PREPRINT · OCTOBER 2026 · CONTESTABILITY.AI

From Notice to Contestable Grounds: Necessary Informational Conditions for Algorithmic Contestability

Kerry Thornhill

Graduate Student, Interdisciplinary Studies
Anuradha and Vikas Sinha College of Artificial Intelligence and Advanced Analytics
University of North Texas

EMAIL kerrythornhill@my.unt.edu
WEB contestability.ai
ORCID 0009-0007-8936-0019

Abstract

Notice, explanation, and an appeal procedure can leave an affected party unable to investigate an institution's determination. Where contestability is independently required or promised, this article develops the Necessary Informational Conditions framework to identify what information must enable. It connects a material objection to the question that must be answered, the operations that could answer it, and a test of whether any adequate route can dispense with a candidate operation. Six functions organize the inquiry: determination and consequence; sociotechnical roles and evidence control; data, source, and quality; governing-standard-operation fit; evidential warrant, uncertainty, and threshold; and contextual applicability. Necessity is conditional on the task, an independently justified adequacy floor, and the feasible routes compared. Different documents or expert arrangements may perform the same operation; an adequate route that omits it defeats its necessity. The framework also specifies arrangements for making the work possible: baseline orientation, contemporaneous preservation, inquiry-activated access, protected mediation, and usable support. These must let affected parties discover what warrants investigation before they are expected to substantiate an objection. Meaningful contestability also requires timely institutional consideration through which evidence and reasons can affect the determination. K.W. v. Armstrong illustrates how obstacles to inquiry and review can interact. A proposed research program tests omission and substitution, compares rival classifications, and develops a contestability toolkit. The contribution is a common framework for diagnosing failures of inquiry, comparing remedies, and designing procedures through which affected people and their representatives can investigate consequential judgments and bring evidence and reasons to bear on them.

Keywords: algorithmic contestability; informational conditions; institutional decision-making; evidentiary access; accountability; functional necessity

1. Notice without contestable grounds

The notices in *K.W. v. Armstrong* gave participants in Idaho's Medicaid Developmental Disabilities Waiver program budget amounts calculated by a statistical budget tool from assessment data, but did not explain the decreases sufficiently for participants to prepare a response (*K.W. ex rel. D.W. v. Armstrong*, 789 F.3d 962, 967, 971, 974 [9th Cir. 2015]). The Ninth Circuit affirmed the extension of a preliminary injunction forbidding such reductions without adequate notice (789 F.3d at 966). On summary judgment in 2016, the district court addressed flawed data, undefined review criteria, and restrictions on using assessment materials in appeals (180 F. Supp. 3d 703 [D. Idaho 2016], official slip op. at 11–12, 14–15, 22–24).

The participants knew that their budgets had been reduced, but they lacked a usable way to establish why and whether the reduction could be defended. This article explores that problem. I use effective incontestability to describe a determination that remains formally open to challenge, yet practically beyond meaningful challenge, because the affected party cannot adequately investigate a material objection or obtain a reasoned response. The *K.W.* record documents obstacles capable of producing a condition of effective incontestability. The litigation itself also shows that a further route to challenge was available; it does not establish that ordinary review was adequate.

An institution's determination does not become meaningfully contestable merely because the system behind it is transparent, its output explained, or a human exists somewhere in the decision loop. Open access to a model may show how it computes a result while still leaving the affected person unable to check the data used in their case or the standard the institution applied. An explanation can be intelligible even when its premises remain unverifiable. A human reviewer may repeat the output without investigating either the premises or the determination. In each case, the institution can retain control of the grounds needed to assess the determination, while the affected person bears the burden of contesting it. This is more than background unfairness. It compounds an existing asymmetry in authority, information, and investigative capacity. That brings us to our question: What must remain investigable when a fallible computational assessment acquires authority over a person's options?

I propose investigative capability as the organizing principle for assessing informational adequacy. The informational requirements of contestability can be specified through the work needed to investigate material objections. A common framework can guide that work across institutions and decision domains with different standards, evidence, and determination processes, provided its claims survive testing in those settings. A critical question remains: once law, policy, contract, professional norm, or an institution's own commitment independently requires or promises an opportunity to contest, which informational operations must remain available to the affected party or an authorized or claimant-directed representative?

The question of what powerful institutions owe those who must live with their decisions reaches into larger disputes over power, dependence, and whose judgment counts. I take up that broader

question in a companion article, *Private Power without Epistemic Preemption: Relational Institutional Answerability in Algorithmic Decision-Making*, currently in preparation. That article develops the normative theory of Relational Institutional Answerability: institutions exercising consequential private power over people through computational systems must not convert epistemic advantages into epistemic preemption but must keep their determinations open to informed challenge and consequential revision. This paper specifies what meaningful contestability requires; the companion article argues why private institutions owe it.

The Necessary Informational Conditions (NIC) framework combines a register of six candidate informational functions, a test of conditional necessity, and a delivery architecture. That architecture helps people discover grounds for challenge, preserves the records needed to investigate them, and provides access, protected mediation, and practical support. People should not have to prove a hidden defect before they can obtain the information needed to investigate it. Meaningful contestability also requires an institutional response: a forum that answers challenges with reasons and can secure correction or warranted reaffirmation.

The article's original contribution is an approach that connects the grounds of a challenge, the investigative work it requires, and the arrangements that make that work possible and consequential. The NIC framework gives affected parties and their representatives a basis for identifying what to investigate and request. It gives institutions and researchers a basis for designing and evaluating the routes through which those inquiries can proceed. The proposed research program combines empirical tests of necessity and transportability with the development of a contestability toolkit to enable affected parties to make informed challenges while preserving warranted reliance on automated assistance.

Sections 2 and 3 establish the literature, object, and standpoint; Section 4 develops necessity, Section 5 delivery, and Section 6 the functions; Section 7 returns to *K.W.* to show how barriers to contestation interact; Section 8 states objections and boundaries; Section 9 sets out tests of omission, substitution, and portability to make the register more robust across institutions and contexts of use.

2. From explanation to investigative capability

Contestability scholarship has already made clear that merely explaining a decision is not enough to make it open to meaningful challenge. Drawing on due process, Kaminski and Urban propose an individual right to contest AI decisions alongside systemic governance and set out four contestation archetypes (2021, 1957). Analyzing submissions to the public consultation on Australia's proposed AI Ethics Framework, Lyons, Velloso, and Miller found calls for notice of decisions and how to challenge them, explanations, accessible review procedures, and clarity about who is accountable (2021, 106:12–106:14). Alfrink et al. treat contestability as a design commitment: AI systems must support scrutiny and intervention throughout their lifecycle, and they identify features and practices that serve that commitment (2023, 614, 628). This article picks up from there and develops a portable framework to test informational necessity: a way to

distinguish, across institutional settings, operations that help people investigate a decision from those an adequate investigation cannot do without.

A portable framework must explain both what changes across settings and what carries over. Ploug and Holm develop informational requirements for contesting AI-supported health-care decisions and ask how those requirements extend to other fields (2020, 2–4). Bayamlıoğlu connects different grounds for contestation to the information and safeguards needed to pursue them (2018, 439–43; 2022, 1059). NIC builds on these accounts by making investigative operations the basis of comparison. Standards, evidence, and professional responsibilities vary across settings; portability depends on identifying the work an adequate investigation requires and how different arrangements can support it.

Cobbe, Lee, and Singh locate reviewability in the wider technical and organizational process, addressing version-specific recording, disclosure for critical understanding, and access by affected parties (2021). Their account makes clear why generic documentation of a system is not enough to review a particular use of it. For the affected party, the question is which grounds an objection requires, how they can be made usable, and whether any that were not retained can be adequately reconstructed.

Legal scholarship connects informational access to the practical work of challenging a decision. Fraser, Snoswell, and Simcock examine how explanations of AI decisions can help claimants establish their case in tort litigation (2022, §1.2). Stewart examines where contestation fails and how affected parties can test the basis of a decision (2026). Cen, Ismael, and Zheng propose a test for deciding when an evidentiary inquiry requires access to particular resources and when a substitute will suffice (2026). These questions inform NIC's focus on the informational operations an adequate investigation requires.

Contestability depends on how an inquiry begins and what happens when it produces a challenge. Springer and Whittaker explore how people can control the timing and complexity of transparency information (2020, 29:24–29:26, 29:28). NIC asks what must already be available for someone to recognize grounds for investigation: people cannot request information about a problem they have no way to discover. Bovens explains accountability through the obligation to answer to a forum that can question and judge conduct (2007, 450–51), while Pi et al. examine how organizations respond to accountability demands (2026). Making information usable for investigation addresses one part of this problem; meaningful contestation also requires someone with the responsibility and authority to act on what the inquiry reveals.

An institution can explain a decision. An institution can disclose records. An institution can reveal its automated decision-making model and provide a formal, human-in-the-loop appeal path. Still, people affected by a consequential decision may be left unable to investigate whether that decision is justified. The Necessary Informational Conditions (NIC) framework provides a portable method for exposing and correcting the informational failures that make decisions effectively incontestable and allow injustice to persist.

Its distinctive move is to organize contestability around epistemic tasks: what a person must be able to discover, test, or establish to assess a material objection. This shifts the focus from the information provided to the judgments that information enables. NIC connects each task to the informational operations that support it, then tests whether adequate inquiry remains possible without them. That makes informational necessity testable and provides a common basis for identifying indispensable capabilities and evaluating alternatives. A request process fails if people cannot discover what to ask; a substitute fails if it cannot support the assessment the challenge requires.

Wherever contestability is required or promised, institutions must make their decisions answerable to investigation. Control over the information needed for a challenge can become control over which objections are discoverable, which can be tested, and which injustices remain beyond reach. NIC makes that power open to examination by identifying the investigative capabilities people need and the conditions that sustain or defeat them. A decision that survives because people cannot investigate it has escaped scrutiny; its survival does not show that it was justified.

3. Object, standpoint, and scope

The object of contestation is a determination. An output becomes a determination when an institution adopts it or uses it to shape a person's treatment, opportunities, burdens, or standing. Calling it advisory does not change this. It may still frame an investigation, set a default, narrow a reviewer's options, or shift the burden of securing a different result. Institutional use gives the output its force. An unused experimental output ordinarily falls outside this account. A determination is also broader than a claim. A claim states a proposition. A calculated budget combines data, a model, policy rules, thresholds, and judgments about need. Any of these elements or their application to the case may be challenged.

Adequacy is judged from the affected party's position. Can that person or a representative they authorize or direct turn reasonable doubt into an informed assessment or challenge? Institutional possession of the records is not enough. The person or an appropriate intermediary must be able to examine and use what the inquiry requires. A clinician, advocate, lawyer, or other specialist may do the investigative work. The affected party must be able to set that work in motion and bring its findings into the process governing the determination. Appointing a representative does not remove barriers the representative also faces.

Not every doubt supports a material objection. A material objection, or contestational ground, is a reason that would matter to the determination if substantiated. It could warrant correction, further inquiry, suspension, reconsideration, or a change in how the determination is weighted or used. It may challenge a conclusion, how it was reached, or whether it applies to the case. It may also challenge a classification, governing standard, or allocation. Such a challenge need not be an epistemic defeater: a reason against a belief or its justification. Dissatisfaction alone is not enough. Neither is curiosity that could make no rational difference. But materiality does not

depend on what the affected party can already demonstrate. A plausibility threshold that people cannot meet without access to hidden grounds can block the very inquiry needed to establish them.

Meaningful contestability is a practical capability with institutional force. An affected person, acting directly or through an authorized or claimant-directed representative, must be able to recognize a determination, discover material grounds for challenge, and investigate the determination as made. They must be able to present reasons or counterevidence in time to matter and receive a reasoned response through a process empowered to secure correction, modification, suspension, systemic change, or warranted reaffirmation. A procedure that leaves people unable to investigate a determination and assess its grounds does not give them the capability to contest it. Cost, inaccessible language, disability barriers, retaliation risk, fragmentation, expertise demands, and delay can prevent people from exercising that capability. A forum that receives objections but cannot act on them supplies no institutional force. Transparency, explanations, disclosures, and human review by themselves do not establish meaningful contestability. Informational necessity must be tested against this capability as a whole: the epistemic work needed to investigate and assess a determination, the practical means to carry it out, and an institutional process able to act on its findings.

4. Conditional functional necessity

Necessity can concern what makes a process meaningfully contestable, what an inquiry needs, what an outcome depends on, what the law requires, or when, to whom, and in what form information must be available. My focus is conditional functional necessity: whether an epistemic task can be adequately performed under stated conditions without a particular operation or its faithful equivalent. The task must be justified independently of the operation being tested. Otherwise, the claim assumes what it is meant to establish. Better outcomes from added information do not by themselves establish either causal or functional necessity.

Rawls distinguishes a liberty from its worth to its holder (1999, §32, 179). An opportunity to contest can likewise be granted without being made usable. The capability defined in Section 3 requires institutional uptake while leaving open the organization of informational functions, the form of disclosure, and who performs the technical work.

A conditional claim begins with a material objection within the regime's scope, the epistemic task needed to assess it, and an independently justified standard of adequacy. Consider an income-based denial. A claimant may ask whether they were entitled to the benefit and what grounds the institution used to deny it. The first concerns the correct disposition; the second concerns the original exercise of authority. Tracing the income value used by the system is one candidate operation for the second inquiry. The task is to establish enough about the original grounds to assess the alleged defect, rather than demand that particular trace. The comparison must reflect practical constraints and feasible routes. The claimant need not already know what went wrong; inquiry may be needed to establish whether a suspected mismatch matters.

The no-substitute test compares ways of completing the same task, including records, queries, reconstruction, audits, and expert processes, used separately or together. At least one feasible route performing the candidate operation must meet the adequacy standard. If none within the comparison can do so while omitting that operation and its faithful equivalents, the operation is conditionally necessary. An adequate alternative defeats the claim even if it is less informative or convenient. An expert who traces the same disputed income value through protected records, however, preserves the operation while changing its delivery. This is a faithful equivalent, or faithful substitute: the distinction follows the investigative work, whoever performs it and whatever records they use.

Specifying the comparison in advance makes it accountable, but a later-demonstrated adequate route still counts. The necessity claim fails wherever that route works. A narrower claim may remain defensible elsewhere, but the original finding must stand. Changing the task or adequacy standard, relabeling the alternative, or appealing to an untested substitute cannot undo the result.

A fresh determination can settle entitlement without settling the historical question. Independently verified income at the original decision date, assessed under the binding eligibility rule, might establish that the denial must be reversed. It would not establish which income value the institution used, how a disputed value entered the record, or whether the error could affect another decision. Tracing the original value may therefore be dispensable for deciding entitlement without removing the separate need to investigate the original grounds.

Those grounds matter because institutions can shape decisions upstream by controlling what counts as relevant evidence or acceptable knowledge. Consider an insurer with data no claimant can access. It may use that advantage to define counterfactuals that favor lower payouts. The claimant sees the denial; the institution controls the knowledge used to produce it. Its incentives can enter the assessment before anyone affected can investigate how it works. A successful appeal changes one result. It need not expose the choices that produced it.

Where upstream choices supply a material ground of challenge, historical inquiry must reach the choices governing the original determination. Otherwise, an institution can change the result while keeping its exercise of authority beyond examination. Those grounds may matter to past loss, discriminatory treatment, recurring error, or the institution’s defense of its conduct. A claimant need not prove such defects before being allowed to investigate them. Where contestation extends to the determination as made, a new answer cannot resolve questions about the original grounds. This specifies what the opportunity to contest must cover once its independent basis is established.

Historical reconstruction requires enough evidence about the original grounds to assess the objection, without demanding every record or a complete replay. An alternative that answers the historical question while omitting the candidate operation still defeats that operation’s necessity. The historical task and its adequacy standard should be specified at the outset, alongside any

entitlement question. Epistemic fidelity concerns whether the evidence reaches the use under challenge; it does not make a particular document indispensable.

Practical adequacy adds timeliness, independence, accessibility, affordability, and support for understanding and using findings. A faithful examination available only after consequential review closes may come too late to matter. A formal entitlement may need material and institutional support to become an effective opportunity (Nussbaum 2003, 37–39). Uptake concerns whether the findings receive consideration through a process with authority to respond. Even an open hearing can be undermined by an unjust credibility or intelligibility deficit (Fricker 2013, §2).

The six functions are a proposed organization of this work, not a classification established by the definition of contestability. They may overlap or leave gaps, and their presence alone does not guarantee meaningful contestability. Their value must be tested against other ways of organizing the same operations: do they improve the questions asked, evidence sought, or corrective routes identified? One capability may remain available while another is blocked. Persistent overlap may justify merging functions; recurring differences may justify splitting them; overlooked objections may require additions or a narrower scope. A case activating only three functions shows a particular profile within the register. Reorganizing the categories leaves an operation's necessity unchanged when the task, operation, adequacy standard, and comparison conditions remain the same. An adequate route that omits the operation defeats the claim however the register is organized.

5. Staged availability of contestable grounds

Specifying required operations leaves open when, to whom, and in what form grounds must be supplied. A request-only system can leave people unable to discover what to ask. Giving everyone everything can overwhelm recipients, bury relevant grounds, and expose protected information. Generic documentation or the current model may not reveal how an earlier decision was made. Where contestability is required or promised, delivery must address these obstacles. Its justification comes from the investigative work it enables, rather than from the sixfold register alone.

A challenge-generating baseline should identify the operative determination, its practical consequence, the adopting institution, automation's material role, and a usable next step. It should also provide cues about data, criteria, uncertainty, or deployment limits where these help make grounds for inquiry recognizable. Success depends on whether the affected person, with assistance where needed, can recognize and pursue a concern. Discovering grounds for challenge is itself part of the informational requirement.

Where inquiry concerns the determination as made, the relevant grounds must remain available or be faithfully reconstructable. Preservation should cover the inputs, rules, model or procedure versions, thresholds, validation materials, and human interventions needed for foreseeable

inquiries. Cobbe, Lee, and Singh connect these records to disclosure suited to critical understanding (2021). The 2014 class-wide injunction in *K.W.* required the agency, on request, to provide unaltered copies of the budget spreadsheets actually used for a participant, including retained versions used for that participant's earlier budgets (*K.W. v. Armstrong*, No. 1:12-cv-00022-BLW, Dkt. 130, at 31–32 [D. Idaho filed Mar. 25, 2014]). Retention should be proportionate to the tasks it supports. Preserving records does not by itself make them accessible or usable.

Access must be sufficient to investigate the concern. A plausible mismatch, conflicting source, exceptional circumstance, or uncertainty about applicability can raise questions about relevance as well as accuracy. Institutions may ask for a reasonable question before producing costly or protected material, but cannot require proof that depends on obtaining it. Access can deepen as the inquiry develops, with thresholds that remain crossable without that proof. Requests and decisions to grant, limit, or refuse access should be recorded. Restrictions and refusals, including findings that a request is irrelevant, must give reasons and remain open to review.

Redaction, secure inspection, and protected mediation can enable inquiry while protecting privacy, security, test integrity, and proprietary interests. Mediation requires a competent, sufficiently independent intermediary who examines the decision-specific grounds and returns findings the affected party can invoke. Representation achieves little if the representative faces the same inaccessible evidence. In 2016, the *K.W.* court required a plan allowing participants to view assessment portions needed to challenge a reduction and present challenged portions on appeal. It noted that protective orders and limited access could address copyright concerns (*K.W.*, 180 F. Supp. 3d at 717). Under *CK v. Magistrat der Stadt Wien*, a controller claiming that information due under GDPR Article 15(1)(h) contains protected third-party data or trade secrets must supply it to the competent supervisory authority or court. That body balances the interests and determines the extent of the data subject's access (C-203/22, ECLI:EU:C:2025:117, ¶¶72–76). These examples support protected routes to access without establishing a general entitlement to model internals. Any expert substitute must still demonstrate its adequacy for the inquiry.

Formal access may still be unusable. Translation, accessible formats, synthesis, representation, and sufficient time may be needed to overcome the barriers identified in Section 3. The test is whether the party can use the grounds to assess a concern and, where warranted, present an informed objection in time. A competent, sufficiently independent forum must also be able to obtain needed evidence and answer the objection with reasons through a process capable of correction or warranted reaffirmation. This institutional uptake lies beyond the six informational functions, but meaningful contestability depends on it.

Baseline orientation, preservation, inquiry-activated access, protected mediation, and usable support form the delivery architecture. They overlap and need not follow a fixed sequence: preservation begins before inquiry, and one answer may raise another question. Their adequacy depends on the operation at stake. Abundant records cannot compensate for a route that leaves a required assessment inaccessible. A responsive forum cannot compensate for lost grounds when

no reconstruction or alternative can answer the question. Existing routes may already provide the needed orientation or support, making another notice or intermediary unnecessary. What matters is an adequate route for each material objection within scope. The functions that follow organize the assessments those routes need to support.

6. Six informational functions

The register groups the operations involved in investigating a determination. Its six functions distinguish questions that may arise in that inquiry; one record or assessment can serve several functions. Table 1 moves from orientation to assessment without prescribing the order of investigation. Across the functions, baseline information makes relevant concerns recognizable and further access enables investigation. A faithful substitute performs the relevant operation and returns findings the affected party can use; historical inquiry must reach the determination as made. Section 4's necessity test applies to each operation.

Table 1. The six informational functions.

Function	Core question	Investigative failure	Route to inquiry or correction
1. Determination and consequence	What did the institution decide, and what followed?	The person cannot tell whether a score led to denial, delay, ranking, or advice.	Identify the determination, its consequence, and the output's institutional role.
2. Sociotechnical roles and evidence control	Who produced, interpreted, adopted, changed, or controls the grounds?	The source of a disputed input or the actor controlling it is unknown.	Reconstruct relevant roles, interventions, and control of evidence.
3. Data, source, and quality	Is the evidence accurate, current, complete, correctly linked, and appropriately constituted?	Suspected record defects or distorted representation cannot be investigated.	Examine provenance, records, transformations, and the wider evidence base.
4. Governing-standard–operation fit	Does the implemented target, proxy, rule, or cutoff correspond to the governing standard?	The implemented criterion cannot be compared with the authorized one.	Compare the operation with the standard and identify mismatches.
5. Evidential warrant, uncertainty, and threshold	Does the evidence support the inference, its weight, and the consequence attached to the threshold?	Validation, uncertainty, or support for the operative cutoff cannot be assessed.	Examine validation, uncertainty, subgroup performance, thresholds, and reliance.

Function	Core question	Investigative failure	Route to inquiry or correction
6. Contextual applicability	Do the system, evidence, rule, and assumptions apply in these circumstances?	The conditions supporting the use cannot be compared with the case.	Assess contextual fit and identify grounds to gather evidence, limit use, or decline application.

Function 1 identifies what the institution concluded, which output it adopted or relied on, and the practical consequence. A person may be able to inspect input data and validation yet remain unable to tell whether a score led to denial, delay, or a recommendation. The missing knowledge concerns what the institution did with the output. The baseline must make that relation intelligible. Where the consequence identifies the determination and the output–action relation is fixed and public, this work is already done. An additional notice would supply no missing capability.

Function 2 reconstructs who generated, supplied, interpreted, adopted, changed, or controls the grounds, including automation's contribution and human interventions. A person may inspect a disputed record without knowing who supplied or altered it, or who can retrieve the underlying evidence. A responsibility map does not establish data accuracy, and identifying a reviewer does not establish the competence, independence, or authority required for institutional uptake. Inquiry concerns what happened in the case, beyond a generic account of how the process should work. The baseline identifies the adopting institution and automation's material role; further access reconstructs the relevant chain and locates control of its evidence. An independent examiner can do this work from contemporaneous records. Where one institution enables the affected party to perform the same task adequately without reconstructing that chain, reconstruction is dispensable.

Function 3 assesses the constitution and integrity of the evidence base: data used, omitted, inferred, linked, transformed, or treated as authoritative in the case and the development population. It examines provenance, identity, accuracy, currency, missingness, labels, proxies as data features, linkage, representation, and whether information collected for one purpose appropriately entered this determination (Nissenbaum 2010). Barocas and Selbst distinguish errors in individual records, systematically poorer records for some groups, and distorted group representation even where individual records are sound (2016, 684–85). Accuracy-maximizing prediction can also reproduce inequality through features that redundantly encode protected status (691–92). Aggregate performance cannot settle the accuracy of a case record or the representativeness of development data. Baseline cues identify relevant values, sources, or features of the wider evidence base. Inquiry may then need to reach training labels, exclusions, transformations, and aggregate evidence beyond the personal record. Findings should explain corrections or why the data remain unchanged. Whether sound data operationalize the governing standard is Function 4's question.

Function 4 compares the governing legal, contractual, professional, or policy standard with the implemented target, construct, proxy, objective, rule, or cutoff. It distinguishes computational fidelity from faithful operationalization: justification must refer to a norm external to the system (Henin and Le Métayer 2022, 1399). Passi and Barocas trace how a project's goal was reformulated through proxies until the optimized quantity reflected the choices made in defining and measuring it (2019, §5). Jacobs and Wallach distinguish an unobservable construct from its measurement model; predictive validity is one aspect of construct validity, distinct from out-of-sample predictive performance (2021, §§2–3.2). Accurate data and strong predictive performance against the implemented target can coexist with a mismatch to the governing standard. Criteria displacement occurs when review adopts the algorithm's criterion in place of the governing standard. Fit also leaves legitimacy open: a faithfully implemented standard can itself be challenged. Baseline cues identify the governing standard and implemented criterion; further access enables comparison with the relevant variables, labels, rules, weights, and thresholds actually used. For a threshold, Function 4 asks whether the cutoff faithfully implements the standard; Function 5 asks whether the evidence warrants the consequence attached to crossing it.

Function 5 assesses evidential support for an inference and the consequence attached to a selected threshold. It examines validation design, calibration, uncertainty, subgroup performance, threshold proximity, competing evidence, and the support connecting inputs to conclusions. Evidential warrant is distinct from the causal path producing an output. A model may rank cases well while giving unreliable risk estimates: discrimination and calibration differ, and calibration should be checked in external validation and the model updated where indicated (Van Calster et al. 2019). Usefulness depends on threshold probability and the consequences of error, not accuracy alone (Vickers and Elkin 2006). Deployment guidance adds prospective validation in the use setting and monitoring as populations and protocols change (Wiens et al. 2019); reporting guidance identifies the version, setting, human factors, and modifications evaluated (Vasey et al. 2022). Accurate inputs and faithful implementation do not establish warrant for the population and case. Baseline cues flag material uncertainty, threshold proximity, or known limits. Further access reaches the validation and case-specific evidence needed to assess reliability, alternatives, uncertainty, and the score–consequence relation. A second assessment is a faithful substitute when it evaluates the disputed inference or threshold use in the original determination and returns usable findings, including uncertainty. An assessment that independently settles entitlement may make that operation dispensable for entitlement while leaving the historical question unanswered. Population-level performance may leave case-specific doubts near a threshold or outside the evidence's strongest range unresolved; a conceded recording error may permit correction without broader validation inquiry.

Function 6 assesses whether the system, evidence, rule, and assumptions apply to the actual person, population, institution, place, period, workflow, and circumstances. Validity established in development may not carry over after population or protocol shifts, policy changes, altered workflows, or repurposing. Selbst et al. describe the portability trap: a solution built for one

social context can mislead or cause harm in another, including across jurisdictions within a domain (2019, 61). Applicability requires attention to the circumstances of use even where the data are accurate and the governing standard faithfully implemented. Baseline cues identify the intended domain, deployment conditions, and known limits. Further inquiry compares development and deployment conditions and introduces context the system did not capture. A competent, sufficiently independent professional may perform that assessment. A fresh assessment may settle what should happen next; historical inquiry must also address the applicability of the original use, without reconstructing every step. Domain knowledge identifies which contextual differences matter. Where validation already addresses the present context, one inquiry can serve both warrant and applicability.

The functions interact. Data exclusions can weaken validation, target choices can change which evidence is relevant, and an inference supported in one setting may fail in another. Functions 3–6 distinguish the integrity of the evidence base, the fit between the governing standard and implemented operation, the warrant for an inference and its consequence, and applicability in the circumstances of use. Coding can record overlap while identifying each task and corrective route. One proposed rival partition merges Functions 3, 4, and 5 into an evidence-adequacy category while retaining their operations as internal distinctions. The test is whether the sixfold organization improves recognition of material concerns, information requests, or identification of corrective routes compared with that equally informative grouping.

An assessment should pair a six-function profile of required tasks, available routes, and their adequacy with a separate assessment of institutional uptake. The profile is not averaged: abundant information cannot compensate for a blocked task or an absent route to consequential review. It should identify the obstruction and help compare remedies. Domain modules can specify the sector-specific evidence, expertise, and procedures needed to investigate each question. Portability testing should establish whether those adaptations preserve investigative capability. Section 7 examines how obstacles to inquiry and institutional review interact in a litigated case.

7. *K.W. v. Armstrong* as a structured demonstration

K.W. shows how a right to appeal can coexist with barriers to investigating the decision. Idaho's statistical budget tool calculated individual allocations for Medicaid waiver participants (789 F.3d at 967). A 2014 class-wide preliminary injunction required court-approved advance notice and an opportunity for a fair hearing before covered service reductions or terminations took effect; the Ninth Circuit affirmed its extension to the class in 2015 (Dkt. 130, at 31; 789 F.3d at 966). The district court's March 2016 opinion examined the tool's data, testing, and review (*K.W.*, 180 F. Supp. 3d 703, slip op. at 11, 36). A June 2016 partial declaratory judgment recognized rights to inspect and copy past, present, and future budget methodologies and supporting records, with protections for identifying information and lawful privilege. It left

access to the specified Scales of Independent Behavior–Revised (SIB-R) materials unaddressed (*K.W. v. Armstrong*, No. 1:12-cv-00022-BLW, Dkt. 301, at 2–3 [D. Idaho June 29, 2016]).

The notices identified eligibility and a budget limiting available services, supplying the orientation described by Function 1. Knowing the amount, however, did not tell participants what to question. The Ninth Circuit upheld the finding that the notices failed to explain the decreases sufficiently to permit a responsive defense (789 F.3d at 967, 971–74). The district court later rejected a third proposed notice whose checkbox for an observed or verified change could leave the basis of verification unexplained (slip op. at 26–29). Supporting evidence did not have to accompany the notice, but the reasons for the reduction did (slip op. at 29 n. 8). Separate delivery of evidence could support investigation only if the notice first made the grounds for inquiry recognizable.

Function 3 concerns the construction and integrity of the evidence base. Of 3,512 participant records initially assembled, 998 were loaded into the statistical software, which could use 733 to build the tool. Erroneous records, unmatched identifiers, and missing information contributed to the exclusions, and one region was underrepresented (slip op. at 11–12). These features of the development sample required investigation. An accurate explanation of the formula could not establish the quality of the data used to build it.

Function 5 asks what supports reliance on the resulting inference. The agency estimated inadequate budgets for 10 to 15 percent of participants, had not recalculated the tool annually, and had increased 62 percent of reconsidered budgets under a discontinued process (slip op. at 12–13). The court rejected average budget-use figures as an answer to individual inadequacy and found that the tool's unreliability violated due process (slip op. at 13, 24). Data integrity and evidential warrant pose different questions: what entered the analysis, and what supports reliance on its results? Individual inadequacy also raises Function 6's question: do the tool's assumptions and supporting evidence fit the participant's circumstances? Regional underrepresentation can matter to both data quality and contextual applicability.

At informal review, managers could increase a budget only if they judged the increase necessary for the participant's health and safety. Neither the statute nor agency regulations defined that standard, and participants whose requests were denied were not told why. The court required written standards, evidentiary guidance, and reasons for denial (slip op. at 14–15, 18–19). Function 4 here concerns the relationship between the review standard and its application. Participants needed to know how reviewers applied the standard and what evidence they required, alongside how the budget itself was calculated.

Function 2 identifies the roles of the department, assessment contractor, assessors, and software in producing the determination (slip op. at 5–7). Access to the grounds raised a further problem. Participants and their representatives could inspect specific SIB-R questions and responses to understand scores but could not copy them, and the department barred use of SIB-R materials at hearings (slip op. at 22–23). The court required a plan permitting inspection of necessary

portions and presentation of challenged analysis (*K.W.*, 180 F. Supp. 3d at 717). This is the protected route discussed in Section 5: access had to support both examination of the assessment and use of the findings in a challenge.

Participants also needed assistance and a forum able to act on their objections. The court required a plan securing a suitable representative's commitment before informal review and confirmation of a reduction. Its four ordered plans addressed assessment access, representation, review standards, and tool improvement with regular testing (slip op. at 19–21, 24–25, 36). Hearing officers could uphold the department's decision or remand for updated assessment documents, budget recalculation, or reconsideration of denied services. They could not themselves modify budgets or approve denied services (789 F.3d at 967; 180 F. Supp. 3d 703, slip op. at 8). Review authority and corrective authority were divided. The route to correction therefore required attention to what the department did after remand.

K.W. exposes the relationship between defects in a determination and barriers to investigating them. Flawed development data and unreliable inferences gave participants grounds to question their budgets. Missing reasons, restrictions on assessment materials, undefined review standards, and inadequate assistance obstructed the work needed to assess and pursue those grounds. The informational functions, delivery arrangements, and institutional response each mattered. Identifying the decision and providing an appeal did not supply the capability to investigate it.

The NIC framework makes that gap diagnosable. It connects the objection to the investigative task, identifies the operations and access needed to perform it, and asks whether the findings can enter a process capable of acting on them. That connection directs repair toward the blocked capability. The next sections examine which operations are indispensable, which distinctions improve investigation, and how the framework carries across institutional settings.

8. Objections and boundary conditions

The sixfold division may place boundaries at the wrong points in a continuous informational field. Those boundaries earn their place if coders can apply them consistently and they improve investigation compared with an equally informative grouping that preserves the same operations. The proposed merger of Functions 3, 4, and 5 provides a direct comparison.

Necessity becomes circular if the task or adequacy floor merely restates the operation being tested. The framework therefore starts with the material objection and the assessment it requires, specifying the task and standard of adequacy independently of the candidate operation. Three failures remain distinct. An adequate route that omits an operation defeats its necessity. A distinction that adds no stable investigative value calls for revising the boundary. A recurring material objection the register cannot represent exposes a coverage gap. Each failure directs a different repair.

Access also creates burdens. Information can overwhelm recipients, expose protected material, or invite manipulation. Staging can reduce volume while preserving discovery: baseline

information must let people recognize what warrants investigation, while fuller detail follows the question. Protected mediation requires a competent, independent intermediary who examines the disputed grounds and returns findings the affected party can use in time. Cognitive, administrative, privacy, security, and strategic costs must be compared across routes that preserve adequate inquiry.

A process that works for a represented professional may fail someone facing disability, language barriers, poverty, time pressure, or retaliation. Usable access can require translation, assistance, representation, interim protection, or collective action. Adequacy must be judged for the people the route serves and the barriers they face.

The affected-party standpoint includes representative and collective agency. Individual disputes may leave population-level discrimination, feedback effects, and institutional objectives unexamined. Unions, advocates, auditors, and regulators can investigate systemic concerns through their mandates and access to evidence. Viljoen's challenge goes further: individual claims and remedies cannot adequately represent the population-level interests generated by data production (2021, 578–79). Contestability belongs alongside audit, prohibition, redistribution, and democratic governance. A checklist of disclosures that ignores who can use information and how institutions respond would reproduce the abstraction from social context criticized by Selbst et al. (2019, 60–62). The framework maps what can be challenged and how. A contestable determination may still be unauthorized, unlawful, or unjust.

Epistemic dependence poses the deepest challenge. Hardwig argues that justified trust in an expert can support belief without possession of the expert's evidence. Expertise may also be needed to judge whether a layperson's objection is sound (1985, 335–37, 342–43). An affected person can therefore reasonably accept an expert's answer without reproducing its reasoning. Lange's account of AI deference addresses a different relationship. A system with an established epistemic advantage should carry strong default weight as a contributory reason while preserving the user's independent reasons. Domain mismatch, evidence undermining reliability, disagreement from a comparably reliable authority, or relevant reasons the system may have missed can defeat that deference (2026, 180–81).

Meaningful contestability can operate through justified epistemic dependence. Representatives also rely on experts, and justified assurance can settle some questions. The test is whether the distributed route adequately answers the material objection. A competent, independent expert can supply that route by examining the disputed grounds and returning findings the affected party can use to assess the objection and pursue consequential review. Repeating the institution's assurance does not answer an objection to its basis. The party must be able to initiate inquiry and use its findings; final authority over every technical judgment is unnecessary. Whether AI advice remains contributory or displaces independent evidence is an empirical question for Section 9.

The obligation to provide contestation rests on law, policy, contract, professional norms, institutional commitments, or an independent normative argument. I develop the argument from

private power in the companion article. Section 9 separates two questions about transfer: whether adapted routes preserve investigative capability as standards, evidence, expertise, and procedures change, and whether the register's distinctions remain useful in the new setting. A failed transfer requires diagnosing whether the problem lies in the register, delivery arrangements, or institutional barriers. That diagnosis determines what must change and where the framework's use is supported.

9. Making the register empirically vulnerable

The research program tests three claims: that the register covers material objections, that its distinctions improve investigation, and that particular operations are indispensable for specified tasks under specified conditions. A varied corpus of successful and failed contestation episodes would link determinations, requests, grounds, and consequences. Initial coding would identify objections and epistemic tasks independently of the register and rival classifications; subsequent mapping would identify what each classification misses. Preregistered comparisons, including the proposed merger of Functions 3, 4, and 5, would use the same cases and equally informative representations preserving the same operations. Development would follow explicit stopping conditions (Nickerson, Varshney, and Muntermann 2013, 343–46), with coding agreement contributing to a broader assessment of interpretive quality (O'Connor and Joffe 2020, 11).

Development would begin with purposively selected credit cases, followed by public benefits and health care. The program would build a contestability toolkit: connected affected-party and reviewer interfaces, a case generator, a record linking grounds, evidence, and disposition, tested modules, scoring tools, and an adaptation guide. Independent users in education and public administration would then adapt the toolkit to synthetic scholarship and permit decisions. Both settings would be held out from development, including scenario authoring and pretesting. Transfer tests would use a frozen version with permitted adaptations specified in advance. They would assess whether investigative capability survives changes in standards, evidence, expertise, and procedures, and whether the local standards remain coherent instances of Function 4. Comparisons with rival classifications would also test whether the register's distinctions remain useful in the held-out settings.

Affected parties and institutional participants would separately assess coverage, burdens, and alternatives. Cognitive interviews would examine whether people understand baseline cues and can formulate useful questions before a request menu supplies them. Confirmatory tests would preregister defects and epistemic tasks specified by independent experts, feasible routes, adequacy floors, precision requirements, and revision rules. Tasks and floors must be justified independently of the operation being tested. The functions would not be treated as interchangeable indicators of one latent trait (Diamantopoulos and Winklhofer 2001, 270–71).

Experiments would compare a complete route, omission of the focal operation and its equivalents, a faithful substitute, and a conventional-reason control matched for length, salience, source, and affordances. Within the specified comparison, conditional necessity is supported

when at least one route performing the operation clears the adequacy floor, and every tested omission route falls below it, with prespecified precision. An adequate omission route defeats necessity regardless of how the complete route performs. Substitute tests address delivery: adequacy asks whether the substitute clears the floor; noninferiority asks whether it performs no more than a prespecified margin worse than the complete route. A substitute can be noninferior while both routes remain inadequate. A substitute that falls below the adequacy floor fails as a delivery route. A nonsignificant difference establishes neither equivalence nor noninferiority.

Comparison with a neighboring task would help distinguish loss tied to the omitted operation from a general reduction in useful support. The necessity judgment would still turn on performance against the adequacy floor. Outcomes would include defect diagnosis, useful requests, and use of evidence. They would distinguish unwarranted allegations from reasonable requests to check a decision in which no defect has been designated. Inquiry need not begin with a correct accusation.

A separate delivery comparison would test whether staged access supports the same task adequately at lower burden than a complete dossier, including whether people can discover what to investigate before receiving a request menu. Field studies using institutional records would trace delays, withheld grounds, evidence entering review, and correction or warranted reaffirmation. Institutional uptake would be measured separately from access. Reviewer simulations would distinguish judgments about evidential warrant, the disposition justified on the merits, and the action available under assigned powers. That separation would identify when an appropriate response requires escalation or referral.

A linked deference study would vary independent evidence and labeled advice while controlling reliability and evidence overlap. Preemption would be assessed through reduced sensitivity to independent evidence, retained adviser influence, and movement toward adviser-only judgment. A prespecified benchmark would integrate evidence and advice according to their reliability and overlap, distinguishing justified agreement from displaced judgment. I predict more discriminating reliance on automated assistance as people become better able to assess grounds for challenge. Accessible participation and group-specific reporting would identify unequal performance. Materials and negative findings would be shared where privacy permits.

Revision rules would merge distinctions that add no stable investigative value over an equally informative grouping preserving the same operations. They would split categories that obscure recurring differences in operations and remedies, and add categories or restrict scope where material objections remain unrepresented. Informative tests require faithful manipulation, independent adjudication, credible alternatives, and sufficient precision to assess the relevant floor or comparison margin.

An adequate omission route established by such a test defeats the necessity claim for those conditions. Changing the task or threshold afterward cannot preserve the original claim. If a well-conducted test leaves the omission route's adequacy unresolved, I propose one diagnostic

replication. If the same claim remains inconclusive across the prespecified settings after that replication, I would downgrade it. Demonstrated dispensability rejects the claim; persistent uncertainty reduces its standing. The program would thus connect each retained claim to the evidence supporting it and direct revision toward the operations, distinctions, or arrangements that fail.

10. Conclusion

The NIC framework's central innovation is to organize informational necessity around epistemic tasks: the work needed to assess a material objection. It connects those tasks to informational operations and tests whether adequate investigation remains possible without them. This move exposes why an accurate explanation, abundant disclosure, and a review procedure can still leave a person unable to investigate the grounds of a decision. It also identifies what would restore that capability.

The approach carries a common question across domains while allowing the answer to change with the setting. Public benefits, health care, credit, and employment involve different standards, evidence, expertise, and procedures. In each, the framework asks what an affected person must be able to establish, what inquiry that requires, and how its findings can enter consequential review. The model's economy lies in this common structure. Institutions can adapt the means of investigation while retaining a shared method for assessing their adequacy.

For affected people, the framework identifies a route from reasonable doubt to informed challenge, directly or through assistance. For institutions, it offers a method for locating barriers, directing resources toward necessary investigative work, and evaluating alternatives that protect legitimate interests while preserving inquiry. It connects informational access to the authority needed for correction or warranted reaffirmation. For researchers, epistemic tasks provide a common basis for comparing failures and remedies across settings, testing necessity claims, and improving the framework itself.

Fairness and justice give this work its purpose. People need to discover whether the grounds of their treatment withstand scrutiny. Institutions need to recognize errors, respond to justified objections, and sustain decisions that warrant reliance. Meaningful contestability connects those interests through inquiry that people can exercise and institutions can act upon. The NIC framework gives that shared commitment a practical, testable form: a way to identify what investigation requires and make its findings matter.

Works cited

Scholarly works

Alfrink, Kars, Ianus Keller, Gerd Kortuem, and Neelke Doorn. 2023. "Contestable AI by Design: Towards a Framework." *Minds and Machines* 33: 613–39. https://doi.org/10.1007/s11023-022-09611-z.

Barocas, Solon, and Andrew D. Selbst. 2016. "Big Data's Disparate Impact." *California Law Review* 104: 671–732. https://doi.org/10.15779/Z38BG31.

Bayamlıoğlu, Emre. 2018. "Contesting automated decisions: A view of transparency implications." *European Data Protection Law Review* 4 (4): 433–46. https://doi.org/10.21552/edpl/2018/4/6.

Bayamlıoğlu, Emre. 2022. "The right to contest automated decisions under the General Data Protection Regulation: Beyond the so-called 'right to explanation.'" *Regulation & Governance* 16 (4): 1058–78. https://doi.org/10.1111/rego.12391.

Bovens, Mark. 2007. "Analysing and Assessing Accountability: A Conceptual Framework." *European Law Journal* 13 (4): 447–68. https://doi.org/10.1111/j.1468-0386.2007.00378.x.

Cen, Sarah H., Hannah Ismael, and Lucia Zheng. 2026. "Barriers to Evidence in AI-Related Cases and the Privatization of Proof." In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, 7793–834. https://doi.org/10.1145/3805689.3812219.

Cobbe, Jennifer, Michelle Seng Ah Lee, and Jatinder Singh. 2021. "Reviewable Automated Decision-Making: A Framework for Accountable Algorithmic Systems." In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 598–609. https://doi.org/10.1145/3442188.3445921.

Diamantopoulos, Adamantios, and Heidi M. Winklhofer. 2001. "Index Construction with Formative Indicators: An Alternative to Scale Development." *Journal of Marketing Research* 38 (2): 269–77. https://doi.org/10.1509/jmkr.38.2.269.18845.

Fraser, Henry, Aaron J. Snoswell, and Rhyle Simcock. 2022. "AI Opacity and Explainability in Tort Litigation." In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 185–96. https://doi.org/10.1145/3531146.3533084.

Fricker, Miranda. 2013. "Epistemic justice as a condition of political freedom?" *Synthese* 190 (7): 1317–32. https://doi.org/10.1007/s11229-012-0227-3.

Hardwig, John. 1985. "Epistemic Dependence." *The Journal of Philosophy* 82 (7): 335–49. https://doi.org/10.2307/2026523.

Henin, Clément, and Daniel Le Métayer. 2022. "Beyond explainability: justifiability and contestability of algorithmic decision systems." *AI & Society* 37: 1397–410. https://doi.org/10.1007/s00146-021-01251-8.

Jacobs, Abigail Z., and Hanna Wallach. 2021. "Measurement and Fairness." In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 375–85. https://doi.org/10.1145/3442188.3445901.

Kaminski, Margot E., and Jennifer M. Urban. 2021. "The Right to Contest AI." *Columbia Law Review* 121 (7): 1957–2048. https://columbialawreview.org/content/the-right-to-contest-ai/.

Lange, Benjamin. 2026. "Epistemic Deference to AI." In *Bridging the Gap Between AI and Reality*, edited by Bernhard Steffen, 174–86. Lecture Notes in Computer Science 16032. Springer. https://doi.org/10.1007/978-3-032-01377-4_9.

Lyons, Henrietta, Eduardo Velloso, and Tim Miller. 2021. "Conceptualising Contestability: Perspectives on Contesting Algorithmic Decisions." *Proceedings of the ACM on Human-Computer Interaction* 5 (CSCW1), article 106: 1–25. https://doi.org/10.1145/3449180.

Nickerson, Robert C., Upkar Varshney, and Jan Muntermann. 2013. "A Method for Taxonomy Development and Its Application in Information Systems." *European Journal of Information Systems* 22 (3): 336–59. https://doi.org/10.1057/ejis.2012.26.

Nissenbaum, Helen. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford University Press.

Nussbaum, Martha. 2003. "Capabilities as Fundamental Entitlements: Sen and Social Justice." *Feminist Economics* 9 (2–3): 33–59. https://doi.org/10.1080/1354570022000077926.

O'Connor, Cliodhna, and Helene Joffe. 2020. "Intercoder Reliability in Qualitative Research: Debates and Practical Guidelines." *International Journal of Qualitative Methods* 19: 1–13, article 1609406919899220. https://doi.org/10.1177/1609406919899220.

Passi, Samir, and Solon Barocas. 2019. "Problem Formulation and Fairness." In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 39–48. https://doi.org/10.1145/3287560.3287567.

Pi, Yulu, Lucas Lichner, Jae Woo Lee, Sijia Xiao, Renwen Zhang, and Jatinder Singh. 2026. "Push and Pushback in Contesting AI: Demands for and Resistance to Accountability." In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, 3590–625. https://doi.org/10.1145/3805689.3812323.

Ploug, Thomas, and Søren Holm. 2020. "The four dimensions of contestable AI diagnostics—A patient-centric approach to explainable AI." *Artificial Intelligence in Medicine* 107, article 101901. https://doi.org/10.1016/j.artmed.2020.101901.

Rawls, John. 1999. *A Theory of Justice*. Rev. ed. Cambridge, MA: Belknap Press of Harvard University Press.

Selbst, Andrew D., danah boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. "Fairness and Abstraction in Sociotechnical Systems." In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59–68. https://doi.org/10.1145/3287560.3287598.

Springer, Aaron, and Steve Whittaker. 2020. "Progressive Disclosure: When, Why, and How Do Users Want Algorithmic Transparency Information?" *ACM Transactions on Interactive Intelligent Systems* 10 (4), article 29: 1–32. https://doi.org/10.1145/3374218.

Stewart, Matthew. 2026. "Beyond Explanation: Evidentiary Rights for Algorithmic Accountability." In *Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency*, 1066–97. https://doi.org/10.1145/3805689.3806532.

Van Calster, Ben, David J. McLernon, Maarten van Smeden, Laure Wynants, and Ewout W. Steyerberg, on behalf of Topic Group "Evaluating diagnostic tests and prediction models" of the STRATOS initiative. 2019. "Calibration: the Achilles Heel of Predictive Analytics." *BMC Medicine* 17, article 230. https://doi.org/10.1186/s12916-019-1466-7.

Vasey, Baptiste, Myura Nagendran, Bruce Campbell, David A. Clifton, Gary S. Collins, Spiros Denaxas, Alastair K. Denniston, Livia Faes, Bart Geerts, Mudathir Ibrahim, Xiaoxuan Liu, Bilal A. Mateen, Piyush Mathur, Melissa D. McCradden, Lauren Morgan, Johan Ordish, Campbell Rogers, Suchi Saria, Daniel S. W. Ting, Peter Watkinson, Wim Weber, Peter Wheatstone, Peter McCulloch, and the DECIDE-AI expert group. 2022. "Reporting Guideline for the Early-Stage Clinical Evaluation of Decision Support Systems Driven by Artificial Intelligence: DECIDE-AI." *Nature Medicine* 28 (5): 924–33. https://doi.org/10.1038/s41591-022-01772-9.

Vickers, Andrew J., and Elena B. Elkin. 2006. "Decision Curve Analysis: A Novel Method for Evaluating Prediction Models." *Medical Decision Making* 26 (6): 565–74. https://doi.org/10.1177/0272989X06295361.

Viljoen, Salomé. 2021. "A Relational Theory of Data Governance." *Yale Law Journal* 131 (2): 573–654. https://yalelawjournal.org/article/a-relational-theory-of-data-governance.

Wiens, Jenna, Suchi Saria, Mark Sendak, Marzyeh Ghassemi, Vincent X. Liu, Finale Doshi-Velez, Kenneth Jung, Katherine Heller, David Kale, Mohammed Saeed, Pilar N. Ossorio, Sonoo Thadaney-Israni, and Anna Goldenberg. 2019. "Do No Harm: A Roadmap for Responsible Machine Learning for Health Care." *Nature Medicine* 25 (9): 1337–40. https://doi.org/10.1038/s41591-019-0548-6.

Cases and orders

K.W. ex rel. D.W. v. Armstrong, 789 F.3d 962 (9th Cir. 2015). https://cdn.ca9.uscourts.gov/datastore/opinions/2015/06/05/14-35296.pdf.

K.W. ex rel. D.W. v. Armstrong, 180 F. Supp. 3d 703 (D. Idaho 2016). https://www.govinfo.gov/content/pkg/USCOURTS-idd-1_12-cv-00022/pdf/USCOURTS-idd-1_12-cv-00022-13.pdf.

K.W. v. Armstrong, No. 1:12-cv-00022-BLW, Dkt. 130 (D. Idaho filed Mar. 25, 2014) (class certification and class-wide preliminary injunction).

K.W. v. Armstrong, No. 1:12-cv-00022-BLW, Dkt. 301 (D. Idaho June 29, 2016) (partial declaratory judgment).

CK v. Magistrat der Stadt Wien (*Dun & Bradstreet Austria*), Case C-203/22, ECLI:EU:C:2025:117 (CJEU Feb. 27, 2025). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62022CJ0203.

AI assistance: During the preparation of this article, I used OpenAI Codex for Toulmin model argument mapping and reference analysis, and Claude for sentence-level editing to improve clarity and brevity. All prompts and outputs are retained in my working notes to track intellectual provenance, and I am solely responsible for the article's content and claims.