

GOVERNED AUTONOMY

The Reckoning: Calibration and Accountability for Autonomous Judgment

A paper in the Governed Autonomy collection

Andrew Durden
Independent research
Publication edition 2026.10.7

Research on authority, evidence, and continuity in autonomous systems.

Content identifier: 30f2410373ee5d9211c2ab5e

Contents

The Reckoning: Calibration and Accountability for Autonomous Judgment	3
Abstract	3
1. Why occurrence evidence is insufficient	3
2. Calibration, discrimination, and outcomes	4
3. What makes a forecast assessable	4
4. Agency is different from manufactured evidence	5
5. Samples, reference populations, and selective reporting	6
6. Changes in the production of judgment	6
7. Worked conceptual scenarios	8
8. Consequences, public proof, and institutional boundaries	9
9. Evidence limits and research questions	10
Source guide	11

The Reckoning: Calibration and Accountability for Autonomous Judgment

Governed Autonomy · Publication 2026.10.7

Research study: conceptual exposition of The Reckoning specification

Source: Public exposition, successor 2026.10.7; historical research and evidence separately qualified

Abstract

An action record can establish what a system did without establishing whether it assessed the uncertainty well. Repeated success may reflect easy tasks, favorable conditions, or excessive confidence that has not yet encountered a difficult case. The Reckoning proposes a separate account of probabilistic judgment: expectations recorded before outcomes, outcomes resolved against defined evidence, and evaluations that retain their sample and context limits. Calibration describes the relationship between stated probabilities and observed frequencies; it is distinct from task success, permission to act, and the economic value of an outcome.

This paper examines the conditions under which institutional calibration could inform decisions: independently assessable outcomes, explicit forecast populations, appropriate reference forecasts, and attribution to the environment that produced the judgment. Four scenarios distinguish calibration from competence, authority, and causal improvement. The retained 10 October 2026 author-directed, model-assisted source examination identified resolution and calibration-related functions in the examined Freehold engine revision. Their implementation does not establish measured calibration performance or the proposed aggregate standing and its consequences. Numerical examples in the source are illustrative rather than measured cohorts, so the effectiveness of the proposed policy remains an empirical question.

1. Why occurrence evidence is insufficient

A system can provide a complete account of an action and still have misunderstood its prospects. It may correctly record a purchase that was authorized, accurately report the subsequent result, and preserve a valid history. None of those facts establishes that the system's initial confidence was justified. Evidence about execution and evidence about judgment answer different questions.

The distinction becomes visible across repeated decisions. Consider two forecasters that succeed equally often. One describes the uncertainty in a way that broadly matches the observed outcomes; the other repeatedly expresses much greater confidence. Their success rates alone do not reveal the difference. Yet that difference matters when people depend on their estimates to allocate resources, compare options, or decide how much contingency to retain.

The Reckoning proposes that expectations should be attributable and fixed before the relevant outcome is known. Later evaluation can then compare a stated probability with the defined result. This records what was reported at the time; it does not establish the forecaster's unobservable private belief. Timing and provenance support accountability without turning a recorded statement into proof of sincerity.

The source places this work alongside institutional execution and public evidence. Its conceptual role is additive: it observes the quality of judgment and may support restrictions, while existing authority remains responsible for whether an action is allowed. The distinction fits the collection's governed autonomy thesis, but the proposed institutional integrations retain their own evidence requirements. [Public review scope](#).

2. Calibration, discrimination, and outcomes

Calibration concerns whether probabilities correspond to observed frequencies in a defined population of forecasts. It is assessed across enough relevant observations to support a comparison. A single unsuccessful forecast does not establish miscalibration, just as a single successful forecast does not establish good calibration.

Discrimination concerns whether forecasts distinguish situations with different prospects. Someone who always gives the same probability can appear well calibrated in an aggregate population while offering little information about which cases are more likely to succeed. A useful assessment should therefore avoid treating calibration as the whole of judgment quality.

Outcome competence concerns whether the system achieves the task. **Decision value** concerns whether choosing the task was worthwhile given its benefits, costs, and alternatives. These can diverge. A low-probability experiment may be worth attempting because of its potential value or the information it produces. A highly predictable activity may be unimportant. Calibration should not be used as a substitute for evaluating either task quality or decision value.

The source discusses established probability-scoring ideas, including Brier and logarithmic losses, calibration error, and comparison with a reference forecast. Their interpretation depends on the event, reporting domain, population, and resolution process. Conditional mathematical properties of an individual scoring rule do not automatically transfer to an aggregate standing assembled from several statistics and policy choices.

A score may be reproducible while the decision to use it as a reputation or resource constraint remains a separate policy hypothesis. Assessing that policy requires evidence about the behavior it induces as well as the statistic it computes. [Public review scope](#); [proper-scoring background](#).

3. What makes a forecast assessable

An assessable forecast needs more than a confident sentence. The outcome must be sufficiently well defined that later evidence can determine what happened. The time horizon matters because “eventually” can evade resolution indefinitely. The subject and setting matter because a result from one kind of work may say little about another.

The source requires a prospective record and an outcome grounded in defined evidence rather than retrospective model judgment. This reduces opportunities to reinterpret a prediction after seeing the result. It also makes the resolution process part of the claim: an evaluation is only as sound as the connection between the defined event and the evidence used to resolve it.

Absence of an event can be meaningful when the proposition explicitly says that the event will occur within a specified horizon and observation is adequate. Absence is less informative when the observation itself is interrupted. Those cases should not be collapsed into a single outcome. A forecast that fails should not be converted into “unknown” merely to improve a record, while a forecast that genuinely cannot be tested should not acquire a fabricated result.

The source distinguishes ordinary non-occurrence from justified indeterminacy and separately records the reasons for the latter. This is a conceptual accounting requirement, not a claim that every external event can be perfectly observed. Revocation, migration, incomplete source coverage, and independent counterparty events can affect whether the original test remains valid.

The relationship between the forecast and its action also needs care. Recording an expectation does not itself authorize the proposed action. Equally, an absent forecast is not automatically evidence that the underlying action lacked permission. Coverage and authority are different dimensions whose relationship must be stated rather than assumed. Public review scope.

4. Agency is different from manufactured evidence

Autonomous systems are often participants in the events they forecast. A system may predict that a planned service will attract demand and then work to make that happen. Excluding every forecast influenced by the forecaster would remove much of the judgment that matters in an acting institution.

The Reckoning instead distinguishes influence over the outcome from unilateral control over the evidence used to declare success. A forecast about genuine customer demand may remain informative even when the system builds and markets the service. A forecast that can be satisfied by creating its own administrative record offers a different kind of evidence. A fabricated transaction or cooperative counterparty that creates a favorable appearance is different again.

This distinction is conceptually useful but difficult to operationalize. Formal separation of accounts does not necessarily establish economic independence. Two counterparties can appear distinct while sharing control or incentives. A measurement can be externally hosted while remaining vulnerable to choices made by the evaluated subject. Independence must be justified in relation to the claim, not inferred from a label.

Action-contingent forecasts also complicate ordinary scoring arguments. The subject can change effort, select interventions, or abandon a plan after making the forecast. Evaluating such forecasts requires an explicit interpretation of what was predicted and under which

intended action. A desirable scoring property for fixed external events is not, by itself, a proof of good incentives for this wider decision process.

The source permits legitimate agency while treating fabricated resolution as a threat. Evidence independence is therefore necessary, but its effectiveness against collusion and strategic event selection remains to be demonstrated. Public review scope.

5. Samples, reference populations, and selective reporting

Forecast evaluation can be misleading even when every included result is correct. A subject may report only the easiest cases, choose a favorable reference population, or combine domains whose differences disappear in an overall average. A clean-looking record can conceal the decisions for which uncertainty mattered most.

Coverage describes how much of the relevant decision population carries an assessable forecast. It helps expose selective reporting, but introduces another policy choice: the definition of the relevant population. If that definition changes after outcomes are visible, the resulting statistic can be made more favorable without improving any judgment. Stable inclusion rules and clear domain boundaries are therefore part of interpretability.

A **reference forecast** represents a baseline against which performance is compared. A sophisticated forecast may add little information when a simple account of the prevailing frequency performs just as well. Conversely, using an inappropriate baseline can make useful work appear weak. The reference must be attributable to information available at the time, with an understandable population and window.

Small samples create a separate problem. A few favorable outcomes can produce impressive raw values without supporting a stable conclusion. Uncertainty intervals, adequate sample size, diversity, and domain-specific reporting can reduce overstatement, but their validity depends on assumptions about the data. Repeated observations from closely related tasks or counterparties may supply less independent information than their count suggests.

The source distinguishes raw observations from the standing it proposes to publish or use downstream. That distinction is valuable even though the particular aggregate policy remains unqualified. An unassessable population should remain unassessed; undefined statistics must not be replaced with favorable defaults. A public badge should not compress uncertainty into a level of confidence that the underlying evidence cannot support. Public review scope.

6. Changes in the production of judgment

Judgment depends on more than a model's name. The material available to it, the way context is selected, the tools it can use, and the conditions under which it generates responses can change its behavior. A familiar model label does not establish continuity of the evaluated process.

The Reckoning proposes attributing forecasts to the wider environment that produced them. This supports a distinction between historical performance and evidence for a changed

system. A favorable record from the earlier environment should not silently become a certificate for its replacement.

Change is not itself proof of regression. A correction may improve performance, a new information source may reduce uncertainty, or a tool change may eliminate a recurring error. The appropriate conclusion is that the basis for comparison has changed and fresh evidence is needed. Conservative restrictions during evaluation can be a policy response, but their thresholds and practical consequences require validation.

Observed drift is also not automatically causal attribution. A decline may coincide with a model change while the task population becomes harder or the external environment changes. Knowing which components changed helps investigation; it does not establish that one of them caused the observed difference. Reliable diagnosis needs comparable populations and attention to confounding factors.

When relevant components cannot be identified, continuity claims become weaker. Under that limitation, the source withholds its stronger protection against carrying reputation across changes. Conservative treatment of unknown components can limit reliance on an earlier record, but it cannot make every meaningful change detectable. [Public review scope](#).

6.1 Evidence during restricted operation

Reassessment concerns two decisions with different owners. The evaluator determines what the new evidence supports about the changed system. The authority responsible for the affected work determines whether a restriction can be changed under the applicable policy. Clearing a calibration restriction does not grant a previously absent permission, revive revoked access, or establish that a separate competence requirement has been met.

The source's proposed reassessment route accumulates prospective, independently resolved forecasts attributable to the changed environment and compares them with an applicable baseline under declared criteria. Historical evidence remains associated with the former environment. Changed conditions call for a new body of evidence; they are not evidence of deterioration by themselves. The source also preserves ordinary resource and gate constraints during the evaluation period. Forecasting remains optional for autonomous participants: a restriction does not authorize prompting or pressuring them to generate forecasts, and the deterministic resolution process supplies no new participant instruction. The available route consequently depends on forecasts arising from permitted work and on adequate resolution evidence. It does not establish that these remaining opportunities always generate enough comparable evidence to clear a restriction.

A restriction changes the population from which evidence is observed. If high-consequence actions are unavailable, the remaining forecasts may concern easier, smaller, or otherwise different tasks. Their favorable scores can support a claim about that restricted population without establishing restored performance on the excluded work. If a promotion restriction prevents the outcome needed to resolve a forecast, the absence of that outcome cannot simply be scored as a failed candidate method. The evaluator must retain the condition that prevented the test and apply the declared resolution semantics.

Assessment setting	Evidence available	Consequence for reassessment
Authorized work continues within the restriction	New forecasts and independently observed outcomes from the changed system	Evidence can inform reassessment for the represented task population, with inclusion rules and changes in difficulty retained.
Restrictions remove a material class of work	Only lower-exposure or different-domain outcomes are observed	The excluded class remains unassessed. Good performance on available work cannot automatically clear a claim about unavailable work.
Observation needed for resolution becomes unavailable	A prospective forecast exists, but its defined test lacks adequate evidence	Preserve the interruption and its cause. Apply justified indeterminacy where appropriate rather than fabricating a success or a failure.
An authorized evaluator proposes a separate assessment	Independently conducted tests or eligible existing forecasts may offer comparative evidence	The assessment needs its own authority, target population, resolution evidence, and decision rule. It cannot direct participant forecasting, bypass a promotion hold, or automatically replace the forecasts required by the existing clearing rule. Such a route is an evaluation option to be validated, not an implemented service established by the source.

Where adequate evidence cannot accumulate through the available route, the scientific conclusion is that the affected claim remains unassessed. The responsible owner needs an explicit reassessment disposition: continue permissible work, arrange a separately authorized assessment, repair the evidence dependency, or determine whether the governing policy requires an authorized change. None of these options permits an institution to clear its own hold or evade the active clearing conditions. A policy change is a separate authority decision; it is not evidence that the changed system has qualified. Elapsed time and silence are not substitute evidence. Requalification through the source's proposed clearing conditions is conditional on adequate comparable forecasts; universal eventual requalification is outside the established evidence.

An evaluation of the policy should measure false restrictions, time and resources needed for reassessment, useful work lost, and cases that remain unassessed. It should compare task difficulty and opportunity before and during restriction, retain excluded and indeterminate cases, and examine whether clearing decisions generalize to the work that was withheld. These observations test the consequences of the restriction policy, rather than merely verifying that its score was recomputed. [Public review scope](#).

7. Worked conceptual scenarios

Four hypothetical scenarios distinguish the conclusions that a forecast record can support.

Illustration 1 — a useful forecast that turns out wrong. A research group considers an uncertain technical approach. It records that success is unlikely but possible and defines the result that would count as success. The attempt fails under the stated conditions. That result

belongs in the record, yet it does not by itself show poor calibration. A series of comparable forecasts is needed to evaluate whether the group has assessed such uncertainty well. Whether the experiment was worthwhile additionally depends on cost and what was learned.

Illustration 2 — high confidence without authority. A service provider has a strong record in a particular domain and predicts that a proposed external action will succeed. The action lacks the required permission. Its favorable calibration history can inform another party's assessment, but it cannot supply the missing authorization. The action remains outside scope until the actual authority owner permits it. This is a denial based on permission, not a judgment that the forecast is inaccurate.

Illustration 3 — interrupted observation and recovery. A forecast depends on an independently observed outcome, but access to that evidence is revoked before the test can be completed. The evaluator first determines whether sufficient evidence remains available through an authorized route. If it does, resolution can continue. If it does not, the record must preserve the limitation and its cause rather than invent success or failure. An externally forced interruption should not by itself become evidence of manipulation. A pattern of suspicious interruptions would require its own investigation and evidence.

Illustration 4 — improved results after a system change. A system changes how it retrieves context. Its next few forecasts perform well. The earlier record remains attributable to the earlier setting, and the new results begin a distinct body of evidence. The short favorable run does not immediately establish stable improvement. Further evaluation considers comparable tasks, coverage, uncertainty, and whether the change or restriction affected which cases were attempted. Adequate evidence may support reconsideration by the appropriate owner within the affected scope. If the available tasks do not test the restricted work, that work remains unassessed and requires an authorized reassessment decision rather than an automatic inference from the easier population.

The scenarios show that calibrated judgment is neither a success label nor a permission token. Its value comes from making uncertainty assessable while retaining the surrounding conditions.

8. Consequences, public proof, and institutional boundaries

The source proposes that calibration may tighten existing institutional limits but cannot grant capabilities, create authority, or substitute for the evidence required by other decisions. Whether the proposed restrictions improve decisions is a separate empirical question.

The distinction also applies to public evidence. A portable record can help another institution evaluate a claim, but recognition of that evidence does not transfer permission. The receiving institution still determines what it accepts, for which purpose, and under which local authority. A favorable record in one domain does not establish competence in all domains.

Public projection creates a privacy problem because forecasts can reveal strategy. Aggregates and selected evidence may be useful without exposing the original private proposition, but repeated or highly detailed disclosure can permit reconstruction. Redaction

should therefore be assessed against the combined information available to the audience, not just the wording of a single claim.

The source also separates an inspection surface from a channel that directs the institution being evaluated. Its proposed public rankings and displays are not intended to become motivational instructions. This constrains the communication channel; it does not prove that allowed information cannot influence subsequent model behavior in undesirable ways.

The relationship to Freehold is explicit in the source's authority ranking and integration sections: The Reckoning is an additive supporting layer. Its relationship to Cuttlefish receipts and revocation is also proposed explicitly. These source contracts describe intended responsibilities. They do not establish that an implementation has adopted them. Broader correspondences with the architecture synthesis should retain that distinction.

9. Evidence limits and research questions

9.1 Implemented calibration functions

The retained 10 October 2026 author-directed, model-assisted source examination identified prospective-outcome, scoring, standing, and resolution-related paths in the Freehold engine. That source fact is recorded separately in the [implementation evidence](#). It does not supply a measured forecast population, an independently evaluated standing policy, or a causal improvement in decisions. An implemented evaluator and an empirically validated institutional policy are different research subjects.

9.2 Statistical and institutional limits

The source specifies a statistical and institutional policy, but supplies no measured autonomy or calibration results. Its numerical records are illustrative. An earlier populated aggregate example was withdrawn because it could not support a consistent reconstructed cohort. The proposal must therefore be evaluated on actual prospective forecasts and independently resolved outcomes before its estimates or consequences can be validated.

Several limits are decisive. The aggregate standing is a policy heuristic whose components do not inherit the incentive properties of an individual probability loss. The debiasing and interval procedures need a fully specified estimator and appropriate treatment of dependence. A record that prevents retrospective editing does not reveal private belief. A calibrated report does not prove that its underlying action was prudent or ethical, permitted or valuable.

The source's proposed protection for exploration is also a hypothesis. A scoring arrangement might reward informative uncertainty, but actors can respond by changing which tasks they attempt or report. Determining whether the policy encourages useful exploration requires studying those choices, not only recalculating scores on a fixed sample. Likewise, using calibration to weight generated training material is a proposed research direction; it is not an established guarantee against degradation from recursive synthetic data.

A substantive evaluation would examine forecast selection, difficult and easy domains, changing populations, independent resolution, repeated counterparties, interrupted tests, and strategic behavior. It would report task outcomes and decision value alongside

calibration, retain unassessable cases, and distinguish estimator uncertainty from operational uncertainty. The collection's evidence guide and glossary provide the surrounding terminology.

The central research question is whether prospective, attributable uncertainty records improve institutional decisions while preserving authority and avoiding misleading reputational compression. The proposal makes that question assessable by separating forecast quality, outcome evidence, and institutional consequences. Their joint effectiveness remains to be tested.

Source guide

This successor guide points to public arguments and their limits. The new private evidence linkage remains excluded. Previously public conceptual source guides, including any retained titles, document ordinals, sections and chapter maps, are not claimed to have been removed; see the correction scope. The historical edition remains available unchanged.

Purpose and scope. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

Scoring and interpretation. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

Prospective records and outcomes. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

Agency and evidence independence. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

Populations and publication. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

Attribution and relationships. Review the relevant argument and assumptions in this paper, the evidence guide, and the research review route. Specification, reported results and production observations have distinct scope.

No source-guide entry establishes a full implementation, a complete historical migration proof, or an independent audit.