Process Essentialism and the Epistemology of Academic Knowledge: Why How We Write Should Not Determine What We Know

Masaya Ochiai

Independent Researcher, Nagoya, Aichi, Japan
ORCID: 0009-0006-6069-419X
Corresponding author: masaya.ochiai.research@gmail.com

Recommended citation

Ochiai, Masaya. 2026. "Process Essentialism and the Epistemology of Academic Knowledge: Why How We Write Should Not Determine What We Know." *Synthese*, 207, 200. https://doi.org/10.1007/s11229-026-05580-3.

Document structure

This PDF contains this notice page, the accepted title page and declarations, and the accepted manuscript body. It does not contain the publisher typeset PDF / Version of Record.

Process Essentialism and the Epistemology of Academic Knowledge:

Why How We Write Should Not Determine What We Know

Masaya Ochiai
ORCID: 0009-0006-6069-419X
Independent Researcher, Nagoya, Aichi, Japan

Corresponding author: masaya.ochiai.research@gmail.com

Statements and Declarations

Use of AI tools: ChatGPT (OpenAI) was used for translation into English and language editing of content conceived by the author. Conceptual design, argument structure, and final source selection were based on the author's own judgment. All final wording was reviewed, revised, and approved by the author, who takes full responsibility for the manuscript.

Author contributions: The author conceived the study, developed the conceptual framework, and wrote the manuscript.

Ethical considerations: Not applicable. This is a conceptual paper; no human subjects research was conducted.

Consent to participate: Not applicable.

Consent for publication: Not applicable.

Competing interests: The author declares none.

Financial support: None.

Data availability statement: This article is conceptual; no dataset was created.

Process Essentialism and the Epistemology of Academic Knowledge:

Why How We Write Should Not Determine What We Know

Abstract

Large language models (LLMs) are increasingly used for drafting, translating, and editing academic prose. Recent debates and editorial policies often treat facts about how a text was produced as epistemically significant. This paper diagnoses the philosophical assumption behind such process-oriented norms: process essentialism, the view that the justificatory status of academic knowledge claims is constitutively dependent on their production process. I argue that strong forms of process essentialism are mistaken once we distinguish justificatory status from ethical questions of credit, accountability, and provenance. In publication contexts, what primarily matters is the quality of the public epistemic record—argumentative structure, evidential support, engagement with relevant work, and, where applicable, data, code, methods, and provenance claims—rather than the causal history of composition. I develop three objections: a genetic-fallacy argument, an argument from the unobservability and unenforceability of private process criteria, and an epistemic justice argument about exclusionary burdens. I conclude by showing how disclosure, contributorship, and provenance practices can be justified within an output-based framework without treating private process facts as constitutive of justification.

Keywords: process essentialism; evidentialism; reliabilism; peer review; large language models; epistemic justice

1. Introduction

The emergence of large language models (LLMs) such as ChatGPT, Claude, and Gemini has generated intense discussion within the academic community about the nature and boundaries of legitimate scholarly practice. Editorial policies have proliferated and continue to evolve: where publishers and journals have issued guidance, the overwhelming majority prohibit listing LLMs as authors while requiring human accountability (for example, among those with explicit guidance in Ganjavi et al.'s (2024) sample, 96% of top publishers and 98% of top journals prohibited listing generative AI as an author), yet they diverge on what uses must be disclosed and what uses are restricted (for example, generative images) (COPE, 2023; EDP Sciences, 2023; Ganjavi et al., 2024; International Committee of Medical Journal Editors [ICMJE], n.d.; Nature Portfolio, n.d.). Alongside these policies, a growing literature debates the permissibility, risks, and implications of AI-assisted academic writing (Stokel-

Walker & Van Noorden, 2023; Lund et al., 2023; Resnik & Hosseini, 2026). Relatedly, epistemologists have begun to analyse conversational AI systems as potential sources of testimony or epistemic authority, raising questions about reliance, responsibility, and trust (Wheeler, 2020; Freiman, 2024; He & Yang, 2025).

A striking feature of the current discourse is the extent to which concerns about LLM use in academia focus on process rather than output. The worry is not primarily that AI-assisted papers contain worse arguments, cite fewer sources, or reach less defensible conclusions. Rather, the anxiety centres on the fact that the manner of production has changed—that authors may not have struggled through each sentence themselves, that the cognitive labour of composition has been partially offloaded, that the traditional pathway from reading through thinking to writing has been altered or bypassed. This process-oriented anxiety, I suggest, reflects an implicit philosophical commitment that deserves explicit examination.

In this paper, I identify and critique what I term process essentialism: the view that the epistemic value or legitimacy of academic knowledge is constitutively determined by the process through which it was produced. Process essentialism holds that how a scholarly text came into being—what tools were used, what cognitive operations the author personally performed, in what sequence ideas were developed—is epistemically relevant to the evaluation of the knowledge claims that text contains. My central contention is that this view is philosophically untenable and that it conflates questions about the sociology of knowledge production with questions about the epistemology of justification.

Two clarifications about scope will be important throughout. First, my primary target is the justificatory status of public academic knowledge claims as presented for publication and communal uptake. This is compatible with denying that the named author thereby personally knows the claim, deserves full epistemic credit for it, or satisfies every ethical condition of authorship. It is also compatible with treating later replies to criticism, revision letters, and corrigenda as part of the broader evaluative record where these materials are available for scrutiny. Second, although my central examples concern conceptually oriented scholarship whose primary evidential currency is argument and citation, the output-based framework is not limited to the humanities. In empirically oriented fields, the relevant ‘output’ normally includes data, code, materials, explicit provenance claims, contributor statements, and other publicly inspectable artifacts; treating those materials as part of the evaluative record preserves the output/process distinction while making room for reproducibility and accountability concerns.

The argument proceeds as follows. In Section 2, I define process essentialism more precisely and identify its manifestations in contemporary academic discourse. In Section 3, I situate the debate within the reliabilism–evidentialism controversy in analytic epistemology, arguing that the epistemological structure of academic evaluation is fundamentally evidentialist even though process-essentialist norms implicitly presuppose a reliabilist framework. Section 4 advances three distinct arguments against process essentialism: epistemological, practical, and justice-based. Section 5 reinterprets peer review as an output-based quality assurance mechanism. Section 6 addresses objections, including concerns about fabrication and the role of transparency. Section 7 concludes.

This diagnosis is intended as a contribution to social epistemology and the epistemology of science. It is not offered as a survey of editorial policies or as a proposal for a single best policy package. Rather, the aim is to (i) introduce a diagnostic category—process essentialism—for a family of process-oriented epistemic norms; (ii) clarify the relevant output/process distinction for academic evaluation; (iii) develop three objections to strong forms of process essentialism (genetic-fallacy, unobservability/enforceability, and epistemic justice); and (iv) show how disclosure and integrity policies can be justified within an output-based framework without treating private tool-use facts as epistemic defeaters.

2. Process Essentialism: Defining the Target

By process essentialism, I mean the following thesis:

(PE) The epistemic value of an academic knowledge claim is constitutively dependent on the process by which it was produced. More precisely: differences in the production process can make a difference to a claim's justificatory status even when the relevant public epistemic record is held fixed.

The constitutive thesis at issue is best formulated as a grounding claim (cf. Fine, 2012; Rosen, 2010). The question is whether private process facts belong to the set of facts in virtue of which publication-level justificatory status obtains. Supervenience matters here only as a diagnostic constraint on difference; it does not itself supply the relevant explanation. On the reading of PE-S and PE-M adopted here, private compositional history remains part of the grounding base for justificatory status even when the relevant public epistemic record is held fixed. That is the claim rejected in this paper. Once the public epistemic record is fixed—including the reasons offered, evidential support, explicit method or provenance claims, and, where available for scrutiny, the dialectical record—no further justificatory role remains for private compositional history. Questions about realization or metaphysical individuation may be important in other contexts, but they are not required for the present argument.

Table 1. Distinguishing constitutive dependence from other ways process can matter

Relation	Characterization	If the relevant public epistemic record is held fixed…
Constitutive dependence	Process facts are part of what grounds justificatory status (a grounding claim, diagnosable by a supervenience-style test).	…they can still differ in epistemic standing solely because the production process differed.
Causal/instrumental relevance	Process affects output quality, but only via its tendency to produce better or worse texts.	…there is no epistemic difference once output quality is held fixed.
Defeasible evidential relevance	Process facts provide defeaters or risk-markers (e.g., reasons to double-check), without being constitutive.	…they matter only insofar as they indicate likely errors or missing evidential work.
Ethical/professional relevance	Process matters for credit, accountability, and integrity norms, even if not for justification.	…ethical evaluation may differ even when epistemic evaluation does not.

Table 1 distinguishes constitutive dependence from other ways in which process can matter (causally/instrumentally, defeasibly as evidence, or ethically/professionally). This framing allows us to distinguish stronger and weaker versions of process essentialism. Strong versions claim that process facts are necessary conditions on justification, while weaker versions treat process information as defeasible evidence or as ethically relevant without being epistemically constitutive.

(PE-S) Necessary-process thesis: a claim is epistemically justified-as-academic-knowledge only if it was produced via an approved process (e.g., without certain forms of AI-generated drafting), even when the relevant public epistemic record is held fixed.

(PE-M) Process-sensitive supervenience thesis: two papers could make the same public case, cite the same evidence, and be equally defensible under scrutiny, yet differ in justificatory status solely because one author arrived at that record through the 'wrong' process.

(PE-W) Process-as-evidence thesis: facts about production are not constitutive but provide defeasible indicators (or defeaters) of likely error, missing work, or lack of accountability; they matter only insofar as they bear on expected output quality.

2.1 Output, Process, and Epistemic Control

The central contrast in this paper is best stated in terms of the public epistemic record rather than in terms of the author's private cognitive history. By 'output' I mean the publicly inspectable package that is presented to a community for assessment and uptake: not only the submitted manuscript, but also whatever supporting materials (data, code, appendices, methodological details, proofs, registered reports, contributor statements, disclosure statements) and responsive materials—where available for scrutiny by relevant evaluators—such as revision letters, point-by-point replies, clarifications, and corrigenda. By ‘process’ I mean the causal history of production that is not itself part of this broader evaluative record—e.g., the order in which the author read sources, drafted paragraphs, or interacted with particular tools—except insofar as the author makes such facts part of the record through explicit claims. Table 2 summarises the working definitions that will be assumed in the remainder of the paper.

Table 2. Working definitions: output, process, and epistemic control

Term	**Working definition (in this paper)**	**Evaluation-relevant upshot**
Output (O)	The publicly inspectable record presented for assessment: the manuscript plus cited reasons/evidence, explicit method or provenance claims, and—where relevant—data, code, materials, registered reports, contribution statements, appendices, and corrigenda.	Output-based evaluation inspects this record and its evidential relations rather than reconstructing private drafting history.
Process (P)	The non-public causal history of production (order of reading, drafting, tool use, etc.), except where such facts are themselves asserted in O.	P can matter instrumentally, ethically, or as defeasible evidence (PE-W), but it is not a constitutive ingredient of justification in PE-S/PE-M’s sense.
Dialectical record	Responses to criticism available for scrutiny during review or after publication: author replies to reviewer comments, revision-stage clarifications, Q&A answers, follow-up clarifications, and post-publication corrections.	Where available for scrutiny, this is part of O. It is a primary way peer review tests understanding, accountability, and the ability to revise in light of reasons.
Epistemic control	The agent’s capacity to take responsibility for the claim by tracking reasons: explaining, defending, revising, and supplying further support when challenged.	Typically assessed through output-level indicators (especially the dialectical record), not by access to the author’s private cognitive steps.
Method/provenance claims	Assertions within O about how evidence was obtained (e.g., that an experiment was run, data were observed, or analyses were executed as described).	If false, O contains a falsity and the purported evidential link fails; auditing such claims targets the truth of output-level assertions, not a freestanding ‘process legitimacy.’

This clarification matters for several objections. For example, peer review commonly includes requests for clarification and revision, and it often tests whether authors can defend their claims in response to criticism. On the present usage, these dialectical interactions are part of the output record to the extent that they are available for scrutiny. Output-based evaluation therefore does not deny that an author's ability to answer objections can be relevant; it denies that justificatory status depends on private provenance facts over and above what is available for scrutiny.

The same point applies to process information that is deliberately externalised. Registered reports, author-contribution statements, analysis notebooks, prompt logs, and similar materials cease to be merely private process facts once they are submitted for public scrutiny. On the present view, such artifacts are legitimate objects of evaluation precisely because they have become part of the public epistemic record. This is one reason why open-methods and provenance practices are often defended in fields concerned about replicability or a replication crisis: once inferential steps and supporting materials are externalised, they can be rerun, challenged, and independently checked.

Process essentialism is not a position that any single author has explicitly endorsed under this name. I use it as a diagnostic label for a family of commitments that are implicit in a range of contemporary academic norms and policy proposals. To keep the target charitable and precise, my main focus is on the stronger theses (PE-S) and (PE-M) above. The weaker thesis (PE-W)—on which process facts function as defeasible indicators of error or risk—can be accommodated within an output-based framework and is not what my arguments aim to reject. Consider the following examples of process-essentialist (or PE-leaning) reasoning, all of which are commonplace in current debates about LLMs in academia.

These norms are not obviously irrational. They are often motivated by reasonable concerns about reliability, accountability, credit, and the social role of scholarly communication. The question is whether such concerns require treating production histories as epistemically constitutive, or whether they are better captured as instrumental considerations, defeasible risk-markers, or ethical/professional constraints that can coexist with output-based epistemic evaluation.

First, there is the norm that scholars should engage seriously with relevant literature in the course of formulating and presenting their arguments. That norm may often be instrumentally valuable and reliability-enhancing. On a PE-style reading, however, a paper whose author first developed a thesis independently and only later integrated it with the literature would be epistemically deficient compared to one whose argument emerged through continuous prior engagement—even if the two papers are identical in their final argumentative quality. My target is not the reasonable claim that literature engagement is often instrumentally valuable, but the stronger claim that a first-principles or literature-later route would by itself make an otherwise equivalent paper epistemically inferior.

Second, there is the norm that authors should compose their texts ‘in their own words’—a requirement that extends beyond the prohibition on plagiarism to encompass a broader expectation about the cognitive provenance of the text. When this norm is applied to LLM use, the concern is not that the resulting prose is of lower quality or that the arguments are less sound, but that the labour of composition was not performed by the named author in the expected way.

Third, there is the expectation that scholars should write in the language of publication without AI translation assistance, or at least that such assistance should be viewed with suspicion. This norm implicitly ties epistemic standing to the linguistic process of composition: a paper originally drafted in English is treated as epistemically distinct from an identical paper translated from Japanese, even when the translated version is indistinguishable in its final form.

Fourth, and more broadly, there is the widespread assumption that certain generative AI tools used for drafting, summarising, or translation bear on the epistemic status of the resulting work—that such use requires disclosure not merely for transparency’s sake, but because it changes how the paper should be epistemically evaluated. Ordinary spell-checking, proofreading, or grammar-support tools are not usually treated this way; the live controversy concerns tools that can reshape substantive prose, assist translation, locate literature, suggest arguments, or generate analytic artifacts.

What unites these examples is a common structure: in each case, the evaluative focus shifts from the content of the scholarly output to the process of its production. The question ‘Is this a good argument?’ is supplemented or replaced by the question ‘Was this argument produced in the right way?’ Process essentialism thus represents a particular answer to the fundamental epistemological question of what makes academic knowledge claims justified: they are justified, at least in part, by virtue of having been produced through certain approved procedures.

It is important to distinguish strong process essentialism (PE-S/PE-M) from three related but distinct positions. First is the reasonable view that certain processes are instrumentally valuable because they tend to produce better outputs. One might hold, for instance, that engaging deeply with the literature before writing typically leads to better arguments, without thereby claiming that the process of engagement is constitutive of the argument's epistemic status. Second is the view that transparency about methods is intrinsically valuable as a professional norm. One can endorse disclosure of AI tool use while maintaining that the disclosed information is not itself an epistemic defeater. Third is the weaker, evidential thesis (PE-W): knowledge of production methods can function as defeasible evidence or as a risk-marker that sometimes rationally triggers extra checking. I return to these distinctions in Section 6.

Two further clarifications will help keep the target in view. First, throughout this paper my target is the narrower epistemic claim that private tool-use facts (or other private provenance facts), over and above the publicly inspectable evidential record, are constitutive of justificatory status (PE-S/PE-M). This is distinct from (and compatible with) ethical and professional process norms concerning credit, responsibility, conflicts of interest, or research integrity. Second, the output/process distinction is not meant to deny that some provenance facts can matter epistemically. They can, but typically only when they are themselves asserted within the output record (e.g., in methods or provenance statements) or when they bear defeasibly on the trustworthiness of that record; in such cases they function as part of the evidence or as risk-markers (PE-W), not as free-standing constitutive conditions.

3. The Epistemological Question: Process or Output?

Process essentialism raises a question that epistemologists have long debated in other contexts: what is the relationship between the justification of a belief and the process by which that belief was formed? The classical debate between reliabilism and evidentialism provides the theoretical framework within which process essentialism can be most productively analysed.

3.1 Reliabilism and the Primacy of Process

Process reliabilism, as developed by Alvin Goldman (1979, 1986), holds that a belief is justified if and only if it is produced by a reliable cognitive process—one that tends to produce a high ratio of true beliefs to false ones. On this account, the justificatory status of a belief is constitutively linked to the causal process that generated it. However, contemporary reliabilism is more heterogeneous than the classical picture suggests. Virtue reliabilists emphasise competences rather than brute processes (Sosa, 2007; Greco, 2010), and Goldman's own later work develops a social epistemology that evaluates institutions—testimony, peer review, and other filtering mechanisms—in terms of their reliability

(Goldman, 1999). This diversity matters because not every reliabilist view supports the strong process-essentialist claim that production history is itself an epistemic condition on the justificatory status of a publicly available text.

Certain transpositions of process reliabilism from individual belief formation to academic knowledge production yield something close to strong process essentialism. On this transposition, an academic claim is justified only if it was produced through a reliable scholarly process—one involving certain approved methods of inquiry, composition, and review—and this reliability is treated as constitutive of epistemic status rather than merely instrumental. Still, it is important not to identify reliabilism with process essentialism too quickly. Even a reliabilist can regard publicly inspectable reasons, evidential support, and critical uptake as central because they are precisely the social mechanisms through which a community filters for reliable outputs (Goldman, 1999). What is distinctive about PE-S/PE-M is the further claim that a text can be epistemically disqualified by its provenance even when it is equally defensible on the basis of the reasons it presents.

Several features of the current LLM debate reflect this reliabilist logic. The concern that AI-assisted writing may be unreliable—that LLMs hallucinate, that they may introduce subtle errors, that the process of using them is not well understood—parallels the reliabilist worry about belief-forming processes whose track record is uncertain. If we cannot verify that the process is reliable, the reliabilist reasons, we cannot trust the output.

It is important to note, however, that the reliabilist transposition faces difficulties that are specific to the academic context. Goldman's original reliabilism concerns the justification of individual beliefs held by individual agents. Academic knowledge claims, by contrast, are public propositions embedded in argumentative structures and evidential networks. The reliability of the process by which a belief was formed may be relevant to the individual agent's epistemic situation, but once a knowledge claim has been publicly articulated and embedded in an argumentative structure, it becomes available for evaluation on its own terms—on the basis of the reasons and evidence adduced in its support. The transition from private belief to public claim introduces a new evaluative context in which the process of formation recedes in relevance and the quality of the articulated argument comes to the fore.

3.2 Evidentialism and the Primacy of Output

Evidentialism, as articulated by Earl Conee and Richard Feldman (2004), offers a starkly different picture. On the evidentialist account, a belief is justified if and only if it fits the evidence available to the believer. The process by which the belief was formed is irrelevant to its justificatory status; what matters is the current evidential relation between the belief and the body of evidence. A belief that fits the evidence is justified regardless of its causal history, and a belief that does not fit the evidence is unjustified regardless of the reliability of the process that produced it.

Transposed to the academic context, evidentialism supports output-based evaluation. An academic knowledge claim is justified insofar as the published text provides adequate evidence for its conclusions: the arguments are valid, the evidence is cogent, the engagement with existing literature is thorough, and the reasoning is transparent. Whether the author used a typewriter or a word processor, whether they drafted in their native language or the language of publication, whether they used AI tools for translation or composition—these facts about the productive process are, on the evidentialist view, epistemically irrelevant to the justificatory status of the claims.

3.3 The Institutional Twist: Peer Review, Policy, and Mixed Norms

Here I want to draw attention to what I regard as a significant tension in contemporary academic practice. The institutional mechanism that academia has developed for evaluating knowledge claims—prepublication peer review—asks reviewers to assess a submitted manuscript's arguments, evidential support, originality, and engagement with relevant work (Heesen & Bright, 2021). Reviewers are not ordinarily tasked with reconstructing an author's private drafting history—for example, verifying the order in which sources were read or confirming that each sentence was composed without mechanical assistance. In this sense, peer review is largely an output-based evaluation.

An immediate worry is that the reliabilism–evidentialism debate concerns the justification of an individual's beliefs, whereas publication decisions concern an institutional status: whether a community should treat a publicly articulated claim as eligible for uptake in a shared body of knowledge. This worry is well taken. My use of the reliabilism–evidentialism frame is therefore not a claim of strict identity between the two contexts, but a way of isolating a structural choice. Institutions can either (i) make justificatory status depend on publicly inspectable reasons and evidence, or (ii) make it depend on private provenance facts over and above that public record. Many desiderata that genuinely matter in institutional settings—reproducibility, accountability, and risk-management—are best handled by expanding and policing the output record (data, code, materials, provenance claims, and the dialectical record), rather than by treating the author's private composition history as constitutive.

Yet the discourse and policy response around LLMs increasingly introduces process-oriented criteria that sit uneasily with this output-based structure. Many journals now require authors to disclose certain uses of generative AI, prohibit listing AI systems as authors, and emphasise human responsibility for accuracy and integrity (COPE, 2023; ICMJE, n.d.; Nature Portfolio, n.d.). At the same time, policies vary widely in what they treat as disclosable or impermissible, and they continue to shift as the technology and its uptake evolve (Ganjavi et al., 2024). Whatever their practical merits, these interventions import process-based considerations into an institution that was designed to evaluate outputs. The result is an epistemological hybrid: a system that officially evaluates knowledge on evidentialist grounds while increasingly regulating the processes by which that knowledge is produced.

Importantly, not every public process-related requirement is process essentialist. Registered reports, author-contribution statements, open methods, and provenance disclosures are often compatible with an output-first framework precisely because they convert accountability-relevant facts into publicly inspectable artifacts. What matters, on the present view, is that these practices are justified as ways of enlarging the public epistemic record—not as evidence that private drafting history is itself constitutive of justificatory status (Chambers, 2013; Nosek & Lakens, 2014; Allen et al., 2014).

This distinction also helps with institutional-status worries. Journals can legitimately impose governance conditions such as responsible authorship, traceability, or bans on ghost and generative co-authors as admission rules for the literature. But such institutional facts do not by themselves transform a weakly supported claim into a justified one. They regulate eligibility, accountability, and credit; they do not supply evidential warrant in the absence of adequate publicly inspectable reasons (Hosseini et al., 2023; Hosseini et al., 2024).

The point can be sharpened by borrowing Searle's (1995) notion of a status function. In an institutional setting, publication eligibility may count as a socially conferred epistemic standing rather than a merely administrative one: acceptance into the literature makes a paper citable, archivable, and presumptively available for communal uptake. In that sense, journals do not merely sort manuscripts; they assign an institutional status within a practice of shared inquiry.

But admissibility and justification still come apart. A manuscript can satisfy all admission rules and yet remain poorly supported, and a manuscript can lose institutional standing or admissibility without thereby altering the strength of the public reasons it presents. This distinction is conceptual, not permissive—it does not recommend rule violation. Institutions may legitimately govern accountability, traceability, and authorship, and those governance conditions may well be epistemically important in a broad social sense. But the evidential relations that make a claim well supported are not created by institutional declaration alone. They depend on the publicly inspectable reasons, evidence, and inferential structure the manuscript makes available.

This creates a tension. If the evidentialist structure of peer review is epistemically appropriate—if the right way to evaluate academic knowledge is by examining its content—then importing reliabilist process criteria risks confusing the question ‘Is this claim well supported?’ with the question ‘Was this claim produced in the right way?’ If, on the other hand, the reliabilist framework is correct and process does matter constitutively, then peer review as currently practised would be incomplete for that stronger task, and additional mechanisms for evaluating scholarly work would be needed. The current discourse has not sufficiently grappled with this dilemma.

3.4 Likelihoodism and the Discovery/Justification Distinction

Elliott Sober’s (2008) work on the logic of evidence is instructive here. Sober argues that evidence for a hypothesis is constituted by the probabilistic relationship between the evidence and the hypothesis—specifically, the likelihood that the evidence would obtain given the truth of the hypothesis. On Sober’s likelihoodist framework, the question of whether observation O supports hypothesis H over hypothesis H′ is settled by comparing the likelihoods: $P(O|H)$ versus $P(O|H')$. Crucially, once O is fixed as the evidential proposition under consideration, this comparison is a function of the content of O, H, and H′ alone. Facts about how O was generated or recorded can matter upstream insofar as they affect whether O is trustworthy or correctly specified, but those concerns are normally handled through output-level materials (e.g., methods or provenance claims, data, replication, and the dialectical record), not by treating private composition history as constitutive of justification.

Applied to academic evaluation, Sober’s framework yields a powerful argument against process essentialism. The evidential support a paper provides for its conclusions is a function of the logical and inferential relationships between the paper’s premises and its claims—the internal structure of the argument. Whether those premises were articulated by an author typing at a keyboard, dictating to a transcription service, or refining a draft with the assistance of a language model does not alter the logical relationship between premises and conclusions. The validity of a deductive argument is a formal property; the strength of an inductive argument is a probabilistic property. Neither is affected by the biographical details of the argument’s composition.

This point can be sharpened by considering the distinction between the context of discovery and the context of justification—a distinction with a long pedigree in the philosophy of science (Reichenbach, 1938). The context of discovery concerns how a hypothesis, argument, or idea was generated; the context of justification concerns whether that hypothesis, argument, or idea is epistemically warranted given the available reasons and evidence. The distinction has also been widely criticised in post-Kuhnian philosophy of science, in part because practices of discovery can shape what later counts as evidence, salience, and even acceptable standards of justification (Hanson, 1958; Kuhn, 1962; Longino, 1990). Even granting these critiques, however, a modest analytical point remains: once a claim is publicly articulated, its justificatory standing is primarily a function of the reasons that can be offered for it, not of its biographical origin story. Process essentialism (in its strong forms) denies even this modest point by insisting that facts about how a paper was written are constitutive of whether its claims are warranted.

4. Why Process Essentialism Fails

Building on the evidentialist framework developed in Section 3, I now advance three arguments against process essentialism. Each targets the thesis from a different angle: the first is epistemological, the second practical, and the third concerns epistemic justice. Together, they constitute a cumulative case for rejecting PE and adopting output-based evaluation of academic knowledge.

4.1 The Epistemological Argument: The Genetic Fallacy and Justificatory Independence

The most fundamental objection to process essentialism is that it commits a version of the genetic fallacy: it confuses the origin of a claim with the grounds for its justification. The genetic fallacy, in its classical formulation, is the error of evaluating a proposition based on its source rather than its content. While the genetic fallacy is not always fallacious—there are contexts in which source-sensitivity is epistemically appropriate (for example, in arguments from expert opinion) (Walton, 1997)—I argue that academic knowledge evaluation is not such a context, or at least not in the way process essentialism supposes.

One way to avoid overstatement is to distinguish source-insensitive from source-sensitive epistemic contexts. In source-insensitive contexts—logical validity, argumentative coherence, adequacy of citation support, or the fit between stated evidence and conclusion—what matters is the publicly available case for the claim, not who first produced the sentence or in what order the author drafted it. In source-sensitive contexts, by contrast, origin can matter because the claim itself depends on the reliability of a source, method, or performance: testimony about what an experiment showed, expert judgment, skilled interpretation, or provenance claims about how data, code, or images were generated. Appeals to origin are legitimate there because origin is part of the evidential content or because source reliability is itself under evaluation. My argument targets the illicit migration of source-sensitive reasoning into source-insensitive cases.

Consider the following thought experiment. Two philosophers independently produce papers with identical arguments, identical citations, identical conclusions, and equally competent responses to reviewer objections. Philosopher A composed her paper through the traditional process: years of reading, seminar discussions, handwritten notes, and laborious drafting. Philosopher B arrived at the same arguments through an unconventional process: she formulated her thesis first, then used AI tools to help identify relevant literature, and employed translation software to render her arguments in English. If process essentialism is correct, these two papers differ in their epistemic status despite being identical in content and dialectical performance. Philosopher A's paper is better justified because it was produced through the 'right' process.

This implication is deeply counterintuitive. If the arguments in both papers are equally valid, the evidence equally cogent, the engagement with existing scholarship equally thorough, then what epistemic deficiency could Philosopher B's paper possibly suffer from? The process-essentialist must claim that there is some epistemic property—call it processual legitimacy—that is present in A's paper and absent in B's, and that this property is relevant to the paper's justificatory status even though it leaves no trace in the paper's content. But this is precisely what the evidentialist denies, and rightly so. The evidential relations that constitute justification are relations between propositions, not between propositions and the biographical histories of their authors.

Jennifer Lackey's (2008) work on testimonial knowledge reinforces this point. Lackey argues that knowledge acquired through testimony—that is, knowledge obtained from the assertions of others—can be fully justified even though the knower did not personally generate the evidence. What matters is that the testimony is reliable and that the hearer has no defeaters. Analogously, a scholarly argument that is well supported by evidence and free of errors is epistemically justified regardless of whether the author personally performed every cognitive operation involved in its production. The scholarly text, like testimony, provides the evidential basis for its claims; the question of who or what performed the cognitive labour is a question about production, not about justification.

Sanford Goldberg’s (2010) analysis of epistemic dependence extends this reasoning further. Goldberg demonstrates that we routinely and legitimately rely on epistemic contributions from others—colleagues, research assistants, editors, translators—without thereby undermining the epistemic status of our knowledge claims. The question is not whether reliance on others occurred, but whether the resulting epistemic product meets the relevant standards. This does not erase the ethical difference between human and non-human contributors: human assistants can deserve acknowledgment, authorship, or other forms of responsibility in ways AI systems cannot. But that difference concerns credit and accountability, not the justificatory status of an otherwise well-supported claim. If reliance on human collaborators does not vitiate epistemic standing, it is unclear why reliance on computational tools should do so, provided the output satisfies the same evidential criteria.

That difference becomes clearer in Sosa’s terms. Aptness is success because of competence. Competence belongs to a subject who can exercise intellectual abilities, answer for how those abilities are exercised, and appropriately receive credit or blame for the result. Human research assistants can, though need not, occupy that role. They can understand reasons, exercise judgement, answer for interpretive choices, and may deserve acknowledgment, authorship, or blame, depending on what they contributed. Current AI systems, by contrast, may extend, scaffold, or partially automate a scholar’s performance, but they are not, on the present framework, the relevant seat of competence or answerability. For that reason, AI assistance may matter to attribution and credit, but it does not introduce a second competence-bearing contributor in the way a human collaborator can.

An assistant may instead be used in a purely constrained and protocol-bound way—for example, to execute a prespecified Boolean search, populate a summary template, or run a routine statistical procedure under detailed instructions. In such limiting cases, the assistant functions less as a second epistemic agent and more as an instrumentally deployed executor of an externally specified workflow. The contrast with a similarly constrained AI system therefore narrows considerably. This does not undermine the present argument; it sharpens it. The relevant question is not whether the executor is biologically human, but whether the delegated contribution is functioning as a competence-bearing, answerability-involving part of the inquiry or as a mechanised subroutine whose epistemic significance is exhausted by the quality, fidelity, and auditability of the resulting output. Where the task is genuinely mechanical, a protocol-bound human assistant and a similarly constrained AI system may be epistemically on a par at the level of publication-level justification, even though labour, credit, and institutional norms may still distinguish them.

Still, this does not vindicate PE-S or PE-M. One might argue that AI-assisted scholarship can be merely accurate without being apt—that the author's success is not because of her competence. This does raise a genuine issue. The crucial question is what counts as the relevant competence in scholarly contexts and what the 'because of' relation requires. If competence is understood narrowly as the ability to generate fluent prose unaided, then AI assistance threatens aptness. But that is not the competence at issue in scholarship. Scholarly competence centrally involves tracking reasons: selecting questions, interpreting sources, evaluating evidence, constructing and revising arguments, and answering objections. Tasks such as framing a new research question or integrating a large body of literature into a novel interpretive claim are therefore unlike merely executing a prespecified search or routine analysis: they are underdetermined and call for endorsement, defence, and revision, not mere faithful execution. These competences can be exercised through tools, provided the agent retains epistemic control—exhibited in the public record by checking claims, verifying citations, revising in light of counterevidence, and supplying further support when challenged. In peer review, that control is commonly tested through requests for clarification, revision, and point-by-point replies, which are part of the broader output record during review. Where an author abdicates these competences and merely submits text she cannot defend, aptness fails; but then the failure tends to manifest in output-level shortcomings (unsupported leaps, uncorrected errors, or inability to respond to critique). Virtue epistemology therefore motivates a distinction between competence-bypassing and competence-extending uses of AI. Aptness remains epistemically significant for evaluating agents, achievements, competence, and entitlement to credit, but it is not a constitutive condition of publication-level justification once the relevant public epistemic record is held fixed in the relevant sense.

PE-M is often motivated by cases in which we seem to distinguish 'lucky' success from epistemically responsible success. Imagine two physicians who deliver the same correct diagnosis: one infers it from the patient's history, tests, and background theory; the other guesses. Intuitively, the guess is epistemically worse. Two clarifications help. First, the primary target of that intuition is the agent (knowledge, competence, responsibility, credit), not the proposition-as-claim once the supporting reasons are publicly available. Second, the academic analogue of the guessing case is typically not a situation in which outputs are held fixed in the relevant sense. When a writer does not understand what she submits, this generally shows up in the dialectical record—poor replies to reviewers, inability to explain inferential steps, or failure to integrate the claim with adjacent results. If, counterfactually, two scholarly output packages were identical—including their arguments, supporting materials, and competent responses to criticism—then there would be no remaining epistemic difference for communal uptake. Any remaining difference would be ethical or credit-theoretic (who deserves authorship or can be held accountable), and should be addressed by integrity norms rather than by treating private provenance as constitutive of justification.

4.2 The Practical Argument: Unobservability and Unenforceability

The second argument against process essentialism is practical but has epistemic implications. Much of what I have called the private writing process is unobservable from the standpoint of evaluation. Reviewers, editors, and readers have access only to the submitted text; they cannot directly observe the sequence of drafting, tool interactions, or cognitive labour by which it was produced. This limitation does not straightforwardly disappear with increased monitoring. A finished text still underdetermines its productive history.

This unobservability has two important consequences. First, any evaluative criterion based on process is difficult to apply consistently. If two papers are textually identical but were produced through different processes, no reviewer can distinguish them. Process-based criteria therefore risk being unevenly enforced and poorly correlated with what is actually evaluated—the text itself. A standard that can be satisfied or violated without any detectable difference in the evaluated object is, at best, only indirectly connected to epistemic appraisal and can easily collapse into a mere social convention.

Second, the unobservability of process means that process-based evaluation necessarily relies on self-report. Authors must disclose their methods, and evaluators must trust these disclosures. But self-report about writing processes is notoriously unreliable even in good faith. Authors may not accurately recall the sequence in which they developed ideas, the extent to which they relied on various tools, or the precise contribution of different cognitive operations to the final text. Requiring honest and accurate self-report about inherently fuzzy cognitive processes creates a system that rewards confidence in describing one’s own methods rather than the quality of one’s arguments.

The contrast with output-based evaluation is instructive. The content of a paper—its arguments, evidence, citations, and reasoning—is fully observable. Reviewers can assess the validity of arguments, check the accuracy of citations, evaluate the cogency of evidence, and judge the quality of engagement with prior work. These are publicly accessible features of the text. An evaluative system based on such features can be consistently applied, publicly justified, and intersubjectively verified. Process-based evaluation lacks all of these virtues.

One might object that the unobservability of process is precisely why disclosure requirements are necessary—that we should make the process observable through mandatory reporting. But this objection misses the deeper point. Even if we could make some process facts more transparent through disclosure regimes, the question would remain: why should process matter? The practical argument is not merely that process-based evaluation is difficult to implement; it is that the difficulty of implementation reveals something about the nature of the evaluative question. The fact that academic evaluation has historically proceeded without process information—and has produced a broadly functional system of knowledge validation—suggests that private composition history is not, and has never been, the relevant object of evaluation.

To be sure, some process facts can be converted into output artifacts: usage logs, prompt sheets, contribution matrices, lab notebooks, or analysis notebooks. Once externalised for scrutiny, these are best understood as additions to the output record rather than as windows onto some independently justificatory private process. This can be useful, especially where uncritical outsourcing of inferential steps to AI tools would otherwise aggravate replicability problems in fields already concerned about a replication crisis. But it also raises trade-offs of privacy, strategic self-presentation, bureaucratic burden, and feasibility. An output-first view therefore supports risk-proportionate requests for such artifacts when they bear on evidence or accountability, not routine surveillance of every stage of composition (Allen et al., 2014).

There is a deeper epistemological point here that deserves elaboration. The unobservability of writing processes is not a mere practical inconvenience; it reflects a structural feature of propositional knowledge. Propositions stand in logical and evidential relations to one another in virtue of their content, not in virtue of their causal history. The proposition ‘All valid arguments preserve truth’ bears the same logical relations to other propositions regardless of whether it was formulated by a logician, a student, or a language model. This content-determination of evidential relations is precisely what makes output-based evaluation possible and appropriate. Process essentialism, by insisting that the causal history of a proposition’s inscription is epistemically relevant, risks mislocating or obscuring the grounds of evidential appraisal.

4.3 The Epistemic Justice Argument: Processual Injustice and Linguistic Inequality

The third argument against process essentialism concerns epistemic justice. Miranda Fricker's (2007) framework of epistemic injustice—especially testimonial and hermeneutical injustice—together with later work on structural epistemic oppression (Medina, 2013), helps to explain how process-based norms can function as mechanisms of exclusion. I use the term 'processual injustice' for a pattern in which scholars suffer credibility deficits or devaluation not because their reasons are weak, but because their work is produced through a dispreferred process. Its characteristic manifestations include diminished uptake and the displacement of assessment of public reasons by process cues.

Process-essentialist norms, whatever their intention, disproportionately burden scholars who do not fit the traditional model of the well-resourced, anglophone university-based researcher. Consider the expectation that scholars should write in English without translation assistance. This norm imposes little burden on native English speakers but creates systematic additional costs for scholars whose first language is not English. Empirical work suggests that language barriers increase time costs and increase the probability of rejection or revision for reasons that are orthogonal to scientific or argumentative merit (Amano et al., 2023; Ramírez-Castañeda, 2020).

For example, Amano et al. (2023) survey 908 researchers in environmental sciences and report that, among participants who had published only one English-language paper, 38.1% and 35.9% of non-native English speakers (moderate and low English-proficiency nationalities, respectively) had experienced rejection due to English writing, compared to 14.4% of native English speakers—roughly a 2.5–2.6× difference. Ramírez-Castañeda (2020) similarly documents substantial financial costs for translation/editing services (between one-quarter and one-half of a doctoral monthly salary in Colombia) and reports high rates of language-driven rejection or revision. These patterns show that 'write unaided in English' is not a neutral process ideal but a distributive burden.

Moreover, once process policing is implemented via automated detection tools, the burdens can worsen. Liang et al. (2023) argue that GPT detectors frequently misclassify non-native English writing as AI-generated, raising concerns about fairness and robustness. If journals or institutions treat AI suspicion as epistemically disqualifying, they risk systematically discounting precisely those scholars already disadvantaged by linguistic norms.

More broadly, process-essentialist norms tend to privilege researchers with access to the institutional resources that facilitate the 'correct' process: research libraries for extensive literature engagement, seminars and colloquia for intellectual development, editorial assistance for polishing prose, and time for the leisurely unfolding of ideas that traditional scholarly process presupposes. Independent researchers, scholars in resource-poor institutions, and those with caregiving or disability-related constraints are all disadvantaged by process-based criteria—not because their work is of lower quality, but because their process does not conform to an idealised model. The same point applies to interdisciplinary and translation-mediated research, where epistemic labour is often distributed across collaborators, translators, and tools.

How, then, does processual injustice relate to Fricker's taxonomy? It can manifest as testimonial injustice when process cues (e.g., 'translated,' 'AI-assisted') function as proxies for identity or competence and lead evaluators to assign an unwarranted credibility deficit. It can also manifest as hermeneutical injustice when dominant interpretive resources render translation-mediated or tool-assisted production unintelligible as legitimate scholarly practice, producing a form of willful hermeneutical ignorance (Pohlhaus Jr., 2012). More broadly, it functions as a structural mechanism of epistemic oppression in Medina's sense when a historically contingent model of scholarly production is treated as normatively superior without adequate justification (Medina, 2013). The result is epistemic marginalisation based not on the quality of one's knowledge claims but on the manner of their production.

I want to be clear about the scope of this argument. I am not claiming that concerns about AI use in academia are motivated by a desire to exclude marginalised scholars. The argument is structural, not motivational: whatever the intentions behind process-based norms, their effects can be exclusionary. Moreover, some process interventions—especially disclosure requirements and integrity safeguards—may be justifiable on instrumental or ethical grounds (Section 6). The target here is the stronger, disqualificatory idea that deviations from an idealised production process are epistemic defects in themselves (PE-S/PE-M). Given the distributive burdens documented above, such disqualification would require especially strong epistemic justification. But, as argued in Sections 4.1 and 4.2, the relevant justificatory work is done by the publicly available reasons and by output-based scrutiny, not by provenance.

These distributive considerations are not limited to the present organisation of universities. If scholarly activity becomes more decoupled from conventional academic employment, output-first evaluation becomes, if anything, more important, because it allows independent or non-traditionally situated researchers to be assessed by the same public standards. The scenario is therefore not merely a boundary condition but a test case for the justice argument. Under PE-S or PE-M, two comparably strong contributions could be assigned different epistemic standing solely because one was produced outside the resource-rich pathways that conventional academic employment makes easier to access. That would show that the exclusionary pressure at issue is structural rather than incidental: once the public epistemic record is comparably strong, attaching a further deficit to nonstandard process simply reimports social advantage at the level of epistemic appraisal. A related qualification concerns improvements in machine translation. Even if such tools continue to mitigate some linguistic burdens, they do not remove broader asymmetries of time, training, archival access, or review capital. Among comparably protective evaluative regimes, the less exclusionary one is therefore pro tanto preferable.

5. Peer Review as Output-Based Quality Assurance

Having argued that justificatory status tracks the public epistemic record rather than private provenance, I now turn to the institutional question of how such justification is screened in practice. If process essentialism is rejected, what mechanism ensures the quality of academic knowledge? The answer, I suggest, is the one that has been operative throughout the history of modern academia: peer review. But peer review needs to be understood correctly—not as an inquiry into private cognitive history, but as a mechanism for testing publicly available reasons.

The standard peer review process, whatever its imperfections, is designed to evaluate the content of scholarly work. Reviewers assess whether arguments are valid, whether evidence is adequate, whether the contribution is significant, and whether the engagement with existing literature is appropriate. Crucially, these are all features of the output—the submitted manuscript—not of the process by which it was produced. A reviewer who rejects a paper because its arguments are unsound has identified a genuine epistemic deficiency; a reviewer who rejects a paper because the author used AI translation has imposed a process-based criterion that is external to the paper's epistemic content.

Helen Longino's (1990) account of science as social knowledge is illuminating here. Longino argues that the objectivity of scientific knowledge is secured not by the individual cognitive processes of scientists but by the social practices of critical scrutiny—particularly, the practice of subjecting knowledge claims to peer criticism. What makes scientific knowledge objective, on Longino's view, is that it has survived a process of intersubjective evaluation. This evaluation is directed at the content of claims, not at the biographical details of their producers. Longino's framework thus supports an output-based understanding of academic quality assurance: knowledge claims are validated by their ability to withstand critical scrutiny, not by the pedigree of the processes that generated them.

Longino’s position is also sometimes taken to be more radical: social practices can shape not only which claims are accepted, but also what counts as admissible evidence and what standards of inference are appropriate in the first place. If that is right, then ‘evidence’ is not a purely pre-social category. Acknowledging this point does not support PE-S or PE-M. It shifts attention from the private causal history of an individual manuscript’s composition to the public, community-level practices that constitute standards of criticism and uptake. Those constitutive social processes—on Longino’s view—operate by subjecting publicly available outputs to transformative criticism. So even on the more radical reading, the relevant evaluative focus is the claim as it enters a shared critical space, not the author’s private writing process.

Philip Kitcher’s (1993) analysis of the social organisation of science makes a complementary point. Kitcher argues that the reliability of scientific knowledge production depends on the division of cognitive labour within the scientific community—on the fact that different researchers pursue different approaches and subject each other’s work to critical examination. This division of labour operates at the level of outputs: researchers evaluate each other’s published claims, not each other’s working habits. The reliability of the system as a whole does not require that each individual researcher follow a prescribed process; it requires that the community’s evaluative mechanisms are capable of identifying and filtering out inadequate work.

I do not wish to idealise peer review. Remco Heesen and Liam Kofi Bright (2021) argue, from an epistemic consequentialist perspective, that the epistemic benefits of prepublication peer review may be weaker than commonly assumed, and they emphasise familiar limitations of the practice, including inconsistency, bias, and conservatism. Lee et al. (2013) review the empirical literature on bias in peer review and highlight both evidence of bias and significant methodological limitations in studying it. These are serious problems that merit attention and reform. But they are problems with the implementation of output-based evaluation, not reasons to abandon the output-based framework in favour of process-based criteria. The appropriate response to the limitations of peer review is to improve the mechanism of output evaluation, not to supplement it with process-based requirements that, as I have argued, lack epistemic justification.

Indeed, the addition of process-based criteria is likely to exacerbate rather than ameliorate the known problems with peer review. If reviewers are asked to evaluate not only the content of a paper but also the process by which it was produced, this introduces additional dimensions of subjective judgement, additional opportunities for bias, and additional sources of inconsistency—all without any compensating epistemic benefit. A reviewer who is biased against AI-assisted work may reject a paper on process grounds that would be indefensible on content grounds. This is a cost with no offsetting gain.

It is worth noting that the history of academic publishing already contains numerous examples of process-independent innovation that were initially met with resistance but ultimately judged by their outputs. The adoption of word processors, the use of bibliographic management software, the employment of statistical computing packages, and the reliance on professional editing services all changed the process of scholarly production without thereby altering the epistemic criteria by which scholarly work is evaluated. In each case, the academic community eventually recognised that the relevant question was not whether the new tool was ‘properly’ scholarly, but whether the work produced with it met the established standards of argumentative rigour and evidential adequacy. The current debate about LLMs follows the same structural pattern, and there is good reason to expect the same resolution: where tool use is source-insensitive, outputs should be judged by their merits rather than downgraded because of the tools used.

6. Implications and Objections

Having reinterpreted peer review as output-based quality assurance, I now address the main objections that such a view invites. The rejection of process essentialism raises several important questions and invites several objections. In this section, I address the most significant of these.

6.1 The Fabrication Objection

Perhaps the most natural objection to the output-based view is the case of fabricated data. If process is epistemically irrelevant, the objector reasons, then we cannot distinguish between a paper based on genuine experimental results and one based on fabricated data—since the two papers might be identical in their textual content. This would seem to reduce the output-based view to absurdity.

The objection, however, partly trades on an overly thin understanding of ‘output.’ In empirically oriented disciplines, the evaluable output is not merely a narrative report. It includes (i) the paper’s explicit and implicit provenance claims (e.g., that certain numbers result from specified observations and analyses) and (ii) whatever supporting artifacts make those claims answerable to scrutiny (raw data, code, materials, preregistrations, and disclosure statements where applicable). Fabrication is therefore not a case in which process and output diverge while output is held fixed. It is a case in which the output record contains falsehoods or missing evidential links: the paper represents itself as backed by observational evidence when it is not. That is an epistemic defect in the claim-package, even if detecting it sometimes requires audits, requests for materials, replication, or other forms of critical scrutiny.

To see this more clearly, consider the strongest variant of the objection: a researcher fabricates numbers that, by coincidence, match what honest experimentation would have produced. Even here, the fabricated report does not itself provide observational evidence, because no observation occurred and the provenance claims in the output record are false. Any warrant for the central hypothesis would have to come from elsewhere—subsequent replication, converging lines of evidence, or independently accessible data. Importantly, this aligns with PE-W rather than PE-M. Provenance facts can operate as defeaters or risk-markers for reliance on a given paper, but they do so by bearing on whether the evidential relations the paper purports to instantiate really obtain; they do not show that two genuinely evidentially equivalent output records can differ in justificatory status merely because their private production histories differ.

This is a source-sensitive case in the sense introduced in Section 4.1. Where a paper claims that data were observed, an experiment was run, or an image records a phenomenon, the truth of those provenance claims is part of what makes the output evidentially apt. Fabricated data and synthetic evidence therefore matter not because private process is always constitutive of justification, but because the public record now includes claims whose content concerns a reality-tracking process; if those claims are false, the evidential relation fails (Resnik, 2014; Resnik et al., 2025).

John Greco’s (2010) virtue-theoretic epistemology remains useful here, but in a more limited way. Greco distinguishes between the epistemic status of a claim and questions of epistemic credit—whether the agent deserves credit for getting it right. Even if a hypothesis were later vindicated by independent evidence, the fabricator would not deserve epistemic credit. But we need not—and should not—conclude that fabrication leaves justificatory status untouched. Output-based evaluation can condemn fabrication because the paper’s evidential support is not what it purports to be and because it undermines testimonial and institutional entitlements on which scholarly communication depends.

6.2 The Disclosure Question

If process is epistemically irrelevant, why should authors disclose their use of AI tools at all? This question is pressing because many journals now require such disclosure, and the rejection of process essentialism might seem to undermine the rationale for these policies.

My response is that the value of disclosure can be justified on grounds that are independent of process essentialism. First, disclosure supports collective learning and risk management: it allows the scholarly community to track correlations between particular tools, workflows, and recurring error profiles (e.g., fabricated citations) over time. Second, disclosure underwrites accountability. Because AI systems cannot take responsibility for published claims, policies often require human authors to remain responsible for accuracy and to document how AI tools were used (Hosseini et al., 2023; Hosseini et al., 2024; ICMJE, n.d.; Nature Portfolio, n.d.). Third, disclosure can serve ethical norms of honesty and appropriate credit without implying that AI assistance is an epistemic defeater (Resnik & Hosseini, 2026; Cleland et al., 2026). Importantly, many contemporary policies already reflect the distinction between constitutive and merely instrumental roles for process information—for example, by treating ‘AI-assisted copy editing’ differently from substantive generative use (Nature Portfolio, n.d.). This asymmetry suggests that the issue is not assistance as such but the selective treatment of certain generative tools as epistemically suspect. On this approach, disclosure should track uses that bear on the epistemic chain and accountability (for example, where tool use affects claims about evidence, analysis, or provenance), whereas purely linguistic assistance (translation, grammar, and style) should not automatically function as an epistemic defeater.

An output-first audit test may help to operationalise this distinction in editorial practice (Resnik & Hosseini, 2026; Cleland et al., 2026). The test can be stated in four steps. First, ask whether the relevant AI involvement is source-sensitive or source-insensitive in the sense introduced in Section 4.1, that is, whether AI-generated material enters the evidential chain by generating, transforming, selecting, or summarising evidence, analysis, code, images, or provenance claims on which the paper relies. Second, if it does, require a short structured disclosure in the manuscript and treat the relevant artifacts—code, prompts, logs, notebooks, or image-generation details—as part of the public epistemic record. Third, if the use is limited to linguistic assistance that does not alter evidential content or provenance claims, disclosure may be optional or minimal and should not, by itself, count against the paper epistemically. Fourth, treat disclosures as risk alerts, not verdicts: they justify targeted output checks such as citation spot-checks, requests for raw data or code, reruns of analyses, and replication or methodological clarification where anomalies appear—for example, fabricated references, absent raw data, opaque code, or unexplained mismatches between claims and supporting materials. Such checks matter in part because uncritical outsourcing of inferential steps can aggravate replicability problems, especially in fields already concerned about a replication crisis.

These cases differ, however, in epistemic structure. Where AI generates evidential materials —for example, synthetic data, images, or experimental stimuli—the central question is whether the provenance claims embedded in the public record are true and whether the resulting materials are fit for the inferential role the paper assigns to them. Where AI transforms or selects evidential materials—for example, by producing analysis code, summaries, or classificatory outputs—the central question is one of fidelity and auditability: can the transformation be inspected, rerun, and checked against the underlying materials? Where AI supplies an inferential step or interpretation, the central question is whether the author retains epistemic control by being able to endorse, defend, and integrate that step into the paper's public argumentative record. In all three cases, what is epistemologically salient is not the bare process fact that AI was used, but the specific output-level anomaly or vulnerability that such use may introduce. Process-reliabilist considerations, in Goldman's sense, matter here as defeasible indicators of likely failure modes; but their epistemic significance is assessed through output-level checks—replicability, rerunnability, provenance checks, and the author's ability to answer criticism—rather than through bare knowledge that AI was used. Disclosure therefore matters because it helps locate which kind of output-level vulnerability is in play and what sort of scrutiny it warrants—not as evidence of deficiency.

Concrete cases help. AI-written statistical code that generates the reported numbers enters the evidential chain and therefore requires disclosure and auditability. AI-generated images used as experimental stimuli, synthetic data augmentation, or tool-generated literature summaries on which substantive claims depend do likewise. By contrast, translation, grammar correction, or stylistic smoothing of already-authored claims ordinarily do not. The relevant question is not 'Was AI used somewhere in the writing process?' but 'Did AI use alter what counts as evidence, analysis, or provenance within the public record?'

The same point applies to large-scale literature identification, organisation, and routine analytic workflows in empirically oriented research. Choices of databases, search strings, inclusion and exclusion rules, summary templates, coding categories, and statistical procedures may contain substantive judgement at the protocol-design stage; but once those parameters are externalised, parts of the task can be executed by research assistants, scripts, or carefully prompted AI systems in a way that is increasingly analogous to the automation of hand calculations, bibliographic retrieval, or other routine scholarly subroutines. In such cases, the central epistemic questions concern coverage, fidelity, rerunnability, auditability, and authorial endorsement. Search protocols, database choices and dates, prompts, model versions and settings where relevant, extraction rules, logs, and analytic specifications can therefore be treated much like code, methods, or other provenance-bearing materials: not because the private process is itself justificatory, but because externalising these materials makes the evidential chain inspectable and open to targeted checking. The more a task is protocolisable and auditable, the less work is done by the human/machine distinction as such; what matters is whether the resulting public record is accurate, sufficiently complete for its stated purpose, and open to challenge.

By contrast, when the delegated step is to formulate a new question, interpretation, or inferential synthesis from a large corpus, the task is no longer purely mechanical. Variance across human and AI outputs in such cases is not itself an objection to tool use; it shows that the relevant safeguard is retained authorial epistemic control: the ability to endorse, defend, and revise the resulting inferential step under communal scrutiny.

In short, disclosure is compatible with the rejection of process essentialism. What is not compatible is the inferential step from disclosure to evaluation: the move from 'this paper was produced with AI assistance' to 'this paper is therefore epistemically deficient.' The first statement is a fact about production; the second is an epistemic judgement that should be grounded in the paper's content, not its provenance.

6.3 The Reliability Bootstrapping Objection

A more sophisticated objection draws on the reliabilist tradition more carefully. One might argue that even if output-based evaluation is the ideal, we cannot evaluate outputs without some information about the reliability of the process that produced them. This is because the reliability of evidence depends, at least in part, on the process that generated it. If we learn that a new AI tool is prone to systematic errors—for example, that it tends to generate plausible-sounding but inaccurate citations—then knowing that a paper was produced with this tool gives us relevant information about the likely reliability of the paper's content.

This is a more serious objection, and I grant that it has some force. Knowledge of the production process can function as a heuristic for predicting output quality, just as knowledge of a manufacturer's quality control processes can serve as a heuristic for predicting product reliability. But heuristics and constitutive criteria are different things. The fact that process information has predictive value for output quality does not establish that process is constitutive of epistemic status. It merely establishes that, in the absence of direct output evaluation, process information can serve as a proxy. In academic publishing, however, peer review (and post-publication scrutiny) provides at least partial output evaluation, even if it is imperfect and time-limited. Accordingly, the most defensible role for process information is as a defeasible risk-marker that can trigger additional output-checking, not as a constitutive condition on justificatory status.

Different scholarly processes may indeed differ systematically in reliability, and this much should be granted. A literature-later workflow, for example, may raise the initial risk of missed objections, false novelty, or incomplete engagement with adjacent work. Reliability differences of that kind fit naturally with PE-W: they mark certain processes as higher-risk routes, and papers produced through them remain liable to output-level screening, correction, or rejection insofar as peer review functions as an output check. But that point does not establish PE-S or PE-M. The relevant vulnerabilities are precisely of a kind that can be caught at the level of the public record, and once they have been identified and repaired, no residual justificatory defect remains solely in virtue of the private route by which the record was produced. Two output records that are genuinely equivalent in reasons, evidence, and dialectical adequacy—that is, that make the same public case and hold up equally well under scrutiny—therefore do not differ in justificatory status merely because the private route to them was different.

6.4 The Epistemic Labour Objection

A final objection worth addressing concerns the intrinsic value of epistemic labour—the cognitive effort involved in producing scholarly work. One might argue that the process of sustained intellectual engagement—reading, thinking, struggling with difficult ideas, revising one's views—is not merely instrumental to the production of good scholarship but constitutive of the scholarly enterprise itself. On this view, scholarship is not merely the production of texts; it is a form of intellectual practice, and the value of that practice is inseparable from the cognitive labour it involves.

Scholarship has long involved distributed and unevenly credited labour. Supervisors may rely on assistants for literature review, teams may divide composition from data work, and translators or editors may shape the final text. This history clarifies rather than undermines the present argument: epistemic labour may be distributed, or even ethically misallocated, without that fact alone altering the justificatory status of an otherwise well-supported claim.

Modern contributor statements are valuable because they make this division of labour visible and allocate credit and accountability more fairly. But the mere fact of such division does not show that a claim’s justificatory status depends on each named author having performed the same share of composition; it shows why credit and responsibility should be separately tracked (Allen et al., 2014).

I have some sympathy with this objection, but I think it confuses two distinct questions. The first question is whether epistemic labour has intrinsic value as a form of human intellectual development. The answer may well be yes: there is something valuable about the process of struggling with ideas, and this value may not be fully captured by the quality of the resulting text. But this is a question about the value of intellectual practice, not about the epistemology of knowledge claims. The second question is whether the presence or absence of epistemic labour affects the justificatory status of the knowledge claims expressed in a scholarly text. This is the epistemological question, and I have argued that the answer is no.

To conflate these questions is to make a category error. The value of epistemic labour as a form of intellectual practice belongs to the ethics and aesthetics of scholarly life, not to the epistemology of academic knowledge. A monastic tradition may value the process of hand-copying manuscripts as a spiritual practice without thereby claiming that hand-copied manuscripts are epistemically superior to printed ones. Similarly, the scholarly community may value sustained intellectual engagement as a professional ideal without thereby claiming that work produced through such engagement is better justified than work produced through alternative means.

7. Conclusion

The integration of LLMs into academic practice confronts the scholarly community with a fundamental epistemological question: should academic knowledge be evaluated on the basis of production process, or on the basis of publicly available evidential support? I have argued that strong process essentialism (PE-S/PE-M) is mistaken. It risks a genetic-fallacy style conflation of origins with grounds, relies on criteria that are structurally hard to observe and enforce, and—when operationalised as a disqualificatory norm—can generate exclusionary burdens that amount to processual injustice.

The alternative—output-based evaluation—is not a novel proposal. It is implicit in the logic of evidential support and in the institutional practice of peer review, which asks reviewers to assess arguments, evidence, and engagement with relevant literature rather than to reconstruct private writing histories.

Rejecting strong process essentialism does not imply that ‘anything goes.’ Process information can be instrumentally useful, defeasibly evidential, and ethically significant for credit, accountability, and integrity. Where process facts are externalised as contribution statements, registered reports, methods/provenance claims, or supporting artifacts, they can legitimately be assessed as part of the evaluative record. The challenge is therefore to justify disclosure and integrity policies without treating private provenance as epistemically disqualifying in cases where the publicly available reasons are equally strong.

I have focused mainly on LLMs as writing, translation, and editing aids in conceptually oriented scholarship. Cases in which generative tools contribute to the evidential chain—by generating or transforming data, images, code, or statistical analyses—do require disclosure and auditing, precisely because they alter the evidential and provenance claims internal to the public record, not because private drafting history is intrinsically disqualifying.

In short, academic knowledge is justified by publicly inspectable reasons—arguments, evidence, and engagement with existing scholarship—not by the provenance of the prose.

References

Allen, L., Scott, J., Brand, A., Hlava, M., & Altman, M. (2014). Publishing: Credit where credit is due. Nature, 508(7496), 312–313. https://doi.org/10.1038/508312a

Amano, T., Ramírez-Castañeda, V., Berdejo-Espinola, V., Borokini, I., Chowdhury, S., Golivets, M., González-Trujillo, J. D., Montaño-Centellas, F., Paudel, K., White, R. L., & Veríssimo, D. (2023). The manifold costs of being a non-native English speaker in science. PLOS Biology, 21(7), e3002184. https://doi.org/10.1371/journal.pbio.3002184

Chambers, C. D. (2013). Registered reports: A new publishing initiative at Cortex. Cortex, 49(3), 609–610. https://doi.org/10.1016/j.cortex.2012.12.016

Cleland, J., Driessen, E., Masters, K., Lingard, L., & Maggio, L. A. (2026). When and how to disclose AI use in academic publishing: AMEE Guide No. 192. Medical Teacher, 48(4), 542–553. https://doi.org/10.1080/0142159X.2025.2607513

Conee, E., & Feldman, R. (2004). Evidentialism: Essays in epistemology. Oxford University Press.

COPE. (2023). Authorship and AI tools (COPE position statement, February 13, 2023). Retrieved February 17, 2026, from https://publicationethics.org/guidance/cope-position/authorship-and-ai-tools

EDP Sciences. (2023, March 14). EDP Sciences aligns with COPE's position on AI and authorship in research publications. Retrieved February 17, 2026, from https://www.edpsciences.org/fr/actualites/2848-edp-sciences-aligns-with-cope-s-position-on-ai-and-authorship-in-research-publications

Fine, K. (2012). Guide to ground. In F. Correia & B. Schnieder (Eds.), Metaphysical grounding: Understanding the structure of reality (pp. 37–80). Cambridge University Press. https://doi.org/10.1017/CBO9781139149136.002

Freiman, O. (2024). AI-testimony, conversational AIs and our anthropocentric theory of testimony. Social Epistemology, 38(4), 476–490. https://doi.org/10.1080/02691728.2024.2316622

Fricker, M. (2007). Epistemic injustice: Power and the ethics of knowing. Oxford University Press.

Ganjavi, C., Eppler, M. B., Pekcan, A., Biedermann, B., Abreu, A., Collins, G. S., Gill, I. S., & Cacciamani, G. E. (2024). Publishers' and journals' instructions to authors on use of generative artificial intelligence in academic and scientific publishing: bibliometric analysis. BMJ, 384, e077192. https://doi.org/10.1136/bmj-2023-077192

Goldberg, S. C. (2010). Relying on others: An essay in epistemology. Oxford University Press.

Goldman, A. I. (1986). Epistemology and cognition. Harvard University Press.

Goldman, A. I. (1999). Knowledge in a social world. Oxford University Press.

Goldman, A. I. (1979). What is justified belief? In G. S. Pappas (Ed.), Justification and knowledge (pp. 1–23). Springer. https://doi.org/10.1007/978-94-009-9493-5_1

Greco, J. (2010). Achieving knowledge: A virtue-theoretic account of epistemic normativity. Cambridge University Press.

Hanson, N. R. (1958). Patterns of discovery: An inquiry into the conceptual foundations of science. Cambridge University Press.

He, J., & Yang, C. (2025). Testimony by LLMs. AI & Society, 40, 6201–6213. https://doi.org/10.1007/s00146-025-02366-y

Heesen, R., & Bright, L. K. (2021). Is peer review a good idea? The British Journal for the Philosophy of Science, 72(3), 635–663. https://doi.org/10.1093/bjps/axz029

Hosseini, M., Rasmussen, L. M., & Resnik, D. B. (2024). Using AI to write scholarly publications. Accountability in Research, 31(7), 715–723. https://doi.org/10.1080/08989621.2023.2168535

Hosseini, M., Resnik, D. B., & Holmes, K. (2023). The ethics of disclosing the use of artificial intelligence tools in writing scholarly manuscripts. Research Ethics, 19(4), 449–465. https://doi.org/10.1177/17470161231180449

International Committee of Medical Journal Editors (ICMJE). (n.d.). A. Use of AI by Authors. Retrieved February 17, 2026, from https://www.icmje.org/recommendations/browse/artificial-intelligence/ai-use-by-authors.html

Kitcher, P. (1993). The advancement of science: Science without legend, objectivity without illusions. Oxford University Press.

Kuhn, T. S. (1962). The structure of scientific revolutions. University of Chicago Press.

Lackey, J. (2008). Learning from words: Testimony as a source of knowledge. Oxford University Press.

Lee, C. J., Sugimoto, C. R., Zhang, G., & Cronin, B. (2013). Bias in peer review. Journal of the American Society for Information Science and Technology, 64(1), 2–17. https://doi.org/10.1002/asi.22784

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. Patterns, 4(7), 100779. https://doi.org/10.1016/j.patter.2023.100779

Longino, H. (1990). Science as social knowledge: Values and objectivity in scientific inquiry. Princeton University Press.

Lund, B. D., Wang, T., Mannuru, N. R., Nie, B., Shimray, S. R., & Wang, Z. (2023). ChatGPT and a new academic reality: Artificial intelligence-written research papers and the ethics of the large language models in scholarly publishing. Journal of the Association for Information Science and Technology, 74(5), 570–581. https://doi.org/10.1002/asi.24750

Medina, J. (2013). The epistemology of resistance: Gender and racial oppression, epistemic injustice, and resistant imaginations. Oxford University Press.

Nature Portfolio. (n.d.). Artificial intelligence (AI). Retrieved February 17, 2026, from https://www.nature.com/nature-portfolio/editorial-policies/ai

Nosek, B. A., & Lakens, D. (2014). Registered reports: A method to increase the credibility of published results. Social Psychology, 45(3), 137–141. https://doi.org/10.1027/1864-9335/a000192

Pohlhaus Jr., G. (2012). Relational knowing and epistemic injustice: Toward a theory of willful hermeneutical ignorance. Hypatia, 27(4), 715–735. https://doi.org/10.1111/j.1527-2001.2011.01222.x

Ramírez-Castañeda, V. (2020). Disadvantages in preparing and publishing scientific papers caused by the dominance of the English language in science: The case of Colombian researchers in biological sciences. PLoS ONE, 15(9), e0238372. https://doi.org/10.1371/journal.pone.0238372

Reichenbach, H. (1938). Experience and prediction: An analysis of the foundations and the structure of knowledge. University of Chicago Press.

Resnik, D. B. (2014). Data fabrication and falsification and empiricist philosophy of science. Science and Engineering Ethics, 20(2), 423–431. https://doi.org/10.1007/s11948-013-9466-z

Resnik, D. B., & Hosseini, M. (2026). Disclosing artificial intelligence use in scientific research and publication: When should disclosure be mandatory, optional, or unnecessary? Accountability in Research, 33(2), 2481949. https://doi.org/10.1080/08989621.2025.2481949

Resnik, D. B., Hosseini, M., Kim, J. J. H., Epiphaniou, G., & Maple, C. (2025). GenAI synthetic data create ethical challenges for scientists. Here's how to address them. Proceedings of the National Academy of Sciences of the United States of America, 122(9), e2409182122. https://doi.org/10.1073/pnas.2409182122

Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. In B. Hale & A. Hoffmann (Eds.), Modality: Metaphysics, logic, and epistemology (pp. 109–136). Oxford University Press.

Searle, J. R. (1995). The construction of social reality. Free Press.

Sober, E. (2008). Evidence and evolution: The logic behind the science. Cambridge University Press.

Sosa, E. (2007). A virtue epistemology: Apt belief and reflective knowledge, Volume I. Oxford University Press.

Stokel-Walker, C., & Van Noorden, R. (2023). What ChatGPT and generative AI mean for science. Nature, 614(7947), 214–216. https://doi.org/10.1038/d41586-023-00340-6

Walton, D. (1997). Appeal to expert opinion: Arguments from authority. Penn State University Press.

Wheeler, B. (2020). Reliabilism and the testimony of robots. Techné: Research in Philosophy and Technology, 24(3), 332–356. https://doi.org/10.5840/techne202049123