Plausibly Deniable Harm: How Everyday Engineering Practices Produce Hidden Risk

Jennifer R. Ayres

Independent Engineering Consultant, Adelaide, South Australia.

ORCID ID 0000-0002-4538-6512.

Abstract

Harm in engineering and infrastructure organisations is commonly attributed to failure, misconduct, or technical breakdown. This paper argues that harm can also arise through ordinary, rule-following practices that appear responsible on paper. Drawing on science and technology studies scholarship, it introduces the concept of plausibly deniable harm to describe situations in which operational and experiential knowledge is systematically excluded through formally legitimate organisational processes, producing foreseeable adverse outcomes that remain administratively unclaimable.

Through conceptual analysis informed by 25+ years of commissioning practice, the paper identifies five mechanisms through which this exclusion occurs: representational abstraction (knowledge translation losses), procedural compliance (substituting process for engagement), temporal deferral (normalising delay), credibility asymmetries (privileging formal over experiential authority), and interactional ambiguity (softening dissent through professional courtesy). These mechanisms interact to produce harm that is formally compliant yet substantively consequential.

By reframing harm as an emergent feature of compliance-oriented governance rather than exceptional failure, the paper extends science and technology studies (STS) debates on responsibility, accountability, and care in sociotechnical systems. It offers an analytical framework for understanding how organisations maintain procedural accountability while remaining substantively unresponsive, with implications for infrastructure governance, organisational design, and future empirical research on epistemic inequality in high-stakes environments.

Keywords: Plausibly deniable harm; Epistemic inequality; Governance and accountability; Operational knowledge; Invisible work; Sociotechnical systems; Science and technology studies

1. Introduction

Across contemporary engineering and infrastructure systems, organisational harm is often explained as the result of failure, misconduct, or technical breakdown. Yet science and technology studies (STS) scholarship has long shown that system reliability and safety emerge from sociotechnical arrangements that depend on forms of labour, judgement, and coordination that are routinely undervalued or rendered invisible (Star and Strauss 1999). This work demonstrates that what keeps systems functioning is not simply technical design, but the ongoing, situated practices through which people interpret, negotiate, and sustain complex infrastructures (Suchman 1995).

At the same time, STS research has highlighted that not all knowledge within organisations is treated equally. Decisions about what counts as legitimate expertise, whose judgement is trusted, and which concerns are allowed to shape action are structured by institutional hierarchies, representational norms, and governance arrangements (Irwin and Wynne 1996; Jasanoff 2004). These dynamics produce systematic asymmetries in whose knowledge is recognised, recorded, or acted upon — what this paper terms epistemic inequality (Bowker and Star 2008).

While existing literatures illuminate how knowledge becomes marginalised, it has been less explicit about how harm is produced and sustained in contexts where exclusion is neither intentional nor overt. In many engineering organisations, harm does not arise from rule-breaking or negligence. Instead, it emerges through ordinary, formally compliant practices that allow organisations to appear diligent while remaining substantively unresponsive to operational realities (Power 2013). These are environments in which concerns can be acknowledged without being incorporated, and where procedural order can mask the displacement of risk (Mol 2008; Strathern 2000).

This paper develops the concept of plausibly deniable harm to describe this patterned organisational outcome. Plausibly deniable harm refers to situations in which foreseeable adverse consequences become administratively unclaimable because decision-making processes adhere to recognised procedures that systematically exclude certain forms of knowledge (Perrow 2011). The harm is real, but its attribution is diffused through documentation, role separation, and interactional norms that make exclusion appear reasonable, professional, or procedurally justified (Fricker 2007).

The paper asks a straightforward question that anchors the conceptual contribution:

Why do engineering organisations sometimes arrive at harmful outcomes even when they comply with all required procedures, and how does the marginalisation of operational knowledge contribute to this?

To answer this question, the paper offers a conceptual analysis grounded in STS scholarship on invisible work, expertise, governance, and accountability (Jasanoff 2004; Suchman 1995).

The paper proceeds as follows. Section 2 situates plausibly deniable harm within existing STS debates and clarifies its conceptual distinctiveness. Section 3 outlines the analytical approach used to develop the mechanisms. Section 4 elaborates these mechanisms and shows how they interact to produce organisational conditions in which harm becomes both predictable and difficult to contest. Section 5 discusses the implications for STS theories of governance, expertise, and care. Section 6 concludes by reflecting on the conceptual and

practical significance of recognising plausibly deniable harm as a sociotechnical phenomenon.

2. Theoretical positioning: epistemic inequality and organisational harm

STS scholarship has long challenged the assumption that stability, safety, and system performance naturally follow from technical design alone. Foundational work on invisible labour and maintenance demonstrates that much of what keeps sociotechnical systems functioning is easy to overlook precisely because it works (Star and Strauss 1999). Tasks that prevent failure, absorb disruption, or compensate for design limitations often disappear from formal accounts of system operation, becoming visible only when something breaks (Orr 2016; Suchman 1995). This insight has been especially influential in infrastructure studies, where maintenance and operational care are routinely overshadowed by innovation, optimisation, and expansion (Graham and Thrift 2007). Operational knowledge in commissioning is often tacit, situational, and difficult to formalise, making it especially vulnerable to exclusions.

Subsequent scholarship has shown that invisibility is not merely descriptive but political. Labelling certain activities as routine, low-skill, or non-expert reinforces hierarchies of value and authority (Bowker and Star 2008). In regulated engineering contexts, the work of operators, maintainers, and coordinators is often framed as execution rather than judgement, even when it requires complex situational assessment (Irwin and Wynne 1996). As a result, knowledge generated through everyday operational engagement is frequently excluded from formal decision-making processes (Bowker and Star 2008).

Parallel work on expertise examines how organisations distinguish between legitimate and illegitimate knowledge (H. M. Collins and Evans 2002). Expertise is not simply a matter of competence or experience, but of alignment with recognised categories, credentials, and representational forms. Knowledge that can be abstracted, codified, and circulated through documents tends to carry greater authority than experiential or contextual knowledge (Porter 1995). This privileging of abstraction is especially pronounced in large technical organisations, where governance relies heavily on metrics, plans, and compliance artefacts. (Power 2013)

The result is epistemic inequality: a systematic imbalance in whose knowledge counts, when it counts, and on what terms (Bowker and Star 2008). Unlike individual bias or misrecognition, epistemic inequality is embedded in organisational routines, documentation practices, and accountability structures (Hull 2012). It persists even in the absence of discriminatory intent because institutionalised representational forms already shape what counts as valid evidence.

Work on epistemic injustice provides a complementary but distinct lens (Fricker 2007). Testimonial and hermeneutical injustice highlight how individuals and groups can be wronged in their capacity as knowers. However, this literature often focuses on interpersonal encounters or cultural marginalisation rather than the organisational mechanisms that normalise exclusion (Medina 2013). Epistemic injustice concerns interpersonal wrongs; epistemic inequality concerns structural arrangements that determine whose knowledge is institutionally actionable. This distinction is central to the argument developed here (Fricker 2007). While epistemic injustice captures interpersonal credibility deficits, commissioning environments also reveal deeper structural patterns of exclusion that exceed interpersonal dynamics.

2.1 From recognition to accountability: epistemic inequality as organisational harm

Frameworks focused on misrecognition or interpersonal injustice help explain how certain knowers are disadvantaged (Fricker 2007). Still, they do not explain how organisations maintain formal legitimacy while remaining substantively unresponsive to operational knowledge. To address this gap, this section shifts the focus from recognition to accountability, examining how epistemic inequality functions as a mechanism of organisational harm.

In many engineering organisations, decision-making processes are designed to demonstrate that the correct steps have been followed rather than to ensure that all relevant knowledge has been engaged. Documents such as plans, risk registers, and compliance reports act as proxies for due diligence, allowing organisations to claim responsibility without incorporating operational realities (Power 2013). In this context, epistemic inequality does not merely disadvantage certain knowers; it shapes how responsibility is distributed and how harm becomes contestable.

2.1.1 Plausibly deniable harm as emergent

Plausibly deniable harm is conceptualised here as an emergent organisational outcome rather than a deliberate strategy. It arises when routine, formally legitimate practices systematically exclude operational knowledge in ways that make harm foreseeable yet administratively unclaimable. Unlike “normal accidents,” which emerge from system complexity, plausibly deniable harm emerges from epistemic inequality embedded in governance arrangements (Perrow 2011). The term 'deniable' refers not to deliberate denial but to the structural condition whereby harm becomes administratively unclaimable; organisations cannot be held accountable because they followed all procedures correctly.

Plausibly deniable harm describes organisational conditions in which foreseeable adverse outcomes can be disclaimed because decision-making processes have adhered to recognised procedures, even when those procedures exclude critical forms of knowledge. Harm is not denied because it is unforeseeable, but because it falls outside the epistemic boundaries of what the organisation has defined as relevant. The harm is real, but its attribution is diffused through layers of documentation, role separation, and procedural compliance (Strathern 2000).

3. Methodological Orientation: Documents as Sites of Power and Knowledge

This paper adopts a conceptual–analytic approach informed by the author's 25+ years of commissioning practice across water, oil & gas and infrastructure, where documentary forms such as plans, risk registers, and governance records are ubiquitous. Documents are used not as a fixed empirical dataset but as illustrative artefacts that help clarify how epistemic inequality becomes embedded in organisational routines. The aim is not to measure the frequency of particular practices or to offer a systematic empirical account. Instead, the analysis draws on documentary forms that are characteristic of compliance-oriented engineering environments—plans, risk registers, correspondence, and governance records—to demonstrate the plausibility, coherence, and explanatory value of the mechanisms developed (Hull 2012; Smith 1974).

3.1 Documents as sociotechnical artefacts

Documents are approached as sociotechnical artefacts that actively shape organisational life (Asdal 2015). They do not merely record decisions; they structure what decisions are possible, whose concerns are considered relevant, and how accountability is judged (Hetherington 2012). Plans, standards, reports, and formal communications stabilise categories, translate situated activity into abstracted representations, and enable coordination across time and organisational boundaries (Akrich 1992). At the same time, they function as governance tools that allow organisations to demonstrate compliance and control in high-uncertainty environments (Austin 1975).

This dual role—representational and regulatory—makes documents a powerful conceptual lens for examining how epistemic inequality is produced and maintained. The mechanisms developed in Section 4 draw on these documentary dynamics to show how operational knowledge can be softened, reframed, or excluded while decisions remain defensible on paper.

3.2 Analytic focus: illustrative use of documentary practices

The analysis focuses on how documentary practices operate, rather than on specific documents as empirical objects. Planning documents, risk registers, compliance reports, formal correspondence, and governance records are treated as illustrative examples of broader organisational tendencies (Latour and Woolgar 1986). They help clarify how concerns are translated, reframed, deferred, or excluded as they move through organisational hierarchies(Callon and Law 1982).

3.3 Analytic procedure: conceptual synthesis

The analytic procedure is one of conceptual synthesis, informed by documentary practices but not dependent on a formal corpus. The mechanisms are developed through iterative engagement with STS scholarship on documentation, governance, expertise, and accountability, combined with conceptual illustration of how concerns are typically framed, deferred, or rendered out of scope in compliance-oriented environments (Asdal and Hobæk 2016).

Silences—what is not recorded, what is deferred, what is administratively excluded—are treated as conceptual indicators of how epistemic boundaries are drawn. This approach does not assume bad faith. Documents are understood as products of institutional logics that prioritise coherence, defensibility, and manageability over responsiveness (Asdal and Jordheim 2018). By focusing on documentary practices as conceptual resources, the analysis shows how plausibly deniable harm can emerge from legitimate organisational processes without requiring claims to systematic empirical analysis.

3.4 Documents as governance infrastructure

Finally, this conceptual orientation aligns with STS traditions that treat knowledge production as inseparable from governance. Documents are understood as infrastructures of accountability—sites where epistemic authority is negotiated and where the boundaries of responsibility are constructed (Callon 1998; Latour and Woolgar 1986).This makes documentary analysis not an empirical method in this paper, but a theoretically grounded lens for examining how epistemic inequality is enacted in contemporary engineering organisations.

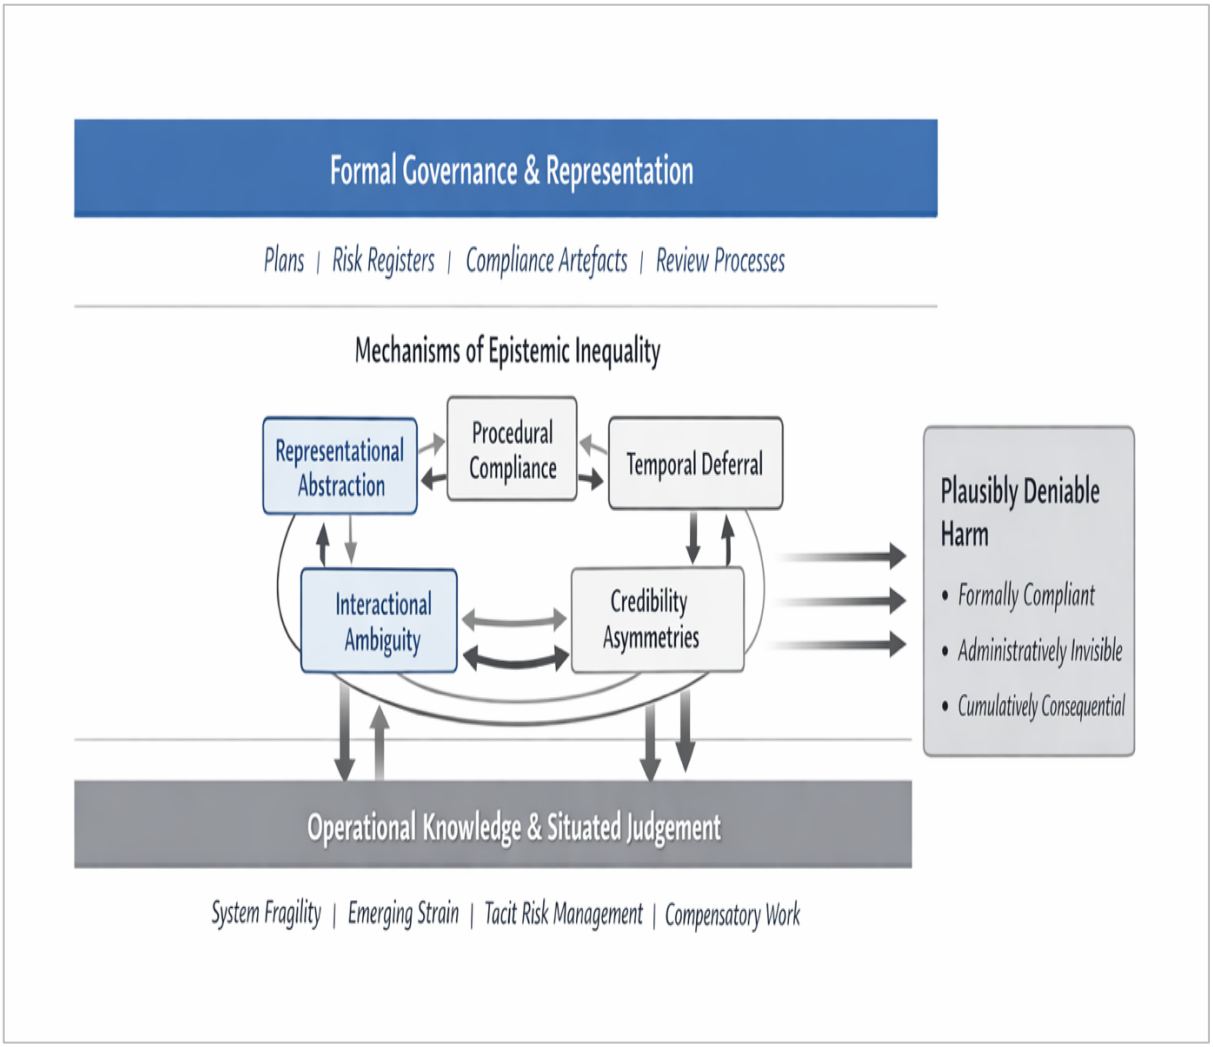

Figure 1. Plausibly deniable harm as an emergent outcome of epistemic inequality

The diagram illustrates how five mechanisms—representational abstraction, procedural compliance, temporal deferral, credibility asymmetries, and interactional ambiguity—operate between formal governance structures and operational knowledge to produce harm that is formally compliant yet substantively consequential.

3.5 Scope and limitations

This paper develops a conceptual framework rather than presenting a systematic empirical analysis. The mechanisms described are analytical constructs derived from theoretical synthesis across science and technology studies, informed by the author's more than twenty-five years of commissioning practice across water, oil and gas, and infrastructure projects. This practice-informed grounding provides observational insight into how documentary practices, procedural compliance, and credibility hierarchies routinely operate in real organisational settings.

The paper does not claim to measure prevalence, establish causation, or empirically validate the proposed mechanisms. Instead, it offers a theoretically grounded account of how epistemic inequality can plausibly produce harm under conditions of formal compliance. Future research should empirically test these mechanisms through

ethnographic studies, comparative organisational analysis, and intervention-based research examining how alternative governance arrangements alter epistemic inclusion and harm outcomes.

4. Mechanisms of Epistemic Inequality and Plausibly Deniable Harm

As prior organisational scholarship shows, formal consultation and participation processes can satisfy procedural expectations without meaningfully incorporating the knowledge they solicit, particularly when engagement is bounded by predefined scopes and timelines (Sharma and Bansal 2020).

Engineering organisations rarely intend to ignore frontline knowledge or make decisions that increase risk. Yet harm can still accumulate through ordinary, compliant, and seemingly reasonable practices. This section explains how that happens. It shows how everyday organisational routines—documentation, planning, consultation, and review—can unintentionally filter out the knowledge needed to prevent harm, while still producing a defensible record of diligence. The mechanisms described here (and summarised in Table 1) are analytically distinct, but in practice they reinforce one another. When they operate together, they create a system where operational insight is softened, delayed, or discounted, while decisions remain entirely defensible on paper. Plausibly deniable harm does not arise from a single weak point but from the interaction of multiple routine practices that make certain forms of knowledge administratively invisible.

4.1 Representational abstraction and knowledge translation

Engineering organisations rely heavily on documents to move information across teams and decision layers. Plans, schedules, risk registers, and performance reports are treated as the authoritative record of what is known. For operational staff, however, knowledge is often contextual, contingent, and grounded in how systems behave under real-world conditions (Bechky 2003). When this knowledge is translated into standard templates, much of its nuance is lost. Templates reward clarity, quantification, and alignment with predefined categories (Sykianakis and Bellas 2011). They leave little room for ambiguity, emerging risk, or the “messy” details that matter most in practice.

These patterns take archetypal forms. For instance, in water treatment commissioning, operational observations of pump behaviour under variable load show that this filtering effect is subtle but powerful. A technician’s observation that a system “sounds different under load” becomes a generic note about “monitoring performance.” A pattern of near-misses becomes a single risk line with a colour code. A complex interaction between ageing infrastructure and seasonal demand is reduced to a simple dependency arrow in a project plan. These transformations make information easier to circulate but harder to interpret. They also create a false sense of completeness: once something is written down, it appears to have been captured, even if the representational form has stripped away the very qualities that made the insight meaningful (Alcadipani et al. 2012).

This is not a failure of individuals but a structural feature of documentation-driven governance. Documents are designed to be stable, portable, and legible to people who were not present when the knowledge was generated. In achieving that stability, they inevitably flatten the situated judgement on which operational work depends. Decisions made based on these sanitised documents appear reasonable and well-informed, even when they miss

critical operational realities. Documentation becomes both a communication tool and a shield: it demonstrates diligence while obscuring the limits of what has been understood.

4.2 Procedural compliance as a substitute for engagement

Most engineering organisations have well-defined processes for consultation, review, and sign-off. These processes are designed to ensure that relevant voices are heard. In practice, procedural compliance can become a substitute for genuine engagement with operational knowledge (Power 2013). Operational staff may be invited to provide feedback, but the process's structure often limits what can be incorporated. Consultation windows may be too short, feedback may need to fit predefined categories, and issues outside the project scope may be ruled out as not applicable. Concerns that require cross-team coordination may be acknowledged but deferred. When harm later occurs, organisations can point to the completed process as evidence that they acted responsibly.

During water treatment facility design reviews, operational staff identify critical sequencing issues—startup procedures requiring simultaneous valve operations exceeding night-shift staffing, or emergency shutdown sequences incompatible with actual access routes. These concerns are raised during formal consultation, recorded, and acknowledged. However, governance structures lack mechanisms to revisit design decisions past procurement milestones. The feedback is categorised as "operational" rather than "design-critical" and deferred out of design. When staffing limitations later become visible, documented consultation serves as evidence of proper process while substantive design limitations persist.

4.3 Temporal deferral and the management of urgency

Time is one of the most powerful and least visible mechanisms shaping how knowledge is valued. Operational concerns often relate to slow degradation, cumulative risk, or capacity limits that do not align with project timelines or reporting cycles (Otsuki 2024). A typical organisational response is deferral. Issues are framed as not urgent, requiring further evidence, better addressed in the next phase, or outside the current decision window. Deferral allows decisions to proceed without explicitly rejecting operational knowledge. It maintains the appearance of prudence—“we need more data”—while postponing action indefinitely (Barua 2024).

Temporal deferral is particularly potent because it is rarely experienced as a decision at all. It is experienced as waiting, monitoring, gathering more information, or aligning with the next governance cycle. Yet these delays shape the conditions under which harm becomes possible. By the time a concern becomes urgent enough to act on, the organisation may have lost the opportunity to intervene effectively.

Temporal deferral operates powerfully through defect classification and warranty systems. During pre-operational testing, an operator notices intermittent pump cavitation—not severe enough to fail acceptance criteria but suggesting bearing wear that will worsen under continuous operation. The observation is recorded and categorised as "monitoring required," and the system is accepted because its current performance meets specifications. Six months later, during warranty review, the observation has been superseded by newer issues, the commissioning team has moved on, and context is lost. When the pump fails 18 months after handover, the warranty has expired, failure is classified as normal operational wear, and the early warning lies buried in closed-out commissioning records. A foreseeable reliability concern becomes normalised into the operational baseline, its progression

documented but never arrested. Harm emerges gradually; attribution diffuses across multiple reasonable deferrals.

4.4 Asymmetries of credibility and authority

Not all knowledge carries the same weight in organisational decision-making. Formal authority, professional status, and proximity to decision structures often matter more than experiential familiarity with system behaviour. This creates predictable credibility asymmetries. Operational staff may understand how systems behave under strain but lack the authority to have their assessments treated as decisive. Analysts or managers may produce abstract assessments that align with recognised expertise, giving their views greater legitimacy. Concerns raised informally or verbally may be discounted because they do not appear in the official record (Alcadipani et al. 2012).

Decisions can be framed as professional judgement rather than exclusion. This is a key pathway through which plausibly deniable harm is produced: the organisation can claim that it acted on the best available information, even when the information that mattered most was never treated as authoritative.

Credibility asymmetries also shape how risk is narrated. A manager's assurance that "the system is within tolerance" carries more institutional weight than a technician's observation that "it's been behaving strangely for weeks." When harm occurs, authoritative assessments serve as evidence of due diligence, even when operational insight was more accurate.

Credibility asymmetries manifest acutely during commissioning handover. A senior plant operator with 20 years' experience notes that the chemical dosing equipment configuration forces technicians into awkward positions near live electrical panels during routine maintenance, inevitably leading to unsafe workarounds or deferrals. This assessment draws on embodied knowledge of how maintenance is performed under time pressure, using available tools. Meanwhile, a consulting engineer produces a formal "maintainability assessment" concluding access meets design standards based on manufacturer specifications and ergonomic guidelines. The written engineering assessment carries institutional weight because it aligns with recognised expertise categories and produces documentable conclusions. The operator's knowledge is acknowledged but not treated as authoritative enough to delay acceptance or require modification. When an injury later occurs—a technician losing balance while reaching across a confined space—the organisation demonstrates due diligence through the professional engineering assessment. The operator's warning existed but lacked epistemic authority to shape decisions.

Table 1. Mechanisms of Epistemic Inequality and Their Contribution to Plausibly Deniable Harm

Mechanism	What it does	How it produces epistemic inequality	How it contributes to plausibly deniable harm
Representational abstraction	Translates situated knowledge into standardised forms	Filters out ambiguity and nuance and contextual detail	Decisions appear informed while critical operational realities are obscured
Procedural compliance	Substitutes process completion for substantive engagement	Limits influence of operational knowledge to predefined scopes	Organisations can claim diligence even when concerns are not acted upon
Temporal deferral	Postpones engagement with operational concerns	Reclassifies urgent issues as premature or out of scope	Harm accumulates while appearing prudently managed
Credibility asymmetries	Privileges abstracted expertise over experiential knowledge	Discounts operational judgement through hierarchical norms	Exclusion appears as professional judgement rather than dismissal
Interactional ambiguity	Maintains courteous interactional norms	Softens or reframes concerns reducing their perceived urgency	Harm is enacted through non-confrontational deflection rather than explicit refusal

Table 1 Shows how each mechanism is individually defensible but collectively produces organisational conditions in which operational knowledge is systematically excluded while decisions remain procedurally accountable.

4.5 Interactional ambiguity and formal politeness

In this paper, “interactional ambiguity” refers to the organisational effects of professional politeness, where concerns are acknowledged in ways that reduce their decisional force rather than their legitimacy.

Operational staff often learn to frame concerns carefully—as questions rather than assertions, as suggestions rather than warnings, as "just wanting to flag" rather than insisting on action. The result is that legitimate concerns reach decision-makers in forms that are easy to acknowledge but difficult to act on.

Conversely, when operational concerns are raised more directly, recipients may reframe them through polite deflection. A statement that "this configuration will create safety risks during maintenance" becomes heard as "some additional training might be helpful." A warning about system fragility becomes "we appreciate your caution and will keep monitoring." The interactional norm—maintain courtesy, avoid conflict, preserve professional relationships—takes precedence over engaging with the substance of the concern.

This mechanism is particularly insidious because it operates through practices that appear respectful and professional. No one is being dismissed or silenced in ways that would be administratively visible. Instead, concerns are acknowledged, thanked, and gently set aside. The harm that later emerges cannot be traced to anyone being rude or unprofessional. The courtesy itself becomes a way of managing dissent without appearing to do so, making exclusion administratively invisible while feeling interpersonally respectful.

4.6 Cascading Effects: How Mechanisms Reinforce Each Other

While analytically distinct, these five mechanisms rarely operate in isolation. In practice, they interact and reinforce one another, creating organisational conditions in which epistemic exclusion becomes increasingly difficult to contest. Each mechanism creates conditions that make the others more likely to occur, forming what can be understood as cascading epistemic closure—a self-reinforcing system through which operational knowledge becomes progressively marginalised while formal governance appears increasingly robust.

Representational abstraction enables procedural compliance by producing documents that appear comprehensive while filtering operational nuance. Once concerns are translated into standardised risk registers, they can be addressed through established review processes ;that demonstrate diligence, however, the abstraction strips away contextual detail that would make them actionable. Temporal deferral then depends on these abstracted representations to appear reasonable—a concern categorised as "monitoring required" seems prudent in documentary form, even as the urgency communicated through direct operational observation becomes diluted when translated into risk ratings.

Credibility asymmetries determine which concerns generate temporal urgency and which can be deferred. When engineers with formal authority identify risks, escalation is immediate; when operators raise similar concerns, they may be categorised as requiring further evidence. These asymmetries make exclusion appear as appropriate prioritisation rather than dismissal. Interactional ambiguity operates across all mechanisms, providing social lubrication that allows exclusion without confrontation. Professional politeness enables representational abstraction by framing translation as collaborative refinement, facilitates procedural compliance by maintaining collegial relationships, and makes temporal deferral appear respectful—"we appreciate your input and will consider it carefully"—rather than dismissive. Through courtesy, deflection becomes harder to name and contest.

4.6.1 An Illustration: Mechanisms Operating in Concert

The interaction of these mechanisms can be illustrated through a commissioning scenario that shows how epistemic exclusion unfolds across multiple organisational processes. During final acceptance testing of a wastewater treatment system, operations staff notice that the automated control system occasionally fails to respond to manual overrides during simulated emergency scenarios. The delay is brief—three to five seconds—and the system ultimately responds correctly, meeting contractual acceptance criteria. However, operators recognise from experience that this delay pattern suggests an underlying timing conflict in the control logic that will worsen as the system ages and as additional control loops are added during future expansions.

When this concern is raised during the formal commissioning review meeting, representational abstraction immediately shapes how it is recorded. The meeting minutes document: "Operations noted occasional response delay during emergency override testing. System performance meets acceptance criteria (Section 4.2.1). Control system functionality confirmed." The situated observation—that the pattern suggests progressive degradation—cannot be captured in the standard acceptance template, which asks only whether current performance meets specifications. The concern is translated into a compliance statement.

Credibility asymmetries reinforce this temporal logic. When the commissioning manager seeks confirmation from the control systems vendor, they receive written assurance that delays are "within normal tolerances" and "no modification is required." This engineering assessment carries institutional weight: produced by recognised experts, formally documented, and aligned with manufacturer specifications. The operators' interpretation—that the pattern indicates a timing conflict—is grounded in experiential knowledge. But this knowledge, lacking institutional markers of formal expertise, cannot override the vendor's assessment. The decision to defer is supported by credible technical judgment based on design specifications rather than operational prognosis.

Eighteen months later, after system expansion has added new control loops, the timing conflict manifests as a critical failure during an actual emergency. The control system's delayed response to manual override contributes to process upset and environmental discharge. The subsequent incident investigation finds that "the system met all acceptance criteria during commissioning" and that "operational concerns were documented and addressed through appropriate monitoring protocols." The procedural record demonstrates compliance. The organisational response appears reasonable at every decision point. Yet the harm was foreseeable to those with situated knowledge of control system behaviour, and it became administratively unclaimable through the interaction of mechanisms that individually appeared defensible.

This is how plausibly deniable harm emerges. Not from a single point of failure, but from the progressive filtering of operational knowledge through multiple organisational processes, each of which can be justified on its own terms. Representational abstraction makes procedural compliance possible. Procedural compliance enables temporal deferral. Temporal deferral is supported by credibility asymmetries. Credibility asymmetries are maintained through interactional ambiguity. Together, these mechanisms create conditions in which harm becomes foreseeable to those whose knowledge is excluded yet remains administratively invisible to those whose decisions shape organisational outcomes.

The cascading nature of these interactions means that intervening at a single point may be insufficient. Improving documentation templates will not address credibility asymmetries. Extending consultation timelines will not overcome representational abstraction if the same filtering occurs over a longer period. Changing interactional norms without addressing credibility structures may simply make exclusion more polite. Reducing plausibly deniable harm, therefore, requires attention to how these mechanisms interact to produce organisational conditions in which epistemic inequality becomes structurally embedded rather than interpersonally enacted.

5 Discussion: What plausibly deniable harm changes for STS

The concept of plausibly deniable harm offers a way to understand how harm emerges from the ordinary functioning of sociotechnical organisations, not from dramatic failures. This section reflects on how the analysis extends existing STS conversations about governance, responsibility, expertise, and care. It shows that harm can be produced through systems designed to demonstrate diligence, and that epistemic inequality is a structural feature of how organisations manage knowledge and accountability, providing vocabulary for analysing forms of harm difficult to see, contest, and attribute

5.1 Reframing harm in sociotechnical systems

Much of the STS literature focuses on moments of crisis, breakdown, or exceptional failure (Perrow 2011). In contrast, plausibly deniable harm shifts attention to the quieter, slower processes through which harm becomes normalised. When organisations rely heavily on documentation and procedural order, they can unintentionally obscure the conditions under which risk is actually increasing (Power 2013). Decisions appear coherent and defensible because they align with established governance practices, even when those practices have filtered out the knowledge needed to prevent harm (Porter 1995).

5.2 Rethinking accountability and organisational knowledge

The findings complicate conventional understandings of accountability. Organisations often distribute responsibility across teams, manage uncertainty through documentation, and construct narratives that demonstrate diligence. These practices create a form of procedural accountability that is easy to evidence but difficult to challenge (Power 2013). When review processes and documentation practices systematically exclude operational knowledge, accountability becomes a matter of demonstrating that the proper steps were followed rather than engaging with the knowledge needed to make safe, informed decisions (Strathern 2000).

Procedural accountability operates as a shared narrative about responsible governance, shaping which forms of knowledge are treated as legitimate(Jasanoff and Kim 2015).

This helps explain why organisations can appear transparent and responsible while remaining unresponsive to the concerns that matter most. It also clarifies why accountability mechanisms often fail to prevent harm: they are designed to verify compliance, not to interrogate the epistemic foundations of decision-making. In many engineering organisations, accountability is retrospective and document-based. It relies on what can be shown, not on what was known (Bowker and Star 2008). Plausibly deniable harm, therefore, offers a way to analyse how accountability can be maintained procedurally while being hollowed out substantively.

This also has implications for how responsibility is distributed. When harm occurs, organisations often point to completed processes, signed forms, and documented consultations as evidence that they acted responsibly. Yet these artefacts may reflect only a narrow slice of the knowledge available at the time. The concept of plausibly deniable harm helps illuminate how responsibility can be displaced onto processes rather than people, and how this displacement can mask the structural conditions that made harm possible.

5.3 Implications for theories of expertise and care

The analysis also contributes to broader STS debates on expertise and care. Scholars have long argued that system reliability depends on forms of maintenance, coordination, and situated judgement that are undervalued within formal governance structures (Suchman 1995). Plausibly deniable harm shows that this undervaluation is not simply a matter of recognition. It is a mechanism through which harm is produced. When operational knowledge is abstracted, deferred, or discounted, the labour of care becomes both essential and institutionally unsupported.

This raises questions about what counts as expertise in engineering organisations. Formal authority often carries more weight than experiential familiarity with system behaviour, even when the latter is more accurate (H. Collins and Evans 2007). The analysis shows how these credibility asymmetries shape not only whose knowledge is heard but whose

knowledge is allowed to matter. It also highlights the emotional and ethical labour carried by operational staff who continue to care for systems even when their concerns are repeatedly softened or deferred. Their expertise is indispensable, yet structurally marginalised.

The concept of plausibly deniable harm also raises fundamental questions about what counts as care work in engineering organisations and whose care work is valued. Care, as theorised by Mol (2008) and extended by Puig de la Bellacasa (2017) involves attentiveness to vulnerability, maintenance of relational connections, and ongoing practices of noticing and responding. In commissioning contexts, operators and technicians engage continuously in precisely these practices—monitoring system behaviour, noticing anomalies, interpreting performance changes, anticipating failure modes—that constitute essential care work for sociotechnical systems. Yet this care remains largely invisible within governance structures that privilege abstracted assessment over situated attention (Star and Strauss, 1999).

When credibility asymmetries discount operational judgment, they simultaneously devalue the care work on which system reliability actually depends. This creates a distinctive double bind for operational staff. They must continue caring for systems—attending to fragility, compensating for design limitations, preventing failures through vigilant interpretation—even as organisational governance structures fail to recognise these activities as requiring their specific form of situated knowledge. The emotional and ethical labour involved in maintaining attention to system vulnerability, particularly when concerns are repeatedly deflected through the mechanisms described here, represents an unacknowledged form of organisational burden that extends beyond conventional accounts of invisible work (Strauss 1985).

This burden is compounded by the epistemic violence enacted when care work is systematically discounted. To have one's professional judgment about system vulnerability repeatedly abstracted into compliance categories, deferred to later phases, or overridden by formal assessments that lack operational grounding is not merely a matter of misrecognition—it is an erosion of the epistemic foundations on which careful practice depends. Operators who observe their warnings documented but not acted upon, acknowledged but not engaged, must decide whether to continue raising concerns or withdraw into compliance with the very documentary systems that filter their knowledge. The former requires sustained emotional labour in the face of recurring deflection; the latter requires accepting complicity in harm that operational knowledge suggests is preventable.

Plausibly deniable harm, therefore, is not only a matter of epistemic inequality but also a question of whose care work is valued and supported within sociotechnical systems. Organisations that rely on operational care while systematically excluding operational knowledge from governance create conditions in which care becomes both essential and unsustainable. This dynamic helps explain patterns of quiet professional attrition, burnout, and disengagement among operational staff who continue to carry responsibility for system safety without the epistemic authority to shape decisions that determine safety outcomes. It also illuminates why post-incident investigations often find that "warning signs were present but not acted upon"—the care work of noticing was occurring, but the organisational structures for acting on operational knowledge were absent or non-functional.

Recognising care as an epistemic practice, not merely an affective orientation, shifts how we understand what constitutes epistemically just governance. It suggests that valuing operational knowledge requires more than including operators in consultation processes; it

requires treating operational care work as a legitimate basis for shaping organisational decisions. This, in turn, challenges conventional divisions between "operational" and "strategic" knowledge, revealing how such categories themselves function as mechanisms of epistemic inequality that enable plausibly deniable harm.

5.4 Conceptual portability beyond engineering

Although this study is grounded in engineering and infrastructure governance, the concept of plausibly deniable harm has broader relevance. Similar patterns can be seen in healthcare, public administration, and algorithmic decision-making, where harm emerges not from overt failure but from governance arrangements that render certain forms of knowledge administratively irrelevant (Jasanoff 2004). The concept, therefore, offers a portable lens for examining how organisations across sectors manage uncertainty, distribute responsibility, and justify decisions that have harmful consequences.

This portability also suggests that plausibly deniable harm may help bridge conversations across subfields of STS. It connects work on organisational knowledge, maintenance and care, accountability, and sociotechnical imaginaries by showing how these dynamics converge in the production of harm (Jasanoff and Kim 2015). It also provides a way to analyse how governance systems create the conditions under which harm becomes both possible and deniable.

In healthcare, similar patterns emerge through electronic health record systems that translate nursing observations of patient deterioration—subtle behavioural changes that experienced nurses recognise as early warnings—into standardised data fields that may not capture the situated assessment giving observations their significance (Mol, 2008). Once abstracted, concerns can be reviewed through clinical governance processes, demonstrating documentation compliance without engaging the interpretive nursing knowledge that generated concern. Credibility asymmetries privilege physicians' diagnostic conclusions over nurses' judgment (Elhihi et al. 2025), even when nurses' continuous patient contact provides distinctive observational access. When adverse outcomes occur, healthcare organisations can demonstrate proper protocols were followed, yet the harm may have been foreseeable to nursing staff whose knowledge was filtered through these governance mechanisms.

In policy contexts, community consultation processes demonstrate procedural inclusion while filtering experiential knowledge into policy-legible categories. Indigenous environmental knowledge, local residents' understanding of neighbourhood dynamics, or community health workers' insights must fit predetermined frameworks that may render certain concerns out of scope (Li 2005; Whyte 2018). Credibility asymmetries privilege technical expertise or academic research over community knowledge, even when lived engagement provides access to dynamics that formal methods miss.

In algorithmic systems, representational abstraction operates through training data and feature selection, in which design choices are embedded, and systems subsequently appear technically neutral (Benjamin 2019). Procedural compliance is demonstrated through fairness metrics and algorithmic audits (Raji et al. 2020), while credibility asymmetries privilege data science expertise over affected communities' accounts of discriminatory outcomes.

These parallels suggest plausibly deniable harm may be pervasive in governance contexts combining documentary accountability, uneven expertise distributions, organisational complexity, and high-stakes harm attribution. The concept offers an analytical purchase for

understanding how harm emerges from formally legitimate arrangements across diverse sociotechnical systems.

5.5 Normative implications: towards epistemically inclusive governance

Finally, the analysis raises normative questions about how sociotechnical governance might be reconfigured to reduce the conditions under which plausibly deniable harm emerges. More epistemically inclusive organisations would recognise operational judgement as a legitimate form of expertise, create pathways for situated knowledge to shape decisions, and treat documentation as a complement rather than a substitute for engagement (Bowker and Star 2008). Acknowledging the structural production of deniability is a necessary first step toward designing systems that are substantively, rather than merely procedurally, accountable.

5.6 Future empirical directions

This framework suggests several empirical research directions. First, detailed ethnographic studies of how documentary practices actually unfold in commissioning environments could test whether the mechanisms operate as theorised. Second, a comparative analysis across infrastructure sectors could reveal contextual variations in how epistemic inequality harms. Third, intervention studies examining what happens when organisations adopt more epistemically inclusive practices could demonstrate practical applications.

6. Conclusion: Rethinking Harm, Knowledge, and Responsibility

This paper has shown how harm can emerge through the ordinary functioning of engineering and infrastructure organisations. Plausibly deniable harm describes a pattern in which operational knowledge is softened, delayed, or discounted through routine practices that appear reasonable, compliant, and defensible (Irwin and Wynne 1996). These practices—representational abstraction, procedural compliance, temporal deferral, credibility asymmetries, and interactional ambiguity—do not operate in isolation. They reinforce one another, creating a governance environment in which harmful decisions can be justified even when they rest on incomplete or distorted understandings of system behaviour.

The concept of plausibly deniable harm contributes to STS by offering a way to analyse how harm is produced not only through breakdowns or misconduct but through the structural features of sociotechnical governance. It highlights the importance of examining how organisations manage knowledge, whose expertise is treated as authoritative, and how accountability is constructed through documentation and process (Bowker and Star 2008). It also underscores the need to attend to the temporal dynamics of decision-making, where delays and deferrals can accumulate into conditions that make harm more likely.

For industry, the analysis offers a practical insight: organisations can meet all procedural requirements and still make decisions that increase risk (Power 2000). Recognising this dynamic opens the possibility of designing more epistemically inclusive systems—systems that value operational judgement, create space for uncertainty, and treat documentation as a tool rather than a proxy for understanding.

During water treatment commissioning, a commissioning engineer observes that filter backwash turbidity spikes suggest media bed channelling—indicating the filter media has

segregated and will require premature replacement. This observation emerges from watching system behaviour across multiple cycles and recognising when patterns signal degradation rather than normal variation. When entered into the standard commissioning template, it becomes "Filter 3: Backwash performance meets design specification (turbidity reduction >95% achieved)." The concern about long-term reliability disappears into green RAG (Red Amber Green) status. The template asks only whether current performance meets specifications, not whether operational behaviour suggests future problems. The system is formally accepted as fully functional, while the operational concern remains unaddressed and undocumented.

Ultimately, plausibly deniable harm invites a shift in how organisations think about responsibility. It suggests that preventing harm requires more than compliance with established processes (Strathern 2000). It requires attention to the epistemic conditions under which decisions are made, the forms of knowledge that are allowed to shape those decisions, and the structural mechanisms that make some concerns visible while rendering others administratively irrelevant. By making these dynamics explicit, the concept provides a foundation for rethinking how sociotechnical systems can be governed in ways that are both procedurally and substantively accountable.

Or, as one industry professional put it: design stays on paper until operators make it work, legitimising the faults they inherit. When accidents happen, everyone follows the process, so a no-blame culture protects the incompetent. We keep assessing compliance with flawed designs rather than actual performance, and the cycle continues until we consult operational staff before finalising designs.

AI use statement

The author used generative artificial intelligence (Grammarly) to assist with language editing and structural clarity during manuscript preparation. All conceptual framing, theoretical development, interpretation, and substantive content are the author's own. The author takes full responsibility for the integrity, accuracy, and originality of the work.

References

Akrich, Madeleine. 1992. "The De-Scription of Technical Objects." *Shaping Technology-Building Society: Studies in Sociotechnical Change*, January.

Alcadipani, Rafael, Farzad Rafi Khan, Ernesto Gantman, and Stella Nkomo. 2012. "Southern Voices in Management and Organization Knowledge." *Organization* 19 (2): 131–43. doi:10.1177/1350508411431910.

Asdal, Kristin. 2015. "What Is the Issue? The Transformative Capacity of Documents." *Distinktion: Journal of Social Theory* 16 (1): 74–90. doi:10.1080/1600910X.2015.1022194.

Asdal, Kristin, and Bård Hobæk. 2016. "Assembling the Whale: Parliaments in the Politics of Nature." *Science as Culture* 25 (1): 96–116. doi:10.1080/09505431.2015.1093744.

Asdal, Kristin, and Helge Jordheim. 2018. "Texts on the Move: Textuality and Historicity Revisited." *History and Theory* 57 (March): 56–74. doi:10.1111/hith.12046.

Austin, J. L. 1975. *How to Do Things with Words*. 2d ed. The William James Lectures 1955. Oxford [Eng.]: Clarendon Press.
Barua, Maan. 2024. *Plantation Worlds*. Duke University Press. doi:10.1215/9781478027744.
Bechky, Beth A. 2003. "Sharing Meaning Across Occupational Communities: The Transformation of Understanding on a Production Floor." *Organization Science* 14 (3): 312–30. doi:10.1287/orsc.14.3.312.15162.
Benjamin, Ruha. 2019. *Race after Technology: Abolitionist Tools for the New Jim Code*. Medford, MA: Polity.
Bowker, Geoffrey C., and Susan Leigh Star. 2008. *Sorting Things out: Classification and Its Consequences*. 1. paperback ed., 8. print. Inside Technology. Cambridge, Mass.: MIT Press.
Callon, Michel, ed. 1998. *The Laws of the Markets*. Sociological Review Monograph Series. Oxford ; Malden, MA: Blackwell Publishers/Sociological Review.
Callon, Michel, and John Law. 1982. "On Interests and Their Transformation: Enrolment and Counter-Enrolment." *Social Studies of Science* 12 (4): 615–25. doi:10.1177/030631282012004006.
Collins, Harry, and Robert Evans. 2007. *Rethinking Expertise*. University of Chicago Press. doi:10.7208/chicago/9780226113623.001.0001.
Collins, H.M., and Robert Evans. 2002. "The Third Wave of Science Studies: Studies of Expertise and Experience." *Social Studies of Science* 32 (2): 235–96. doi:10.1177/0306312702032002003.
Elhihi, E.A., K.L. Aljarary, M. Alahmadi, J.B. Adam, O.A. Almwualllad, M.S. Hawsawei, A.A. Hamza, and I.A. Abdullatif Ibrahim. 2025. "The Mediating Role of Moral Courage in the Relationship between Ethical Leadership and Error Reporting Behavior among Nurses in Saudi Arabia: A Structural Equation Modeling Approach." *BMC Nursing* 24 (1). doi:10.1186/s12912-025-03043-7.
Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford university press.
Graham, Stephen, and Nigel Thrift. 2007. "Out of Order: Understanding Repair and Maintenance." *Theory, Culture and Society* 24 (May). doi:10.1177/0263276407075954.
Hetherington, Kregg. 2012. "*Guerrilla Auditors: The Politics of Transparency in Neoliberal Paraguay* by Kregg Hetherington." *PoLAR: Political and Legal Anthropology Review* 35 (2): 345–47. doi:10.1111/j.1555-2934.2012.01207.x.
Hull, Matthew S. 2012. *Government of Paper: The Materiality of Bureaucracy in Urban Pakistan*. Berkeley: University of California Press.
Irwin, Alan, and Brian Wynne, eds. 1996. *Misunderstanding Science?: The Public Reconstruction of Science and Technology*. 1st ed. Cambridge University Press. doi:10.1017/CBO9780511563737.
Jasanoff, Sheila, ed. 2004. *States of Knowledge*. 0 ed. Routledge. doi:10.4324/9780203413845.
Jasanoff, Sheila, and Sang-Hyun Kim, eds. 2015. *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. Chicago ; London: The University of Chicago Press.
Latour, Bruno, and Steve Woolgar. 1986. *Laboratory Life: The Construction of Scientific Facts*. Princeton, N.J: Princeton University Press.
Li, Tania Murray. 2005. "Beyond 'the State' and Failed Schemes." *American Anthropologist* 107 (3): 383–94. doi:10.1525/aa.2005.107.3.383.

Medina, José. 2013. *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and the Social Imagination*. 1st ed. Oxford University PressNew York. doi:10.1093/acprof:oso/9780199929023.001.0001.
Mol, Annemarie. 2008. *The Logic of Care*. 0 ed. Routledge. doi:10.4324/9780203927076.
Orr, Julian Edgerton. 2016. *Talking about Machines: An Ethnography of a Modern Job*. Collection on Technology and Work. Ithaca: Cornell University Press.
Otsuki, Kei. 2024. "Infrastructural Violence and Its Temporalities." In *Handbook of Infrastructures and Cities*, edited by Olivier Coutard and Daniel Florentin, 240–54. Edward Elgar Publishing. doi:10.4337/9781800889156.00026.
Perrow, Charles. 2011. *Normal Accidents: Living with High Risk Technologies - Updated Edition*. Princeton, NJ: Princeton University Press. doi:10.1515/9781400828494.
Porter, Theodore M. 1995. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton (NJ): Princeton university press.
Power, Michael. 2000. "The Audit Society — Second Thoughts." *International Journal of Auditing* 4 (1): 111–19. doi:10.1111/1099-1123.00306.
—— 2013. *The Audit Society: Rituals of Verification*. Repr. Oxford: Oxford University Press.
Puig de la Bellacasa, María. 2017. *Matters of Care: Speculative Ethics in More Than Human Worlds. University of Minnesota Press. 265 Pages.*
Raji, Inioluwa Deborah, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. "Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing." In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44. Barcelona Spain: ACM. doi:10.1145/3351095.3372873.
Sharma, Garima, and Pratima (Tima) Bansal. 2020. "Cocreating Rigorous and Relevant Knowledge." *Academy of Management Journal* 63 (2): 386–410. doi:10.5465/amj.2016.0487.
Smith, Dorothy E. 1974. "The Social Construction of Documentary Reality[1]." *Sociological Inquiry* 44 (4): 257–68. doi:10.1111/j.1475-682X.1974.tb01159.x.
Star, Susan Leigh, and Anselm Strauss. 1999. "Layers of Silence, Arenas of Voice: The Ecology of Visible and Invisible Work." *Computer Supported Cooperative Work (CSCW)* 8 (1–2): 9–30. doi:10.1023/A:1008651105359.
Strathern, Marilyn. 2000. "The Tyranny of Transparency." *British Educational Research Journal* 26 (3): 309–21. doi:10.1080/713651562.
Strauss, A. 1985. *Social Organization of Medical Work*. Chicago (Ill.) London: University of Chicago press.
Suchman, Lucy. 1995. "Making Work Visible." *Communications of the ACM* 38 (9): 56–64. doi:10.1145/223248.223263.
Sykianakis, Nikos, and Athanasios Bellas. 2011. "Organization Politics and the Role of Accounting." *EUROPEAN RESEARCH STUDIES JOURNAL* XIV (Issue 3): 119–34. doi:10.35808/ersj/330.
Whyte, Kyle P. 2018. "Indigenous Science (Fiction) for the Anthropocene: Ancestral Dystopias and Fantasies of Climate Change Crises." *Environment and Planning E: Nature and Space* 1 (1–2): 224–42. doi:10.1177/2514848618777621.