The Continuity Trap in Data Science Health Research

Clement Adebamowo, Sally Nneoma Adebamowo, Adeola Akintola, Peter Ikhane, Simisola Akintola, Temidayo Ogundiran, Ayodele Jegede, Olusegun Adeyemo, Shawneequa Callier, Muhammad Imam-Thamim, Ibrahim Uthman, BridgELSI BridgELSI Project as part of the DS-I Africa Consortium

Submitted to: Journal of Medical Internet Research
on: April 17, 2026

Table of Contents

The Continuity Trap in Data Science Health Research

Clement Adebamowo[1, 2] BMChB Hons, SCD; Sally Nneoma Adebamowo[1, 2] MBBS, MSc, SCD; Adeola Akintola[3] MSc; Peter Ikhane[1, 4, 5] MSc, PhD; Simisola Akintola[6, 4] BL, LLB, LLM, PhD; Temidayo Ogundiran[4, 3, 7] MBBS, MHSc; Ayodele Jegede[4, 8] PhD; Olusegun Adeyemo[3] MSc; Shawneequa Callier[9, 10] JD, MA; Muhammad Imam-Thamim[11] PhD; Ibrahim Uthman[12] PhD; BridgELSI BridgELSI Project as part of the DS-I Africa Consortium[3]

[1]Epidemiology and Public Health School of Medicine University of Maryland, Baltimore Baltimore US
[2] University of Maryland Marlene and Stewart Greenebaum Comprehensive Cancer Center Baltimore US
[3] Department of Research Center for Bioethics and Research Ibadan NG
[4]Department of Bioethics and Medical Humanities Faculty of Multidisciplinary Studies University of Ibadan Ibadan NG
[5] Deaprtment of Research Center for Bioethics and Research Ibadan NG
[6] Faculty of Law University of Ibadan Oyo NG
[7]Department of Surgery College of Medicine University of Ibadan Oyo NG
[8]Department of Sociology Faculty of Social Sciences University of Ibadan Ibadan NG
[9]Department of Clinical Research and Leadership School of Medicine and Health Sciences George Washington University Washinton DC US
[10] Center for Research on Genomics and Global Health National Human Genome Research Institute Bethesda US
[11] University of Ilorin Ilorin NG
[12] Faculty of Islamic Studies University of Ibadan Oyo NG

Corresponding Author:
Clement Adebamowo BMChB Hons, SCD

University of Maryland Marlene and Stewart Greenebaum Comprehensive Cancer Center
660 W Redwood St
Baltimore
US

Abstract

Secondary use is now the ordinary course for data and biospecimens in health research. Clinical records collected for care become training data for prediction tools; archived images become foundation models; legacy biospecimens become renewable cell lines; and large corpora are repurposed to build health-related language models. Governance nevertheless continues to privilege the most obvious signals of persistence such as provenance logs, repository approvals, broad-consent forms, locality-preserving architectures, and documented ingestion pipelines, as if they were sufficient to establish legitimacy. They are not. In this article, we define the Continuity Trap as a continuity-specific form of proxy closure: a review-stage governance error in which a salient continuity signal in one domain is treated as sufficient reason to stop inquiry into whether semantic, authorization, and relational continuity have also been preserved. The concept is narrower than generic proceduralism, ethics washing, or proxy failure, because it isolates a specific inferential mistake in secondary-use review; and it is distinct from Goodhart's and Campbell's laws, which describe the dynamic corruption of measures once they become targets. The Continuity Trap can occur at an earlier stage, even in good-faith review. Continuity of ethical governance in data science health research must therefore be assessed across four domains, provenance, semantics, authorization, and relational standing, that we previously developed in our Representational Veracity framework. These domains can diverge as data are linked, transformed, modeled, and redeployed. We use vignettes from polygenic risk scores, legacy induced pluripotent stem cell derivation, federated learning, and health-related large language models to illustrate the problem. The policy implication is not universal re-review, but triggered continuity review whenever visible continuity is likely to be overread.

(JMIR Preprints 17/04/2026:98699)
DOI: https://doi.org/10.2196/preprints.98699

Preprint Settings

1) Would you like to publish your submitted manuscript as preprint?
✔ **Yes – Make my manuscript publicly available as a preprint immediately.**

No – I do not wish to publish a preprint.

2) If accepted for publication in a JMIR journal, would you like the PDF to be visible to the public?

✓ **Yes, please make my accepted manuscript PDF available to anyone at any time (Recommended).**

No. Please do not make my accepted manuscript PDF available to anyone. I understand that if I later pay to participate in <a href="https:/

Original Manuscript

The Continuity Trap in Data Science Health Research

Clement Adebamowo BM ChB Hons, ScD[1,2], Sally N. Adebamowo MBBS, MSc, ScD[1,2], Adeola Akintola MSc[3], Peter Ikhane PhD[1,4], Simisola Akintola LLB, BL, LLM, PhD[4,5], Temidayo Ogundiran MBBS, MHSc[3,4,6], Ayodele Jegede PhD[4,7], Olusegun Adeyemo MSc[3], Shawneequa Callier JD, MA[8,9], Muhammad K. Imam-Thamim PhD[10], Ibrahim Uthman PhD[11], BridgELSI Project as part of the DS-I Africa Consortium[3]

[1] Department of Epidemiology and Public Health, University of Maryland School of Medicine, Baltimore, Maryland. [2] Greenebaum Comprehensive Cancer Center, University of Maryland School of Medicine, Baltimore, Maryland. [3] Department of Research, Center for Bioethics and Research, Ibadan, Nigeria. [4] Department of Bioethics and Medical Humanities, Faculty of Multidisciplinary Studies, University of Ibadan, Ibadan, Nigeria. [5] Department of Private and Property Law, Faculty of Law, University of Ibadan, Ibadan, Nigeria. [6] Department of Surgery, College of Medicine, University of Ibadan, Ibadan, Nigeria. [7] Department of Sociology, University of Ibadan, Ibadan, Nigeria. [8] Department of Clinical Research and Leadership, School of Medicine and Health Sciences, The George Washington University, Washington, DC. [9] Center for Research on Genomics and Global Health, National Human Genome Research Institute, National Institutes of Health, Bethesda, Maryland. [10] Department of Islamic Law, Faculty of Law, University of Ilorin, Ilorin, Nigeria. [11] Department of Arabic and Islamic Studies, University of Ibadan, Ibadan, Nigeria.

Corresponding author: Clement Adebamowo, University of Maryland School of Medicine, 660 West Redwood Street, Baltimore, MD 21201, USA. Email: cadebamowo@som.umaryland.edu

Keywords: Ethical governance continuity, Machine learning governance, Dataset shift, Concept drift, Model accountability

Abstract

Secondary use is now the ordinary course for data and biospecimens in health research. Clinical records collected for care become training data for prediction tools; archived images become foundation models; legacy biospecimens become renewable cell lines; and large corpora are repurposed to build health-related language models. Governance nevertheless continues to privilege the most obvious signals of persistence such as provenance logs, repository approvals, broad-consent forms, locality-preserving architectures, and documented ingestion pipelines, as if they were sufficient to establish legitimacy. They are not. In this article, we define the Continuity Trap as a continuity-specific form of proxy closure: a review-stage governance error in which a salient continuity signal in one domain is treated as sufficient reason to stop inquiry into whether semantic, authorization, and relational continuity have also been preserved. The concept is narrower than generic proceduralism, ethics washing, or proxy failure, because it isolates a specific inferential mistake in secondary-use review; and it is distinct from Goodhart's and Campbell's laws, which describe the dynamic corruption of measures once they become targets. The Continuity Trap can occur at an earlier stage, even in good-faith review. Continuity of ethical governance in data science health research must therefore be assessed across four domains, provenance, semantics, authorization, and relational standing, that we previously developed in our Representational Veracity framework. These domains can diverge as data are linked, transformed, modeled, and redeployed. We use vignettes from polygenic risk scores, legacy induced pluripotent stem cell derivation, federated learning, and health-related large language models to illustrate the problem. The policy implication is not universal re-review, but triggered continuity review whenever visible continuity is likely to be overread.

Introduction

Data science health research is built on the secondary use of data that were rarely generated with present applications in view.[1, 2] Electronic health records collected for clinical care are repurposed for sepsis alerts, deterioration scores, and generative summarization tools. Imaging records are assembled to become training materials for foundation models. Stored biospecimens and genomic datasets collected under one research horizon are redeployed for new analytic tasks, new commercial pathways, and new inferential claims. Even when data never leave the originating institution, as in federated learning, they can still be functionally reoriented toward purposes that neither patients nor institutions originally contemplated.[3-8]

In these settings, governance understandably privileges what can be verified. Such verifiable entities include chain-of-custody records, access approvals, signed consent documents, provenance metadata, documented preprocessing steps, model cards, and the “data stayed local” assurances.[9-12] But these answer only some of the questions that ethical review of secondary data use in data science health research should ask. These entities can certify lineage, documentation, and sometimes confidentiality safeguards, but they do not, by themselves, establish that current use remains ethically continuous with what participants, patients, clinicians, or communities were asked to entrust to researchers and clinicians when the primary data was collected.[8, 13-18] An ethics, data governance, or data access committee may treat a visible sign of persistence of ethical governance as if it means the present use of secondary data remains ethically continuous with earlier stages of the same data or material. That inference is unsafe.[19]

In this article, we define this inferential mistake as the *Continuity Trap - an ethical governance error of proxy closure where a salient signal of continuity in one domain is allowed to do the work of legitimacy for other domains, and ethical review closes prematurely on that basis.* Two operational clarifications are needed. First, “recognition” of cross-domain strain requires more than generic procedural caution. A review body recognizes cross-domain strain when it explicitly identifies which domain, provenance, semantics, authorization, or relational standing, has weakened and documents that identification in its review record. Imposing additional documentation requirements without specifying the domain at risk is precisely the kind of undifferentiated procedural response that allows proxy closure to persist. Second, “targeted” means that the safeguard imposed must be functionally matched to the domain that is under strain: if semantic continuity is the weak link, the response must address descriptor validation, not simply procure another provenance certificate.

The concept is narrower than generic critiques of proceduralism, ethics washing, or organizational decoupling.[20-23] It is also analytically distinct from Haggerty’s ethics creep, which describes expansionary scope of procedural review, and from Bosk and de Vries’s bureaucracies of mass deception, which identify substitution of compliance for genuine deliberation.[20, 24] These are important critiques, but the Continuity Trap identifies something different: not why more activities come under review, but why review, even when it occurs, may close prematurely on the strength of a single domain signal. Nor does the Continuity Trap replicate Goodhart’s and Campbell’s laws: those describe dynamic corruption of measures once they become targets, whereas the Continuity Trap can occur earlier, in good-faith review, before any gaming or optimization pressure exists.[25, 26] Once institutions elevate provenance, locality, or broad consent into compliance targets, Goodhart/Campbell dynamics then intensify the trap, but the trap precedes and does not require them.

The concept also fills a gap in the data science literature. Dataset shift, concept drift, representation bias, transportability failure, and contextual integrity all identify real problems, but they do not explain why those problems become institutionally underweighted.[27-30] The Continuity Trap explains how these problems are masked when visible continuity elsewhere is allowed to terminate inquiry.[11, 13]

From Proxy Failure to the Continuity Trap

Proxy closure, as described in Strathern's audit literature, Campbell's law, Goodhart's law, and the broader proxy-failure literature—refers to the institutional practice of relying on a partial proxy for a complex value, treating the proxy as adequate, and stopping inquiry into what the proxy leaves out (Table 1).[25, 26, 31-33] The Continuity Trap is a continuity-specific subtype of proxy closure with three distinctive features. First, the proxied object is not "quality," "fairness," or "compliance" in the abstract, but ethical governance continuity in secondary data use. Second, the unit of failure is the review episode: an access decision, approval, derivation decision, transfer authorization, or clearance for deployment. Third, the structure of the error is cross-domain closure: a sign of continuity in one domain—most commonly provenance, but sometimes broad consent, locality preservation, repository approval, or input documentation—is taken as sufficient reason to stop inquiry into other domains.

Table 1. Comparison of the Continuity Trap, Proxy Closure, and Goodhart/Campbell Dynamics

Concept	What is being overread or optimized?	Unit of failure	Why it matters here
Proxy closure	A partial proxy is treated as if it exhausts a more complex moral or institutional object.	Any governance or evaluative setting	Broader genus: legibility substitutes for fidelity.
Continuity Trap	A continuity signal in one domain is treated as sufficient evidence of ethical continuity overall.	Secondary-use review episode	Continuity-specific subtype: provenance, broad consent, locality, or documentation closes inquiry into semantics, authorization, and relational standing.
Goodhart/Campbell dynamics	A measure becomes a target, and optimization pressure degrades its	Incentive or compliance	Dynamic intensifier: once provenance completeness, locality, or model

Concept	What is being overread or optimized?	Unit of failure	Why it matters here
	connection to the underlying goal.	regime	documentation become targets, actors optimize the visible marker rather than the underlying ethical object.

The Continuity Trap Distinguished from Goodhart's Law

Goodhart's and Campbell's laws describe dynamic corruption: once a measure becomes a target, actors optimize toward the measure, and the measure becomes less reliable as an indicator of the underlying goal.[25, 26, 32] The Continuity Trap is different because it occurs at the review stage, where a visible continuity signal is treated as sufficient, even before gaming, strategic adaptation, or metric manipulation occurs. A federated-learning proposal approved because "the data never moved" can exhibit the Continuity Trap even if every actor is acting transparently and no one is gaming anything. The mistake lies in assuming that locality settles authorization, semantic commensurability, accountability, and distributive fairness. Goodhart dynamics enter later, when institutions convert locality preservation, provenance completeness, or model-card presence into checklist targets. This distinction means one can diagnose the Continuity Trap in a conscientious review committee without alleging manipulation, while also explaining why audit-heavy regimes predictably worsen the problem over time.[6, 7, 23, 25, 26, 32, 34]

Domains of Ethical Governance Continuity

Ethical governance continuity in secondary data use is not monolithic. It is relation-specific across four domains, developed fully in the companion Representational Veracity (RV) framework.[19] We summarize them here to orient the diagnostic logic of the Continuity Trap.

Throughout this article, we use the term *data-stage* to refer to the specific state of a dataset, biospecimen, model, or derivative at the point of review - the concrete object in its present form (a linked dataset, a fine-tuned model, a derived cell line, an aggregate embedding space, a deployed classifier, or a commercially licensed derivative) - rather than the abstract dataset as originally collected. The governance question is *which ethical continuity relations still hold strongly enough at this stage to support the decision that must be made now*.

Each domain must be distinguished from the visible signal that may be offered as evidence for it: a chain-of-custody record is not provenance continuity itself; a consent form is not authorization continuity itself; a model card is not relational accountability. These are evidence artifacts. The ethical question is whether the underlying continuity relation still holds.

Provenance continuity concerns lineage. Where did this data or material come from, what transformations have occurred, and whether the present stage can be traced to the earlier stages. *Semantic continuity* concerns whether descriptors, ontologies, labels, and proxies still correspond to the constructs they are taken to represent. *Authorization continuity* concerns whether the permissions and expectations being relied upon remain normatively coherent with current use. *Relational standing* concerns whether persons, communities, and source institutions retain meaningful capacity to shape conditions, contest uses, or share in value when secondary use generates group-level claims or downstream benefit.[14, 19, 23, 35-40] These domains can diverge: a model may be fully traceable yet semantically unstable; a dataset may be legally accessible yet relationally extractive; a federated architecture may preserve locality while functionally repurposing data in place (Table 2).

Table 2. Domains of Ethical Continuity and the Proxies That Commonly Mislead Reviewers

Domain	What it tracks	Visible proxy that can mislead	Load-bearing review question
Provenance	Lineage, handling,	Audit trail,	What does traceability

Domain	What it tracks	Visible proxy that can mislead	Load-bearing review question
	and transformation	repository approval, ingestion log	establish—and what does it not establish—about meaning, authorization, and accountability?
Semantics	Descriptors, ontologies, labels, and inferred constructs	Stable variable names, legacy categories, harmonized labels	Do current categories still map onto the same constructs, persons, or groups?
Authorization	Permissions, expectations, and justificatory scope	Broad consent, legacy IRB approval, license terms	Is the present use still within the justified scope of entrustment?
Relational standing	Voice, recourse, and claim to benefit or contestation	Naming a source group, token consultation, diversity language	Who can object, shape conditions, or share in value if group-level claims or harms arise?

Operational Conditions for Relational Standing

The relational domain requires discipline: community standing should not be invoked whenever individual consent looks thin. It becomes governance-relevant when four conditions are met: (i) group-level claims are foreseeable from the proposed use; (ii) an affected collectivity is reasonably identifiable; (iii) likely harms or benefits are more than trivial; and (iv) some plausible route to representation exists or can be constructed.

The fourth criterion creates an apparent paradox: communities with weaker organizational infrastructure would, on a strict reading, fail to trigger relational review precisely because no plausible route to representation exists. Yet these are often the communities most vulnerable to extractive secondary data use. We propose that when the first three criteria are clearly met but the fourth is not, the absence of a representational route should itself be treated as a governance finding.

Indeterminacy should constrain rather than enable high-stakes downstream claims. Review bodies should impose precautionary limits on group-level claims, commercialization, or deployment in identifiably affected but representationally absent populations, and should flag the construction of representational pathways through community advisory boards, consortium-level governance, or standing partnerships with local research institutions as a condition for proceeding beyond precautionary constraints.[41]

Communitarian moral philosophy provides a distinctive normative foundation for this domain. In Metz's formulation, ubuntu ethics holds that moral status is constituted through relationships of identification, solidarity, and reciprocal recognition.[42] Molefe's account of relational personhood deepens this by arguing that moral value is realized through active participation in relationships of mutual recognition and shared benefit.[43] The implication for data governance is direct: if personhood is partly constituted through communal relationships, extracting value from community-generated data without reciprocal benefit violates not merely autonomy but the conditions of personhood itself. Relational standing therefore persists even when individual consent is formally in place, because what is at stake is whether communal relationships through which moral value is constituted remain intact, respected, and reciprocally sustained.[35, 42-45]

Institutional Mechanisms That Produce the Continuity Trap

The Continuity Trap is not an isolated lapse in judgment. It is produced by recurrent institutional mechanisms. Meyer and Rowan's foundational analysis of institutionalized organizations demonstrated that formal organizational structures often function as legitimacy-conferring ceremonies rather than instruments of substantive coordination—a process they termed "decoupling," in which visible compliance structure becomes detached from actual work practices it is supposed to govern.[22] When review bodies treat provenance logs, consent documentation, or model cards as sufficient grounds for approval, they engage in precisely this kind of ceremonial compliance: the formal structure of review is preserved, but substantive inquiry into semantic, authorization, and relational continuity is decoupled from the approval decision.

Four specific mechanisms drive this decoupling.

Provenance privilege arises because lineage is highly legible, serializable, and auditable. Provenance evidence receives disproportionate weight not because reviewers believe it is sufficient, but because it is the most readily available, standardized, and defensible in *post hoc* accountability inquiries. Regulatory frameworks such as HIPAA's data-use agreement infrastructure, GDPR's processing-record requirements, and repository accreditation standards reinforce this priority while providing comparatively little guidance on semantic validation, authorization scope assessment, or relational standing evaluation.[16, 46, 47]

Descriptor sedimentation occurs when old labels become embedded in forms, datasets, and analytic pipelines and acquire an appearance of naturalness that makes them resistant to re-examination. Population categories such as "African," "Black," "European," and "ancestry" are not interchangeable, but once encoded in variable names, metadata schemas, and pipeline documentation, they travel between studies, institutions, and jurisdictions without triggering the semantic review that their repurposing warrants.[36-39]

Authorization fossilization occurs when formal permissions outlive the ethical coherence of the uses they now justify. Broad-consent instruments, legacy IRB approvals, data-use agreements, and material transfer agreements were designed for particular horizons of use. When relied upon to authorize materially different downstream applications, e.g., iPSC derivation from legacy biospecimens, commercial licensing of genomic datasets, or training of predictive algorithms on clinical records, the authorization has fossilized: it retains formal legal standing but no longer tracks the normative expectations that grounded the original entrustment. This is the decoupling mechanism most directly anticipated by Brunsson's analysis of organizational hypocrisy.[35, 48]

Community effacement occurs when groups remain present as sources of diversity, legitimacy, or biological material but disappear as practical participants in downstream governance. Source communities may be named in publications, diversity statements, and grant applications, but have no mechanism for contesting derivative uses, shaping benefit-sharing arrangements, or influencing deployment decisions. The visible markers of community engagement are present; the relational infrastructure that would make contestation or reciprocal benefit possible is absent.[22, 35, 42-45]

Together, these four mechanisms explain why ethical governance of secondary data use can be simultaneously conscientious and insufficient. They are not independent pathologies but reinforcing elements of a single institutional logic: the logic of legibility-based compliance, in which the most auditable, standardized, and documentable aspects of governance receive disproportionate institutional weight, while the less tractable dimensions—shifting meanings, collective stakes, changing expectations, and relational obligations—are implicitly downgraded. The Continuity Trap is the moment that downgrade becomes a decision rule.

The Continuity Trap in Existing Data Science Governance Frameworks

The Continuity Trap is not a rival to contemporary data science or AI governance frameworks. It serves as a diagnostic supplement to them. The NIST AI Risk Management Framework, the National Academy of Medicine's AI Code of Conduct, the FDA's evolving regulatory perspective, and local health-system AI governance structures all emphasize lifecycle oversight, transparency, monitoring, evaluation, and accountability[23, 49-51]. But these frameworks operate at a higher level than the review episode where proxy closure may occur. They tell institutions to govern, measure, monitor, and engage impacted individuals. They do not by themselves tell a reviewer when a repository approval, model card, broad-consent instrument, or locality-preserving architecture is being allowed to terminate inquiry that should remain open. That is why even sophisticated institutions can be confident about their ethical governance of secondary data use but for the wrong reasons. Nong and colleagues showed that academic medical centers vary substantially in how they govern predictive AI under uncertainty.[23] The Continuity Trap explains one recurrent source of that variation: institutions are structurally drawn toward what is auditable, standardized, and administratively tractable.

Case Vignettes

The following four vignettes illustrate how the Continuity Trap operates across distinct data-stages and governance contexts. Each case is treated at greater length including full analysis within our Representational Veracity framework.[19] Here we focus narrowly on the inferential structure of the trap itself: which continuity domain provides the salient proxy signal, and which domains remain unexamined as a result.

Polygenic Risk Scores

Polygenic risk scores (PRS) illustrate why the Continuity Trap must be distinguished from the technical problem of transportability. Linkage disequilibrium structure, allele frequencies, effect-size heterogeneity, and environmental interaction explain much of the empirical challenge in moving PRS across populations.[38, 39] The Continuity Trap emerges at a different level: when a well-documented genome-wide association pipeline, ancestry adjustment, or validation workflow is treated as sufficient warrant to move the score into new populations or clinical settings without reopening semantic and relational questions.

The proxy closure mechanism is provenance-driven: a computationally reproducible pipeline with rigorous quality-control documentation creates a strong impression of governance continuity. But categories such as “African,” “European,” “Black,” “Yoruba,” and “ancestry” do not do identical work across contexts.[36-39] Computational and provenance continuity can be excellent while semantic continuity is weak: a PRS can be technically elegant yet organized by descriptors that travel poorly. And governance arrangements may give source communities little voice in downstream deployment, commercialization, or performance oversight - a relational standing failure that provenance documentation cannot reveal.

Legacy Biospecimens and iPSC Derivation

Legacy biospecimens transformed into induced pluripotent stem cell (iPSC) lines reveal the Continuity Trap in a limiting case. Here, material continuity is exceptionally strong: the derivative remains biologically traceable to the donor and retains genomic continuity. Precisely for that reason, reviewers may be tempted to regard the derivative as an ethically continuous extension of the original specimen. That inference is unsafe. Reprogramming changes what the material can become - organoids, disease models, gene-edited derivatives, screening platforms, and potentially commercially valuable products - thereby expanding the horizon of possible uses beyond what may have been contemplated at collection.[35]

It was the specific analysis of iPSC derivation that prompted the naming of the Continuity Trap concept in earlier work.[1] Once iPSC lines are derived, characterized, and distributed to repositories and research laboratories worldwide, the biological material enters a state of indefinite self-renewal. Unlike a dataset that can in principle be recalled, deleted, or access-restricted, a distributed iPSC line cannot be unilaterally withdrawn from global circulation. The cell line propagates, genomic information persists within it, and downstream derivatives multiply the number of biologically active entities carrying the donor’s genetic identity. After derivation, characterization, and distribution, ordinary withdrawal rights become progressively weaker as viable practical instruments for restoring donor or community control. This circumstance is important not because it proves the Continuity Trap is inevitable, but because it shows that sustained engagement, local scientific participation, and benefit-oriented governance can interrupt it.[35]

Federated Learning

Federated learning presents a different variant because the privileged continuity signal is locality.

The statement that raw data stay at the local institution is often true and privacy-relevant.[6, 7] But locality answers only one class of questions. It does not by itself settle whether clinical records generated for care may legitimately function as distributed training inputs for a different purpose; whether diagnostic categories are commensurable across sites; whether smaller institutions bear unequal burdens in producing a global model; or whether patients and contributing sites have meaningful recourse when downstream products are commercialized or deployed asymmetrically.[6, 7, 34]

The technical governance community has not ignored these concerns: differential privacy integration, secure aggregation protocols, gradient auditing techniques, and formal data-use agreements governing aggregation servers represent meaningful responses to privacy and confidentiality risks.[6, 7, 34] These mechanisms are real and important. The argument is rather that they are designed primarily to address confidentiality and privacy risks which lie in the provenance domain, but remain insufficient for the authorization and relational governance problems the Continuity Trap identifies. Differential privacy can bound information leakage from individual records, but it cannot determine whether the purpose for which clinical data are now being used falls within the scope of the trust under which those data were originally generated. In federated learning, the “data never moved” mantra, reinforced by increasingly sophisticated privacy-preserving mechanisms, can become a proxy for “the governance problem has been solved.” That is the Continuity Trap.[6, 7, 34]

Health-Related Large Language Models

Large language models intensify the Continuity Trap because the distance between input documentation and downstream influence is very difficult to govern. An important distinction applies: a model trained on a clearly governed, research-consented clinical dataset presents a different ethical governance problem from one trained retrospectively on a large, opaque, cross-border dataset that incidentally contains health-relevant information. The frontier governance challenge is the latter.[52-54]

In such settings, developers may be able to document broad categories of source material without being able to show how particular persons or communities influenced model behavior, how downstream harms can be traced back to source conditions, or what route exists for contesting use. Machine unlearning is an active technical field and may improve; yet even when partial unlearning is feasible, technical removability is not the same as governance adequacy.[52-54] The Continuity Trap lies in treating input documentation as if it were already an ethically sufficient account of authorization, relational standing, and accountability.

Toward Trigger-Based Continuity Review

We propose a rigorous but not bureaucratically inflationary response to Continuity trap. Continuity review should be triggered when one or more of the following conditions are present:

Box 1. Trigger Conditions for Continuity-Sensitive Review

- Irreversible or hard-to-reverse transformation of source data or material
- Repurposing of descriptors across populations, sites, ontologies, or historical periods
- Materially expanded uses, including commercialization or cross-border transfer
- Production of outputs likely to generate group-level claims, sensitive inferences, or deployment asymmetries
- Absence of a credible pathway for accountability, contestation, or benefit-sharing

Not all secondary use triggers continuity review. Routine reuse within the scope contemplated by the original data-access agreement, e.g., a new statistical analysis of a dataset already approved for the investigator's research program, a replication study using the same variables and population descriptors, or internal quality-improvement analyses that do not generate group-level claims or commercial outputs, would not ordinarily meet any trigger condition. The triggers are designed to identify cases where the data-stage has undergone a qualitative transformation in purpose, scope, or inferential reach, not cases where the same data are being used for closely related analytic ends within an existing governance authorization. The distinction is between quantitative extension (more analyses of the same kind, within the same authorization) and qualitative transformation (new inferential claims, new populations, new commercial pathways, or new derivative objects). Continuity review is triggered by the latter, not the former.

The Continuity Statement

For triggered cases, investigators or developers should provide a short continuity statement answering the following four questions:

Box 2: Continuity assessment survey

- Question 1. What is the present data-stage under review—source dataset, linked dataset, trained model, derivative, or deployment?
- Question 2. Which continuity domain is doing most of the justification for ethical governance continuity?
- Question 3. Which continuity domain is the least secure—provenance, semantics, authorization, or relational standing?
- Question 4. What feasible safeguard targets the endangered domain?

The safeguard that is recommended must match the endangered domain. If semantic continuity is weak, the response should include descriptor justification, uncertainty disclosure, or local validation before deployment—not simply a provenance certificate. If authorization continuity is weak, renewed notice, re-engagement, or explicit restrictions on derivative applications may be sufficient. If relational standing is weak, advisory review, joint oversight, benefit-sharing provisions, contestability pathways, or deployment limits are appropriate. If provenance itself is weak, reconstruction of lineage and transformation auditing are the relevant next steps.

The continuity statement is a structured disclosure, not a self-certification instrument. Investigators draft it, but reviewers evaluate it and retain the authority to disagree with the investigator's domain assessment. When a trigger condition has been met but the investigator asserts that all domains are adequately secure, reviewers should treat that mismatch itself as warranting scrutiny, similar to how a conflict-of-interest committee would scrutinize a nil disclosure from an investigator with substantial industry ties. Question 2 ("Which domain is doing most of the justification?") and Question 3 ("Which domain is the least secure?") require the investigator to articulate where the load-bearing governance work is concentrated, and a claim that no domain is under strain when a trigger condition is clearly met should itself raise a flag.

Dispute Resolution and Epistemic Gaps

When investigators and reviewers disagree about which domain is least secure, a *precautionary attention default* should apply: the domain that either party identifies as vulnerable should receive targeted attention in the continuity statement. This conservative default is justified because the cost of over-reviewing a domain that is actually secure is administrative inconvenience, while the cost of under-reviewing a domain that is actually weak is governance failure. This default does not function as a veto. If targeted review concludes that the flagged domain is adequately secured, the continuity statement documents that finding and the review proceeds.

Semantic continuity evaluation poses a distinct epistemic challenge. Traditional research ethics committees are well positioned to assess authorization and provenance but may lack the informatics, epidemiological, or sociological expertise needed to evaluate whether descriptors, ontologies, or population categories remain valid across data-stages. This gap does not require a new review body. It does require that data-access committees and AI governance bodies incorporate relevant domain expertise when continuity statements identify semantic continuity as the endangered domain. This eventuality is provided for in many ethics guidance documents.[55] In consortium settings, shared descriptor metadata repositories and harmonization protocols can reduce the evaluative burden on any single committee.[36, 37, 41]

Institutional Roles and Burden

Institutional roles must be assigned realistically. Developers and investigators should characterize the present data-stage, justify descriptors, and map downstream uses. Health systems, research ethics committees, and data-access committees should insist on local validation, disclosure of uncertainty, and contestability where semantic or relational continuity is weak. Repositories, funders, journals, and contracts can carry derivative-use restrictions, downstream reporting obligations, and, in some cross-border settings, benefit-sharing or joint-governance conditions.[6, 7, 16, 34, 44, 45] No single committee has plenary authority in adjudicating this. Continuity-sensitive governance works only if duties are distributed across the institutions and bodies that actually control access, deployment, and value capture.

The continuity statement itself should be a concise document of one to three pages structured around the four questions above, submitted alongside existing ethics review or data-access applications. Evaluation should fall to the existing institutional body with jurisdiction over the relevant access decision: the IRB, data-access committee, material transfer authority, or institutional AI governance body. We estimate that a well-structured continuity statement would add approximately two to four hours of preparation time for investigators and one to two hours of evaluation time for reviewers. These are burdens comparable to reviewing a data-management plan and substantially less than a full re-review.

This burden-sensitive approach matters particularly in low-resource settings where overextended research ethics committees should not be asked to perform full sociological reconstruction for every secondary-use request. Triggered review should therefore be parsimonious, use existing forms where possible, rely on consortium-level descriptor metadata, and prefer standing advisory structures over *ad hoc* consultation whenever feasible. In settings where institutional capacity is limited, consortium-level governance mechanisms can distribute the burden across participating institutions rather than concentrating it at any single review site.

Conclusion

The Continuity Trap defines and names a type of failure in the ethical governance of data science health research where institutions infer too much from what they can most easily see. Provenance logs, repository approvals, consent documents, model cards, and locality-preserving architectures all matter, but none of them alone can certify legitimacy. Ethical continuity in data science health research is relation-specific across provenance, semantics, authorization, and relational standing. Those relations can diverge even in technically sophisticated, well-documented systems.

Naming the trap is useful because it identifies a correctable institutional error. It explains how governance can be simultaneously conscientious and insufficient, technically advanced and normatively stale. The policy implication is not a new bureaucracy for every reuse, but a trigger for renewed inquiry whenever visible continuity is likely to be overread. Data science health research will often be most vulnerable not when documentation is absent, but when documentation is so reassuring that institutions stop asking the harder questions. Provenance should therefore be treated as the beginning of ethical review rather than the end of it.

References

1. Adebamowo C, Callier S, Maduka OC, Akintola S, Kukucka J, Arima CG. Ethical Oversight of Data Science Health Research in Africa. NEJM AI. 2024;1(6):AIpc2400033. doi: 10.1056/AIpc2400033.

2. Adebamowo CA, Callier S, Akintola S, Maduka O, Jegede A, Arima C. The promise of data science for health research in Africa. Nat Commun. 2023;14(1):6084. doi: 10.1038/s41467-023-41809-2.

3. Goldstein BA, Navar AM, Carter RE. Moving beyond regression techniques in cardiovascular risk prediction: applying machine learning to address analytic challenges. Eur Heart J. 2017 Jun 14;38(23):1805–14. PMID: 27436868. doi: 10.1093/eurheartj/ehw302.

4. Johnson AE, Ghassemi MM, Nemati S, Niehaus KE, Clifton DA, Clifford GD. Machine Learning and Decision Support in Critical Care. Proceedings of the IEEE Institute of Electrical and Electronics Engineers. 2016 Feb;104(2):444–66. PMID: 27765959. doi: 10.1109/JPROC.2015.2501978.

5. Litjens G, Kooi T, Bejnordi BE, Setio AAA, Ciompi F, Ghafoorian M, et al. A survey on deep learning in medical image analysis. Medical image analysis. 2017 Dec;42:60–88. PMID: 28778026. doi: 10.1016/j.media.2017.07.005.

6. Pati S, Kumar S, Varma A, Edwards B, Lu C, Qu L. Privacy preservation for federated learning in health care. Patterns (New York). 2024;5(7):100974.

7. Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. NPJ Digit Med. 2020;3:119. PMID: 33015372. doi: 10.1038/s41746-020-00323-1.

8. Kaye J, Whitley EA, Lund D, Morrison M, Teare H, Melham K. Dynamic consent: a patient interface for twenty-first century research networks. Eur J Hum Genet. 2015 Feb;23(2):141–6.

PMID: 24801761. doi: 10.1038/ejhg.2014.71.

9. Wilkinson MD, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, et al. The FAIR Guiding Principles for scientific data management and stewardship. Sci Data. 2016 Mar 15;3:160018. PMID: 26978244. doi: 10.1038/sdata.2016.18.

10. Mitchell M, Wu S, Zaldivar A, Barnes P, Vasserman L, Hutchinson B, et al., editors. Model cards for model reporting. Proceedings of the conference on fairness, accountability, and transparency; 2019.

11. Yeung K. Algorithmic regulation: A critical interrogation. Regulation & Governance. 2018 2018/12/01;12(4):505–23. doi: https://doi.org/10.1111/rego.12158.

12. Power M. The audit society: Rituals of verification: OUP Oxford; 1997. ISBN: 0191638927.

13. Grady C, Eckstein L, Berkman B, Brock D, Cook-Deegan R, Fullerton SM, et al. Broad Consent for Research With Biological Samples: Workshop Conclusions. Am J Bioeth. 2015;15(9):34–42. PMID: 26305750. doi: 10.1080/15265161.2015.1062162.

14. Nissenbaum H. Privacy as contextual integrity. Washington Law Review. 2004;79(1):119–58.

15. Rocher L, Hendrickx JM, de Montjoye YA. Estimating the success of re-identifications in incomplete datasets using generative models. Nat Commun. 2019 Jul 23;10(1):3069. PMID: 31337762. doi: 10.1038/s41467-019-10933-3.

16. Nong P, Hamasha R, Singh K, Adler-Milstein J, Platt J. How academic medical centers govern AI prediction tools in the context of uncertainty and evolving regulation. Massachusetts Medical Society; 2024. p. AIp2300048.

17. Elliott A. Better, broader, safer: using health data for research and analysis (the Goldacre review). IOP Publishing; 2022. p. 030101.

18. Steinsbekk KS, Kåre Myskja B, Solberg B. Broad consent versus dynamic consent in biobank research: is passive participation an ethical problem? European Journal of Human Genetics. 2013;21(9):897–902.

19. Adebamowo AC, Adebamowo S, Adeyemo O, Ikhane P, Ogundiran T, Callier S, et al. A Persistence-Based Ethical Governance Framework for Data Science Health Research. SSRN. 2026 Apr 03 2026.

20. Haggerty KD. Ethics creep: Governing social science research in the name of ethics. Qualitative sociology. 2004;27(4):391–414.

21. Floridi L. Soft ethics, the governance of the digital and the General Data Protection Regulation. Philos Trans A Math Phys Eng Sci. 2018 Oct 15;376(2133). PMID: 30322997. doi: 10.1098/rsta.2018.0081.

22. Meyer JW, Rowan B. Institutionalized Organizations: Formal Structure as Myth and Ceremony. American Journal of Sociology. 1977;83(2):340–63.

23. Nong P, Hamasha R, Singh K, Adler-Milstein J, Platt J. How academic medical centers govern AI prediction tools in the context of uncertainty and evolving regulation. NEJM AI. 2024;1(3):AIp2300048. doi: 10.1056/AIp2300048.

24. Bosk CL, De Vries RG. Bureaucracies of mass deception: Institutional review boards and the ethics of ethnographic research. The Annals of the American Academy of Political and Social Science. 2004;595(1):249–63.

25. Mattson C, Bushardt RL, Artino Jr AR. When a measure becomes a target, it ceases to be a good measure. Journal of Graduate Medical Education. 2021;13(1):2–5.

26. Poku M. Campbell’s law: implications for health care. Journal of health services research & policy. 2016;21(2):137–9.

27. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447–53. doi: 10.1126/science.aax2342.

28. Lemmon J, Guo LL, Posada J, Pfohl SR, Fries J, Fleming SL, et al. Evaluation of Feature Selection Methods for Preserving Machine Learning Performance in the Presence of Temporal

Dataset Shift in Clinical Medicine. Methods of information in medicine. 2023 May;62(1-02):60–70. PMID: 36812932. doi: 10.1055/s-0043-1762904.

29. Suresh H, Guttag J. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. Proceedings of the 1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization; --, NY, USA: Association for Computing Machinery; 2021. p. Article 17.

30. Raji ID, Smart A, White RN, Mitchell M, Gebru T, Hutchinson B, et al. Closing the AI accountability gap: defining an end-to-end framework for internal algorithmic auditing. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency; Barcelona, Spain: Association for Computing Machinery; 2020. p. 33–44.

31. Strathern M. 'Improving ratings': audit in the British University system. European Review. 1997;5(3):305–21. doi: 10.1002/(SICI)1234-981X(199707)5:3<305::AID-EURO184>3.0.CO;2-4.

32. John YJ, Caldwell L, McCoy DE, Braganza O. Dead rats, dopamine, performance metrics, and peacock tails: Proxy failure is an inherent risk in goal-oriented systems. Behavioral and Brain Sciences. 2024;47:e67. doi: 10.1017/S0140525X23002753.

33. Higgins WC, Gillett AJ, Kaplan DM, Ross RM. Proxy failure and poor measurement practices in psychological science. Behavioral and Brain Sciences. 2024;47:e75. doi: 10.1017/S0140525X2300287X.

34. Suver C, Thorogood A, Doerr M, Wilbanks J, Knoppers B. Bringing Code to Data: Do Not Forget Governance. Journal of medical Internet research. 2020 Jul 28;22(7):e18087. PMID: 32540846. doi: 10.2196/18087.

35. Ikhane PA, Yusuf T, Adeyemo O, Ogundiran TO, Adebamowo SN, Adebamowo CA. Community and bioethicists' perspectives on iPSC research with biobanked samples collected using broad consent. Stem Cell Reports. 2025 Nov 20:102721. PMID: 41270747. doi: 10.1016/j.stemcr.2025.102721.

36. Khan AT, Adebamowo C, Fullerton SM, Hirbo J, Konigsberg IR, Kraft P. A data model for population descriptors in genomic research. American journal of human genetics. 2025;112(7):1504–14.

37. Smith JL, Adebamowo CA, Adebamowo SN, Darst BF, Fullerton SM, Gogarten SM, et al. Recommendations for responsible use of population descriptors in polygenic risk score development. Nature Genetics. 2025 2025/12/01;57(12):2962–71. doi: 10.1038/s41588-025-02395-9.

38. Kachuri L, Chatterjee N, Hirbo J, Schaid DJ, Martin I, Kullo IJ. Principles and methods for transferring polygenic risk scores across global populations. Nature Reviews Genetics. 2024;25(1):8–25.

39. Martin AR, Kanai M, Kamatani Y, Okada Y, Neale BM, Daly MJ. Clinical use of current polygenic risk scores may exacerbate health disparities. Nature Genetics. 2019;51(4):584–91. doi: 10.1038/s41588-019-0379-x.

40. Suresh H, Guttag J, editors. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. Proceedings of the 2021 ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization; 2021; New York, NY: ACM.

41. Tindana PO, Singh JA, Tracy CS, Upshur REG, Daar AS, Singer PA. Grand challenges in global health: community engagement in research in developing countries. PLOS Medicine. 2007;4(9):e273. doi: 10.1371/journal.pmed.0040273.

42. Metz T. Toward an African Moral Theory. Journal of Political Philosophy. 2007 2007/09/01;15(3):321–41. doi: https://doi.org/10.1111/j.1467-9760.2007.00280.x.

43. Molefe M. An African Philosophy of Personhood, Morality and Politics2019. ISBN: 978-3-030-15561-2.

44. Salami AO. Artificial intelligence, digital colonialism, and the implications for Africa’s future development. Data & Policy. 2024;6:e67. doi: 10.1017/dap.2024.75.

45. Couldry N, Mejias UA. The costs of connection: How data is colonizing human life and

appropriating it for capitalism. The costs of connection: Stanford University Press; 2019.

46. Anderson DR. Health and Human Services Standards for Privacy of Individually Identifiable Health Information. US Att'ys Bull. 2002;50:16.

47. Cohen IG, Mello MM. HIPAA and Protecting Health Information in the 21st Century. Jama. 2018 Jul 17;320(3):231–2. PMID: 29800120. doi: 10.1001/jama.2018.5630.

48. Brunsson N. The organization of hypocrisy: Samfundslitteratur Frederiksberg; 2019. ISBN: 8763003724.

49. AI N. Artificial intelligence risk management framework (AI RMF 1.0). URL: https://nvlpubs nist gov/nistpubs/ai/nist ai. 2023:100–1.

50. Adams L, Fontaine E, Matheny M, Krishnan S. An Artificial Intelligence Code of Conduct for Health and Medicine. 2025.

51. Singh V, Cheng S, Kwan AC, Ebinger J. United States Food and Drug Administration regulation of clinical software in the era of artificial intelligence and machine learning. Mayo Clinic Proceedings: Digital Health. 2025;3(3):100231.

52. Bender EM, Gebru T, McMillan-Major A, Shmitchell S, editors. Proceedings of the 2021 ACM conference on fairness, accountability, and transparency; 2021.

53. Bourtoule L, Chandrasekaran V, Choquette-Choo CA, Jia H, Travers A, Zhang B, et al., editors. Machine unlearning. 2021 IEEE symposium on security and privacy (SP); 2021: IEEE.

54. Wang W, Tian Z, Zhang C, Yu S. Machine unlearning: A comprehensive survey. arXiv preprint arXiv:240507406. 2024.

55. National Code for Health Research Ethics, (2006).

Acknowledgments

The authors acknowledge this study’s participants, the Data Science for Health Discovery and Innovation in Africa Consortium (DS-I Africa), and members of the BridgELSI (Bridging Gaps in the Ethical, Legal, and Social Implications of Data Science Health Research in Nigeria) project. Funding for this project was received from the US National Institutes of Health (BridgELSI, U01 MH127693) grant. Additional support for CA and SNA was received from the Maryland Department of Health’s Cigarette Restitution Fund Program (CH-649-CRF) and the University of Maryland Greenebaum Comprehensive Cancer Center Support Grant (P30CA134274). The funding agencies did not play any role in the publication.

Data Availability

No primary data were generated or analyzed in this study. All sources cited are publicly available.

Authors’ Contributions

Conceptualization – CA

Formal analysis – CA (lead), SNA (equal).

Funding acquisition – TO (lead), CA (equal).

Conceptual development – CA (lead), AA (supporting), PI (supporting), SNA (supporting), SA (supporting), TO (supporting), AJ (supporting), OA (supporting), SC (supporting), MI (supporting), IU (supporting), The BridgELSI Project as part of the DS-I Africa Consortium (supporting).

Methodology – CA

Project administration – CA (lead), TO (supporting), SNA (supporting).

Resources – CA (lead), TO (supporting), The BridgELSI Project as part of the DS-I Africa Consortium (supporting).

Supervision – CA (lead), TO (supporting), SA (supporting).

Validation – CA (lead), SNA (supporting).

Writing – original draft – CA (lead) SNA (supporting).

Writing – review & editing – AA (supporting), PI (supporting), SA (supporting), TO (supporting), AJ (supporting), OA (supporting), SC (supporting), MI (supporting), IU (supporting), The BridgELSI Project as part of the DS-I Africa Consortium (supporting).

Conflicts of Interest

None declared.

Abbreviations

LMIC: low- and middle-income country

Supplementary Files

Untitled.
URL: http://asset.jmir.pub/assets/386cbaac08b6ccda94266045378b89d4.pdf

TOC/Feature image for homepages

The Continuity Trap illustrated.