorised conduct or training departure.

A political ad may disclose sponsor but not data-driven targeting or message-testing influence.

The Trace Economy's novelty lies in treating authorship, provenance, and influence as a unified public-reliance architecture.

9.6 Will This Chill Speech?

A trace framework could chill speech if applied too broadly or coercively. That is why the trigger must remain public reliance, not ordinary expression.

A person should be free to express opinions, speculate, joke, criticise, create, and explore without heavy trace burdens. Not every social media post requires formal provenance. Not every personal recommendation needs institutional disclosure. Not every creative work requires evidentiary documentation.

The trace burden must scale with consequence.

The framework becomes relevant where others are being asked to rely on a claim in a meaningful way: commercially, legally, professionally, politically, institutionally, medically, financially, evidentially, or publicly.

This makes trace compatible with free speech.

The principle is:

Trace does not suppress speech. It contextualises reliance.

People may still speak. Organisations may still advocate. Funders may still support work. Platforms may still recommend products. Governments may still make public claims. But where reliance consequences arise, relevant authorship, provenance, and influence should be visible.

9.7 Who Decides the Trace Burden?

Another objection is that different contexts require different evidence standards. This is correct.

The Trace Economy should not impose one rigid burden across every domain. A casual product comment, funded policy report, professional medical recommendation, police investigation, and war claim do not require identical trace.

The framework therefore uses a scaling principle:

The greater the consequence of reliance, the higher the trace burden.

Specific trace burdens should be developed through domain-specific standards, professional rules, regulatory guidance, institutional policy, licensing terms, and public-interest governance.

For example:

AI-assisted creative work may require authorship and AI-use declarations.

Commercial recommendations may require sponsorship and affiliate disclosure.

Funded research may require funding, methodology, limitation, and conflict disclosure.

Professional advice may require evidence, competence, and human sign-off.

Police investigations may require evidence, authority, conduct, training, and disclosure trace.

War claims may require evidence category, authority, influence, dissent, and retrospective audit.

The Trace Economy supplies the common logic. Domain standards supply the detail.

9.8 Does This Require Platforms to Reveal Trade Secrets?

Not necessarily.

Influence disclosure does not require full exposure of proprietary models, internal parameters, ranking algorithms, or trade secrets. It requires disclosure of material influence factors relevant to reliance.

A platform may not need to reveal the entire ranking system. But it should disclose whether ranking was materially shaped by paid placement, affiliate revenue, platform self-preferencing, behavioural optimisation, prior user data, sensitive inferred traits, or institutional constraints.

A model provider may not need to publish every system instruction. But where an output is constrained by commercial, governmental, employer, or institutional logic, that fact may be reliance-relevant.

A government may not need to disclose classified sources. But where classified claims justify public action, there should be structured evidence categories, authority trace, legal sign-off, oversight, and retrospective audit.

The framework distinguishes between total transparency and reliance-relevant disclosure.

The public does not need to see everything. It needs to see enough to calibrate reliance.

9.9 What About Confidentiality and Security?

Some contexts require confidentiality: legal advice, medical records, trade secrets, investigative operations, national security, personal data, whistleblower protection, and internal deliberation.

The Trace Economy must accommodate this.

Trace does not mean universal public disclosure. It means accountable pathway.

In confidential contexts, the evidence trail may be held privately, institutionally, or under seal, while the public trace marker identifies that a reliance claim exists. Auditors, courts, regulators, professional bodies, or authorised reviewers may access deeper records where appropriate.

For example:

A legal filing may disclose AI use and human sign-off without exposing privileged drafts.

A police investigation may preserve operational secrecy while maintaining auditable records of authority, evidence, and conduct.

A national-security claim may protect sources while still recording evidence categories, confidence levels, legal authority, and retrospective review mechanisms.

A company may preserve trade secrets while disclosing that a recommendation was commercially influenced.

The principle is:

Trace does not require reckless disclosure. It requires accountable disclosure.

9.10 Can Institutions Resist or Co-Opt It?

Yes. Institutions may resist trace because opacity protects power. They may also co-opt the language of trace while avoiding substantive accountability.

This is a predictable risk.

Institutions may adopt superficial trace labels. Platforms may create vague disclosures. Governments may narrow trace obligations. Corporations may treat trace as brand reputation rather than public accountability. Professional bodies may resist additional burdens. Political actors may disclose formal funding while hiding practical influence.

The answer is not to assume institutional goodwill. The answer is to make the protocol specific, claim-based, auditable, and licence-governed.

Trace claims must be tied to defined categories.

Declarations must be standardised.

Audit pathways must exist.

False or misleading trace claims must carry consequences.

Public reliance burdens must scale with consequence.

The governance body must preserve protocol integrity.

Users, journalists, regulators, courts, insurers, and civil society must be able to challenge claims.

Co-option cannot be eliminated, but it can be made more difficult.

9.11 Could Trace Become Bureaucratic?

Yes, if poorly implemented.

A trace framework could become a paperwork exercise where organisations fill out forms without improving accountability. This is the “tick-and-flick” danger. A person might record prompts without meaningful authorship. A company might provide vague AI-use disclosures. A platform might label content while hiding influence. A funder might disclose support without revealing how incentives shaped outcomes.

The antidote is to focus on materiality.

Trace should not ask for endless procedural records. It should ask for the material factors relevant to reliance.

Who stands behind the claim?

What evidence or source pathway matters?

What systems shaped it?

What interests materially influenced it?

What human judgment occurred?

What reliance consequence exists?

This keeps trace substantive rather than bureaucratic.

The principle is:

Trace should measure reliance-relevant contribution and influence, not procedural noise.

9.12 Does Public Trace Create Legal Risk for Users?

Potentially, yes. Any public declaration can create legal, reputational, or professional consequences. But that is also part of accountability.

The Trace Economy should not encourage careless public claims. Users must understand that trace means standing behind something. A trace tag should not be used casually for work that the claimant is not willing to defend.

This is why education, claim types, and licensing guidance are important. A trace claim should specify what is being asserted. A person may claim authorship contribution, not factual certainty. A company may claim disclosure compliance, not product superiority. A researcher may claim funding transparency, not absolute neutrality. A platform may claim influence disclosure, not total absence of bias.

Clear claim types reduce unnecessary risk.

9.13 Is the Framework Too Broad?

The framework is broad because public reliance is broad. However, breadth does not mean every context receives the same rule.

The Trace Economy is not a single rigid template imposed everywhere. It is a common architecture that adapts to context.

Its universal principle is:

If people are being asked to rely on something, they should be able to see who stands behind it, where it came from, and what shaped it.

The application differs by domain.

This is no different from how evidence, disclosure, professional responsibility, and conflict-of-interest principles already operate. They have common logics but domain-specific rules.

The Trace Economy's breadth is therefore not conceptual overreach. It reflects the breadth of reliance itself.

9.14 What Trace Cannot Fix

The Trace Economy is not a cure for all institutional failure.

It cannot guarantee honesty.
It cannot eliminate corruption.
It cannot prevent all manipulation.
It cannot make all users critical thinkers.
It cannot force hostile regimes to become transparent.
It cannot replace courts, regulators, journalists, auditors, professional bodies, or democratic institutions.
It cannot make complex evidence simple.
It cannot remove the need for judgment.

What it can do is make reliance conditions more visible.

It can create public anchors.
It can connect claims to evidence.
It can identify responsibility.
It can reveal material influence.
It can support audit.
It can reduce hidden persuasion.
It can strengthen institutional accountability.
It can provide a common language for authorship, provenance, and influence.

That is enough to make it valuable.

9.15 Summary

The Trace Economy must be presented with disciplined limits.

It is not surveillance.
It is not truth by hashtag.
It is not censorship.
It is not total transparency.
It is not a replacement for law, regulation, or expertise.
It is not immune to gaming or co-option.

It is a public-reliance protocol.

Its purpose is to mark the moment when claims, outputs, recommendations, decisions, or institutional representations become consequential to others and to make the relevant authorship, provenance, influence, and accountability conditions visible.

The framework's strength lies in its simplicity and restraint:

Private cognition remains private.
Public reliance becomes traceable.
Influence becomes visible.

The final section concludes by situating the Trace Economy as a public-reliance operating system for the AI era and beyond.

10. Conclusion: From Opacity to Traceable Reliance

The central problem addressed in this paper is not artificial intelligence alone. Nor is it authorship alone, provenance alone, misinformation alone, advertising alone, or institutional accountability alone. These are all expressions of a deeper structural issue: people are increasingly asked to rely on claims, outputs, recommendations, rankings, reports, decisions, and public narratives without sufficient visibility over who created them, where they came from, or what interests shaped them.

This is the public-reliance problem.

AI has made the problem more visible because large language models can generate fluent, personalised, and apparently neutral outputs at scale. But the same underlying issue exists across social media, funded research, politics, professional advice, policing, commerce, government decision-making, media, institutional reports, and war-justifying claims. In each domain, reliance may be shaped by hidden authorship, hidden provenance, hidden funding, hidden incentives, hidden authority, hidden omissions, or hidden influence.

The Trace Economy responds by offering a unified public-reliance framework built on three pillars:

Authorship identifies the contributor.

Provenance traces the pathway.

Influence reveals the forces.

These three pillars must operate together. Authorship without provenance can become assertion. Provenance without authorship can become mechanical process without responsibility. Authorship and provenance without influence can conceal the interests, incentives, funding, behavioural design, platform logic, institutional pressure, or political forces that materially shaped the result.

Influence is therefore not an optional add-on. It is the missing third layer.

This matters especially in the AI era, where influence may be embedded not merely in advertisements, search rankings, product placements, or social media feeds, but inside personalised reasoning itself. The concept of sponsored reasoning captures this shift. If search engines must label sponsored links, and product placement is recognised as advertising blended into narrative, then AI systems should disclose when commercial,

behavioural, institutional, or political influence is blended into advice, recommendations, rankings, summaries, or decision-support outputs.

The same logic applies beyond AI. Grant-funded research should reveal material funding influence. Political campaigns should reveal donor and targeting influence. Social media recommendations should reveal paid, affiliated, or strategically amplified influence. Police investigations should trace not only evidence but authority and conduct. Professional advice should preserve accountable human judgment. Public-sector decisions should not hide inside automated process. War-justifying claims should carry the highest burden of trace because they ask populations to rely on claims with life-and-death consequences.

The Trace Economy does not seek to monitor thought. It does not require every prompt, draft, note, conversation, or private creative step to be exposed. Such a system would be intrusive and incompatible with genuine cognition. The creative process is often non-linear. Ideas emerge through uncertainty, correction, failure, surprise, and later recognition. The full picture is often only visible once the brush has stopped moving.

For that reason, the framework's trigger is not private cognition. It is public reliance.

The operating doctrine is:

Private cognition remains private.

Public reliance becomes traceable.

Influence becomes visible.

This doctrine protects privacy while preserving accountability. It allows people to think, draft, prompt, create, test, and revise freely. But once the result is placed into public reliance — published, sold, filed, recommended, cited, licensed, funded, governed, prosecuted, insured, or used to justify action — a visible point of accountability should attach.

The Trace Economy's initial mechanism is deliberately simple. A public trace can begin through protocol hashtags — **#TraceEconomy**, **#PoCW**, and **#Unifaircation** — together with direct tagging of the architect or relevant foundation. This creates a public, timestamped, searchable anchor. The tag does not prove everything. It does not replace evidence. It does not determine truth. It does not certify validity by itself.

It anchors the claim.

That anchor says: this work, claim, recommendation, decision, or output has entered public reliance; a human or institution is standing behind it; a provenance claim is being made; and the deeper evidence trail can be audited if necessary.

The deeper evidence may include drafts, source materials, funding declarations, prompt records, AI-use statements, institutional approvals, chain-of-custody logs,

human sign-off, DOI records, version histories, conflict disclosures, audit files, or licensing records. The trace burden scales with consequence. A casual opinion does not require the same trace as a legal filing. A product recommendation does not require the same trace as a medical recommendation. A funded report does not require the same trace as a police investigation. A public policy claim does not require the same trace as a declaration of war.

The principle is:

The greater the consequence of reliance, the higher the trace burden.

This makes the framework proportionate. It prevents trace from becoming a universal bureaucracy while still recognising that high-consequence reliance requires stronger accountability.

The Trace Economy should therefore be understood as a governance primitive. Its simplicity is its strength. A signature is simple. A timestamp is simple. A citation is simple. A DOI is simple. A disclosure label is simple. These mechanisms matter because they create recognisable points of accountability. The Trace Economy public anchor belongs in that family. It is not the whole proof; it is the public doorway into proof.

This paper has argued that the Trace Economy can operate across multiple domains as a public-reliance protocol, licensing framework, compliance layer, procurement standard, insurance condition, and civic accountability tool. Its relevance to AI is immediate, but its scope is broader. AI is one use case. Public reliance is the category.

The future of governance will not be secured merely by labelling AI-generated content. Nor will it be secured by asking whether humans touched outputs. Nor will it be secured by source citations alone. The deeper question is whether the public can see who stands behind a claim, where it came from, and what shaped it.

That is the shift from fragmented transparency to traceable reliance.

The Trace Economy does not decide truth. It makes reliance conditions visible.

It does not censor speech. It contextualises reliance.

It does not surveil private thought. It anchors public claims.

It does not replace law, regulation, professional judgment, evidence, journalism, audit, or democratic process. It gives those systems a clearer pathway into authorship, provenance, influence, and accountability.

The final proposition is therefore simple:

When people are asked to rely on something, they should be able to see the authorship, provenance, and influence pathway behind it.

That is the public-reliance standard.

That is the foundation of the Trace Economy.

And in an era of AI-mediated outputs, sponsored reasoning, platform influence, institutional opacity, and increasingly consequential public claims, it may become one of the essential accountability layers of the twenty-first century.

Reference List

Australian Competition and Consumer Commission. (2025). *Digital Platform Services Inquiry: Final report — March 2025*. The ACCC's final report reinforces the need for regulatory reform to address digital platform competition and consumer harms, including unfair trading practices and digital platform accountability.

Australian Competition and Consumer Commission. (2025). *Digital platforms and services*. ACCC. This source supports discussion of manipulative conduct, false practices, dangerous goods online, and ongoing competition and consumer-protection concerns in digital markets.

Australian Electoral Commission. (2026). *Financial disclosure*. AEC. This source supports the paper's discussion of political finance transparency, explaining that the Commonwealth disclosure scheme is intended to increase transparency and inform the public about financial dealings of political parties, candidates, and others involved in the electoral process.

Australian Electoral Commission. (n.d.). *Transparency Register*. AEC. This source supports the discussion of political campaign funding disclosure and public registers for election, referendum, donor, party, and associated-entity returns.

Australian Government Department of Finance. (2025). *Implementing Australia's AI Ethics Principles in government*. This source supports the paper's discussion of transparency, accountability, contestability, privacy protection, reliability, safety, and human-centred values in public-sector AI governance.

Australian Government Department of Industry, Science and Resources. (2019). *Australia's AI Ethics Principles*. This source supports the discussion of voluntary AI principles, including transparency, explainability, accountability, privacy protection, reliability, and safety.

Australian Government Digital Transformation Agency. (n.d.). *Pilot AI assurance framework: Step 7 — Transparency and explainability*. This source supports the

argument that people should understand when they are significantly impacted by AI and when an AI system is engaging with them.

Australian Information Commissioner. (2026). *Automated decision-making and public reporting under the Freedom of Information Act*. OAIC. This report supports the paper's public-sector decision-making discussion by assessing transparency in Australian Government agencies' communication about automated decision-making.

eSafety Commissioner. (2026). *Social media age restrictions*. This source supports discussion of platform responsibility and child online safety, including the requirement for age-restricted platforms to take reasonable steps to prevent under-16s from having accounts and the regulatory emphasis on platform accountability.

European Parliament and Council of the European Union. (2022). *Regulation (EU) 2022/2065 on a Single Market for Digital Services and amending Directive 2000/31/EC (Digital Services Act)*. This source supports the paper's discussion of online platform transparency, advertising, recommender systems, systemic risk, and digital service accountability.

European Parliament and Council of the European Union. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. This source supports the paper's discussion of AI transparency obligations, provider/deployer responsibilities, and the movement from AI-content labelling toward broader public-reliance traceability.

European Union. (2022). *Digital Services Act, Article 27: Recommender system transparency*. This source supports the argument that platform recommender systems already face transparency expectations regarding main parameters and user options to modify or influence those parameters.

Federal Trade Commission. (n.d.). *The FTC's Endorsement Guides: What people are asking*. FTC. This source supports the paper's discussion of material-connection disclosure, sponsored influence, endorsement transparency, and responsibility for deceptive or insufficiently disclosed commercial relationships.

Federal Trade Commission. (2024). *16 CFR Part 255 — Guides Concerning Use of Endorsements and Testimonials in Advertising*. Electronic Code of Federal Regulations. This source supports the paper's analogy between endorsement disclosure and AI-mediated influence disclosure, including liability for failure to disclose unexpected material connections.

Information Commissioner's Office & The Alan Turing Institute. (2025). *Explaining decisions made with AI*. UK Government / ICO. This guidance supports the paper's discussion of explaining processes, services, and decisions delivered or assisted by AI to affected individuals.

Information Commissioner's Office. (2023). *Guidance on AI and data protection*. ICO. This source supports the paper's privacy and data-protection discussion, especially the application of transparency and data protection principles to AI systems processing personal data.

National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. U.S. Department of Commerce. This source supports the discussion of AI risk management, trustworthiness, accountability, transparency, and governance across the AI lifecycle.

National Institute of Standards and Technology. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. U.S. Department of Commerce. This source supports the paper's discussion of generative AI governance, transparency risks, accountability, ownership, AI actor roles, and human oversight responsibilities.

OECD. (2019). *Recommendation of the Council on Artificial Intelligence*. OECD Legal Instruments. This source supports the paper's reliance on transparency, explainability, responsible disclosure, human-centred values, accountability, and trustworthy AI principles.

OECD.AI. (n.d.). *Transparency and explainability: OECD AI Principle 1.3*. This source supports the claim that AI actors should provide meaningful information appropriate to context so people can understand when they are engaging with AI and can challenge outcomes.

Office of the Australian Information Commissioner. (2024). *Guidance on privacy and the use of commercially available AI products*. OAIC. This source supports the paper's discussion of privacy obligations when organisations use AI products, including publicly available AI chatbots and productivity tools.

Office of the Australian Information Commissioner. (2024). *Guidance on privacy and developing and training generative AI models*. OAIC. This source supports the discussion of personal information, generative AI model training, and privacy-law obligations relating to collection, use, and disclosure of personal information.

Open Government Partnership. (n.d.). *Beneficial ownership*. This source supports the paper's broader public-reliance argument around transparency of control, conflicts of interest, money laundering risk, public contracting, and the need to identify real interests behind legal or institutional structures.

Open Government Partnership. (n.d.). *Anti-corruption: Lobbying*. This source supports the paper's discussion of lobbying transparency, public accountability, fairness of access, and the importance of helping the public understand how decisions are made.

Thaler v. Perlmutter, No. 23-5233, United States Court of Appeals for the District of Columbia Circuit. (2025). This court opinion supports the paper's discussion of human authorship and AI-generated works in copyright law.

UNESCO. (2021). *Recommendation on the Ethics of Artificial Intelligence*. UNESCO. This source supports the paper's discussion of transparency, explainability, auditability, traceability, oversight, impact assessment, human rights, and due diligence in AI governance.

U.S. Copyright Office. (2025). *Copyright and Artificial Intelligence, Part 2: Copyrightability*. This source supports the paper's discussion of AI-assisted authorship, human contribution, copyrightability, and the distinction between AI-generated material and human-authored contribution.

U.S. Copyright Office. (n.d.). *Copyright and Artificial Intelligence*. This source supports citation of the Copyright Office's broader multi-part AI report series, including Part 2 on copyrightability and generative AI.

World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*. WHO. This source supports the paper's discussion of high-reliance AI domains, health-sector governance, human oversight, transparency, accountability, and public benefit.

World Health Organization. (2025). *Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models*. WHO. This source supports the paper's discussion of generative AI and large multimodal models in healthcare, research, public health, and drug development contexts

AI Use Disclosure Statement

This paper was developed through human-directed AI assistance. The author supplied the originating thesis, conceptual framework, terminology, structure, argument direction, examples, refinements, and final editorial judgment. AI was used to assist with drafting, organisation, expansion, wording, internal consistency, and reference identification.

The core concepts, including the Trace Economy, Proof of Cognitive Work, public reliance, private cognition/public reliance distinction, authorship-provenance-influence trilogy, influence provenance, sponsored reasoning, longitudinal cognitive data, and the three-hashtag trace mechanism, were developed and directed by the author.

The author reviewed and approved the final manuscript and accepts responsibility for the claims, framing, argument, and conclusions. AI assistance was used as a drafting

and reasoning support tool, not as an independent author. No AI system is credited as an author.

Zenodo Description

This paper introduces the Trace Economy as a public-reliance protocol for integrating authorship, provenance, and influence into a unified accountability framework. It argues that artificial intelligence has intensified, but not created, a broader institutional problem: people are increasingly asked to rely on outputs, claims, recommendations, rankings, reports, decisions, public narratives, and institutional representations without sufficient visibility over who created them, where they came from, or what interests shaped them.

The paper develops a three-pillar framework: authorship identifies the contributor; provenance traces the pathway; and influence reveals the forces. It argues that existing transparency mechanisms — including sponsored-search labels, product-placement disclosure, research funding declarations, conflict-of-interest statements, chain-of-custody procedures, and AI-content labelling — remain fragmented because they do not provide a general public-reliance protocol across domains.

The paper introduces influence provenance as the missing third layer of trace. It defines sponsored reasoning as the embedding of commercial, behavioural, institutional, or political influence inside AI-generated advice, recommendations, summaries, rankings, or decision-support outputs. It also introduces longitudinal cognitive data as a distinct category of user-derived data produced through repeated AI interaction, with implications for privacy, advertising, manipulation, authorship, cognitive sovereignty, and regulatory design.

The Trace Economy is presented as a low-friction public anchoring mechanism. When a work, claim, recommendation, decision, or output crosses from private cognition into public reliance, it can be marked through the protocol hashtags #TraceEconomy, #PoCW, and #Unifaircation, together with direct tagging of the architect or relevant foundation. The paper emphasises that the trace tag does not replace evidence; it anchors the claim and creates a public doorway into a deeper evidence trail.

The framework is applied across AI governance, social media harm reduction, grants and research, political campaign transparency, professional accountability, policing, public-sector decision-making, commercial recommendations, media, insurance, and war-justifying public claims. Its central doctrine is: private cognition remains private, public reliance becomes traceable, and influence becomes visible.