SCIENTIFIC CULTURE, Vol. 12, No. 5, (2026), pp. 991-1006
Open Access. Online & Print

www.sci-cult.com
DOI: 10.5281/zenodo.20703580

AUDITABLE SCIENTIFIC GOVERNANCE BEYOND KPIS: AN INFORMATION-CONTROL FRAMEWORK FOR EVENT-INSTRUMENTED SOCIO-ECONOMIC SYSTEMS

Jorge Omar Portilla Jaimes[1*], Oscar J. Suarez[2], José Aguilar Castro[3,4], Oliver López-Corona[5]

[1]*Research Group CICOM, Doctoral Program in Automation, Faculty of Engineering and Architecture, Universidad de Pamplona, Pamplona, Colombia.*
[2]*Mechatronics Engineering Department, Faculty of Engineering and Architecture, University of Pamplona, Pamplona, Colombia.*
[3]*IMDEA Networks Institute, Madrid, Spain.*
[4]*CEMISID, Universidad de Los Andes (ULA), Mérida, Venezuela.*
[5]*Investigadores por México (IxM-Secihti), Instituto de Investigaciones en Matemáticas Aplicadas y Sistemas (IIMAS), Universidad Nacional Autónoma de México (UNAM), Mexico City, Mexico.*

ABSTRACT
Digital traces, process mining and artificial-intelligence-enabled analytics make many socioeconomic processes more visible, but visibility alone does not make them governable. This study develops a generic information-control framework for event-instrumented socio-economic systems: bounded local regimes where traces, admissible interventions, decision cadence, institutional authority and reviewable consequences coexist. Using a structured critical review across control, complex systems, process mining, causal inference, statistical verification, responsible artificial intelligence and public-sector governance, the paper reframes intervention as auditable knowledge generation rather than unrestricted automation. The result is a practical evidence structure for linking local claims, authorized actions, comparators, uncertainty, constraints, provenance and governance decisions. The paper contributes four elements: a local-regime definition of event-instrumented socio-economic systems; an epistemic interpretation of intervention as a falsifiable query; the auditable local knowledge bundle as a minimum evidence unit; and a weakest-link maturity rule requiring governance use to remain bounded by observability, falsifiability, statistical verification, legitimacy and auditability. The framework clarifies how accountable decisions can be supported through statistically checkable and reviewable local evidence.

KEYWORDS: *complex adaptive regimes, bounded steerability, process mining, causal evaluation, statistical verification, responsible artificial intelligence, innovation policy, evidence bundles*

INTRODUCTION

The objective of this article is to develop a concise and implementation-independent framework for scientific governance in socio-economic processes instrumented by digital events. Hospitals, tax and inspection offices, procurement platforms, social-service agencies, innovation-support programs and regulatory back offices increasingly produce time-stamped records. These traces enable prediction, prioritization, allocation and feedback. They also create a governance hazard: what can be counted, ranked or visualized may be treated as if it were already sufficiently known.

The article addresses that hazard from a control-oriented perspective, but without reducing socio-economic governance to a conventional engineering problem. Public and institutional decisions should not rely on three weak substitutes for evidence: isolated key performance indicators (KPIs), expert judgement without explicit contrast, or black-box AI prediction detached from causal, institutional and audit conditions. KPIs summarize observations, but they are not claims. Expert judgement is indispensable, but it is not a validity protocol. Artificial intelligence (AI) can estimate risks and recommend actions, but prediction alone does not establish that an intervention is justified, lawful, effective, fair or transferable (Thomas and Uminsky, 2022; Manheim, 2023; Janssen et al., 2020; Papagiannidis et al., 2025).

The proposed alternative is auditable scientific governance. A decision is scientifically defensible when it is tied to a local claim, an observation regime, an admissible intervention, a comparator, statistical uncertainty, constraint checks, provenance and accountable review. This does not mean that socio-economic systems become classical plants. They remain complex adaptive regimes: actors learn, categories drift, incentives respond to measurement, institutions change and objectives are contested (Anderson, 1972; Battiston et al., 2021; Preiser et al., 2018; Garmestani et al., 2023; Kolt et al., 2025). Control becomes useful when it organizes action as bounded inquiry, not when it pretends to command a social system as a fixed mechanism.

The object class is the event-instrumented socioeconomic system (EISES): a bounded local governance regime in which relevant interactions leave digital traces, authorized interventions exist, feedback cadence is meaningful, constraints can be stated, authority is institutionally located and consequences can be reviewed. Examples include portfolio-intake workflows, inspection routing, procurement monitoring, milestone supervision, health service queues and public innovation support processes. These settings are neither classical physical plants nor generic social systems. They are local regimes where knowledge claims can be made operational and auditable.

The article makes four contributions. First, it defines EISES as a local-regime object for scientific governance. Second, it reframes intervention as an epistemic query: an action tests whether a local regime behaves as claimed under specified conditions. Third, it proposes the auditable local knowledge bundle (ALKB) as the minimum evidence unit for governance learning. Fourth, it develops KAM-EISES, a knowledge audit maturity model whose strict rule is

$$G \leq \min\{O, F, V, L, Q\}, \quad (1)$$

where governance use G must not exceed the weakest support among local observability O, falsifiability F, statistical verification V, legitimacy L and auditability Q. The rule is not a substitute for democratic judgement. It is a guard against epistemic overreach.

1. METHODS

1.1. Review design and analytical purpose

The study is a structured critical review with theory building intent. It does not estimate an empirical effect size and is not presented as a PRISMA meta-analysis. Its purpose is to construct a usable unit of judgement for feedback enabled socioeconomic governance: when can an intervention produce local knowledge that is falsifiable, statistically verified, legitimate and auditable enough to support a governance consequence?

The review used a critical-sufficiency corpus. Sources were retained when they supported at least one load-bearing element of the framework: local-regime definition, endogenous observability, supervisory restraint, process-trace evidence, uncertainty reasoning, sensitivity and identifiability, complex-regime interpretation, causal verification, runtime assurance, AI governance, legitimacy or auditability. Re- cent peer-reviewed articles and annual-review venues were prioritized; older sources were retained only when they remained concept-defining.

1.2. Search strategy and reproducibility

The search used disciplinary families rather than a single bibliometric string because the object crosses control engineering, AI governance, public administration, causal inference, complex systems, process analytics and evaluation science. Query families combined terms related to societal control, supervisory control, stochastic adaptive control, complex adaptive systems, process mining, auditability, runtime assurance, responsible AI, causal policy evaluation and innovation

ecosystems.

Inclusion required three conditions: the source was peer- reviewed, archival or a recognized scholarly book; it supported a necessary element of the framework rather than general background; and it clarified a transfer condition, evidentiary requirement or failure mode relevant to governance in instrumented socio-economic regimes. Exclusion removed non-archival commentary, domain-specific work without transferable relevance and sources superseded by stronger recent reviews.

1.3. Evidence families

Table 4 groups the load-bearing evidence families used in the synthesis. The table is descriptive rather than bibliometric: it identifies the function each literature family serves in building a transferable governance framework.

1.4. Synthesis procedure

The synthesis proceeded in four steps. First, operational elements were separated from epistemic elements. Operational elements include traces, state proxies, action records, admissibility checks and decision cadence. Epistemic elements include claims, comparators, uncertainty, validity threats, falsification criteria and governance consequences.

Table 1. Operational use of the framework.

Step	Operational task
1	Delimit the local governance regime: population, institutional context, trace schema, decision cadence, authority and validity interval.
2	Register the governance claim: target indicators, comparator, expected contrast, practical effect, uncertainty and failure condition.
3	Specify admissible interventions: action space, constraint checks, escalation conditions and rollback triggers.
4	Collect trace, outcome and provenance records: baseline state, action record, outcome record, missingness, drift status and rule or model version.
5	Verify the local contrast: estimator, uncertainty interval, sensitivity check, validity threat and subgroup-risk assessment.
6	Assign only a supported consequence: monitor, adjust, test, scale, rollback or escalate, with G bounded by the weakest support axis.

Table 2. Concept map for interdisciplinary readers.

Concept	Plain-language role in the paper
EISES	A bounded process where institutional traces, authorized actions and reviewable consequences coexist.
ALKB	The minimum evidence record linking a claim, an intervention, a comparator, uncertainty, constraints and audit trail.
KAM-EISES	A maturity rule that limits the consequence of a decision by the weakest evidence support.
Bounded steerability	Local ability to move selected indicators under stated constraints, not a claim of full social control.
Endogenous observability	Recognition that traces are shaped by incentives, categories, recording practices and prior governance choices.
Weakest-link rule	A decision cannot be more credible than what can be observed, falsified, verified, authorized and audited.

Second, recurring failure modes were identified: fixed-plant transfer, KPI overreach, endogenous measurement, strategic adaptation, objective drift, unsupported AI delegation, missing provenance and weak inference. Third, each failure mode was converted into a design requirement. Fourth, the requirements were assembled into two artifacts: the ALKB evidence schema and the KAM-EISES maturity model. The framework should be judged by construct clarity, internal coherence, cross-disciplinary coverage and its ability to distinguish defensible governance use from evidentiary overreach.

Table 3. Mini-table for reproducibility of the structured critical review.

Item	Specification
Review type	Structured critical review with theory-building purpose; not a PRISMA meta-analysis and not an effect-size aggregation.
Search window	Searches and backward-forward citation checks were conducted between 15 February 2026 and 15 April 2026.
Search channels	Scopus, Web of Science Core Collection, IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, Taylor and Francis Online, Google Scholar for citation chaining, and publisher pages for DOI verification.
Disciplinary query families	Societal and operational control; supervisory and stochastic adaptive control; complex adaptive and socio-technical systems; process mining and trace auditability; causal inference and policy evaluation; responsible AI and public-sector governance; innovation ecosystems and institutional support.
Inclusion basis	Peer-reviewed article, archival venue or recognized scholarly book; direct support for a necessary framework element; explicit contribution to a transfer condition, evidence requirement or governance failure mode.
Exclusion basis	Non-archival commentary, purely rhetorical AI-governance material, domain-specific sources without transferable relevance, duplicate records and sources superseded by stronger recent reviews.
Corpus endpoint	Final retained sources are the load-bearing references listed in the bibliography; the selection was judged by conceptual sufficiency, disciplinary coverage and relevance to auditable governance use.
Reproducibility record	The project package includes a search-strategy summary and an editable corpus sheet template so that future empirical replications can add database-specific counts, screening decisions and reviewer agreement metrics.

Table 4. Load-bearing evidence families of the review corpus.

Corpus slice	Role	Analytical use
Societal and operational control	Outer scope	Feedback, constraints, cadence and intervention beyond industrial plants (Lamnabhi-Lagarrigue et al., 2017; Jia et al., 2022; Stapleton et al., 2024; Anderies and Mathias, 2025).
Supervisory and adaptive control	Operational core	Admissibility, uncertainty, adaptation and event-triggered action under constraints (Ramadge and Wonham, 1987; Åström, 1983; Kumar, 1985; Mesbah, 2016; Tabuada, 2007).
Complex adaptive and socio-technical systems	Object class	Adaptive, endogenous, nonlinear and partially observable regimes (Preiser et al., 2018; Battiston et al., 2021; Polojärvi et al., 2023; Kolt et al., 2025).
Process mining, observability and audit from traces	Evidence layer	Trace reconstruction, object-centric analytics, predictive monitoring, conformance and process evidence (Teinemaa et al., 2019; Galanti et al., 2023; Rinderle-Ma et al., 2023; Adams et al., 2023; van der Aalst et al., 2024).
Causal inference and statistical verification	Verification layer	Estimands, comparators, uncertainty intervals, sensitivity analysis and sequential evidence (Pearl, 2009; Athey and Imbens, 2017; Howard et al., 2021; Abadie, 2021; Rambachan and Roth, 2023).
Responsible AI and public-sector governance	Legitimacy layer	Accountability, oversight, redress, contestability and public trust (Janssen et al., 2020; Gritsenko and Wood, 2022; Laux, 2024; Papagiannidis et al., 2025).
Innovation ecosystems and institutional support	Stress-test class	Event-rich but normatively contested socio-economic application settings (Cooke et al., 1997; Schot and Steinmueller, 2018; Granstrand and Holgersson, 2020).

2. RESULTS I: EVENT-INSTRUMENTED SOCIO-ECONOMIC SYSTEMS AS LO- CAL REGIMES

2.1. Why the object is not a weakened industrial plant

Digitization can make an administrative process look like a control problem, but the similarity is partial. A hospital pathway, inspection queue, procurement workflow or innovation-support call has states, actions, disturbances and outputs. Yet the meanings of those variables are institutional. Actors learn the rules, categories are revised, documentation practices shift, incentives reshape recorded be- haviour and objectives remain contestable. A control representation that ignores this reflexivity can be formally precise and scientifically invalid.

The implication is constructive. The admissible strategy is supervisory, stochastic, adaptive and evidence-producing, not classical regulation of a stable mechanism. It should specify what is observable, what is authorized, what can be compared, what uncertainty remains and what consequence the evidence can support.

Definition 1 (Event-instrumented socio-economic system). An EISES is a bounded socio-economic complex adaptive system in which: (i) relevant interactions leave time-stamped digital traces with operational meaning; (ii) actionable intervention points exist; (iii) the cadence of measurement, decision and intervention makes feedback meaningful; (iv) part of the admissible constraint set can be stated; (v) the authority to instrument and intervene is institutionally

located; and (vi) consequences of intervention are reviewable through logs, records or formal procedures.

This definition excludes broad claims about controlling society. It applies to local governance regimes within wider systems: a unit, program, workflow, call, queue or regulatory process where evidence can be produced and reviewed.

2.2. Local governance regimes

Definition 2 (Local governance regime). A local governance regime is a time-indexed configuration composed of an operational population or case set, an institutional con- text, an observation and recording regime, an authorized intervention space, an admissible constraint set, a decision cadence, an accountable authority structure and a temporal validity interval.

Locality is not a weakness. It is the condition for hon- est inference. A routing rule validated in one call, legal framework, budget cycle, trace schema or evaluator pool is not automatically valid after reform, migration or incentive change.

2.3. Endogenous observability

Definition 3 (Endogenous observability). Observability is endogenous when the observation channel depends on both the instrumentation regime and participant adaptation; that is, what becomes visible is influenced by prior interventions, incentives, documentation rules, categories and governance choices.

Event logs are institutional inscriptions rather than neutral measurement channels. What becomes inferable depends on event granularity, schema discipline, object linkage, documentation timing, missingness, exception handling and incentives. Process mining provides strong methods for discovering, monitoring and predicting process behaviour from event traces (Teinemaa et al., 2019; Galanti et al., 2023; Rinderle-Ma et al., 2023; van der Aalst et al., 2024). In an EISES, however, the operative question is narrower: can the target state required by a bounded claim be recovered with documented uncertainty?

2.4. Bounded steerability rather than global controllability

Classical controllability is too strong for socio-economic governance. The relevant construct is bounded steerability of selected regime indicators under authority, uncertainty and constraint.

Definition 4 (Bounded steerability). A local regime is boundedly steerable for a claim when authorized interventions can move or maintain selected target indicators within an admissible region, with stated uncertainty and review- able justification, without presuming reachability of the full latent socio-economic state.

This construct remains compatible with network control and target-state estimation (Liu et al., 2011, 2013; Montanari et al., 2022; Motter, 2015), but it does not collapse governability into graph controllability. Nodes may be organizations, citizens, professionals, documents, rules or plat- forms, and their couplings may be strategic, normative, le- gal or informational. Higher-order interactions further limit pairwise reasoning (Battiston et al., 2021). Governance is therefore local regime steering under explicit constraints.

Table 5. From plant-centric control to local-regime scientific governance.

Dimension	Plant-centric reading	Local-regime reading for EISES
Object	Stable plant with fixed transition and measurement semantics	Co-evolving regime of actors, traces, rules, categories, authority and objectives
Measurement	Exogenous observation channel	Institutional trace shaped by schemas, incentives, missingness and adaptation
Intervention	Input applied to a plant	Authorized governance action that also functions as a query to the regime
Objective	Fixed set-point or performance criterion	Bounded, contestable and policy-conditioned target with explicit validity interval
Validation	Repeatable benchmark or experiment under stable assumptions	Local contrast with assumptions, uncertainty, constraint checks and audit trail
Learning	Model refinement for a fixed plant	Accumulation of auditable local knowledge bundles under changing regimes

3. RESULTS II: CONTROL AS AU- DITABLE KNOWLEDGE GENERA- TION

3.1. From control action to epistemic query

In engineering, a control action is often evaluated by regulation, tracking, stability or constraint satisfaction. In socioeconomic governance these criteria are necessary but in- complete. A governance action is also an epistemic query: it asks whether a local regime behaves as claimed when an authorized intervention is applied under specified

conditions. Without a claim, comparator and evidentiary record, action produces activity but not knowledge.

Given a local regime, a claim, an admissible action set and a constraint set, the task is to choose, defer, block or escalate an intervention so that evidentiary value is produced without exceeding legitimacy and auditability. The evidence must answer five questions: what was believed, what was observable, what was allowed, what happened and how strongly the claim survived comparison.

Proposition 1 (Control-generated governance knowledge requires contrast). A feedback intervention in an EISES produces governance-relevant knowledge only if it is linked to a pre-specified or reconstructable contrast between expected and observed local regime behaviour.

Proof sketch. Without contrast, a post-intervention observation can be attributed after the fact to the action, regression toward the mean, exogenous change, selection, measurement drift or strategic adaptation. A contrast does not remove all validity threats, but it creates empirical vulnerability.

3.2. Supervisory restraint as a governance form

The conservative form proposed here separates five operations that are often collapsed in algorithmic governance: estimation, admissibility checking, intervention selection, evidence updating and consequence assignment. It can recommend action, inaction, rollback or escalation, but only at the support level justified by local evidence.

A supervisory decision must make reviewable the estimated condition, uncertainty, risk class, admissible options, blocked options and reasons for escalation or rollback. Event-triggered scheduling is useful because not every observation warrants intervention; some observations warrant monitoring, delayed review or additional evidence (Tabuada, 2007; Gama et al., 2014). The resulting decision menu is intentionally small: monitor, adjust, test, scale, rollback or escalate.

3.3. Analytical tools translated into evidence functions

The theoretical tools used in the framework have operational roles. Information theory asks what uncertainty about a claim is reduced. Complex-systems and statistical-physics reasoning prevents over-atomized interpretation by treating some outcomes as collective regime patterns. Information geometry asks whether the dimensions supporting a decision are identifiable or sensitive under the current observation design. These tools are not presented as implementation requirements; they are used as evidence functions that discipline interpretation.

The framework therefore treats analytic sophistication as insufficient by itself. A refined model may remain unusable if the claim is not falsifiable, if the comparator is missing, if the decision exceeds authority, or if the audit trail cannot reconstruct what happened.

Falsifiable claim → local-regime registration → intervention protocol → trace and constraint record → statistical verification → provenance and audit → governance update → bundle repository.

Figure 1. Control as an epistemic feedback loop. Each intervention is treated as a local query whose evidentiary residue is stored as an auditable local knowledge bundle.

Table 6. Analytical tools translated into governance-evidence functions.

Analytical tool	Contribution	Operational question
Information theory	Uncertainty, discrimination among claims and minimum-assumption reasoning (Jaynes, 1957).	Did the intervention reduce uncertainty about the registered claim without adding unjustified structure?
Complex-systems and statistical-physics reasoning	Coarse-grained states, fluctuations, collective dynamics and higher-order interactions (Anderson, 1972; Thurner et al., 2018; Battiston et al., 2021; Ghavasieh et al., 2020; Kaniadakis et al., 2024).	Is the observed pattern a local regime signal rather than an artifact of isolated indicators?
Information geometry	Local sensitivity, identifiability and model-manifold reasoning (Amari and Nagaoka, 2000; Quinn et al., 2023; Wada and Scarfone, 2024).	Does the decision depend on well-identified directions, or on weakly observable dimensions?
Causal inference	Estimands, comparators, uncertainty intervals and robustness checks (Pearl, 2009; Athey and Imbens, 2017; Abadie, 2021; Rambachan and Roth, 2023).	Against which reference condition can the claim be evaluated?
Responsible AI and public governance	Oversight, redress, contestability, proportionality and public trust (Gritsenko and Wood, 2022; Ulbricht and Yeung, 2022; Laux, 2024; de Fine Licht, 2025).	Is the evidence used within legitimate authority and reviewable procedure?

4. RESULTS III: AUDITABLE LOCAL KNOWLEDGE BUNDLES AND MATU- RITY

4.1. Why event logs are not enough

Event traces alone are operational records, not

scientific evidence. They can show that an action occurred and an out- come followed, but not necessarily which claim was tested, which comparator was intended, which constraints were checked, which uncertainty was known or whether the decision was legitimate. The missing object is an evidence unit that links action, claim, verification and review.

Definition 5 (Auditable local knowledge bundle). An ALKB is a versioned evidence object composed of the local regime, governance claim, intervention and verification protocol, baseline record, action or exposure record, out- come record, evaluated contrast, statistical verification object, constraint-check record, provenance chain and governance decision supported by the bundle.

The bundle is an evidentiary schema, not an implementation description. It specifies what must exist for a decision to be reviewed, criticized, replicated locally or com- pared across regimes. Its practical function is to deter- mine whether evidence supports monitoring, adjustment, controlled testing, targeted scaling, rollback or escalation.

4.2. Auditability, estimands and decision-relative verifica- tion

Auditability is stronger than transparency. A system can publish general rules and still fail auditability if a specific decision cannot be reconstructed. For EISES, auditability requires reconstructing the local regime, identifying the claim and failure criteria, connecting action to evidence, inspecting validity threats and tracing how evidence sup- ported the decision. Process-mining audit research and data-governance literature converge on the need for lineage, role clarity and reconstructable evidence (Fo̊hr et al., 2025; Eulerich et al., 2025; Janssen et al., 2020; Longpre et al., 2024).

A governance claim also requires a local estimand. Let $Y_i(u)$ denote the potential outcome under intervention u in regime R_i, with comparator u_0. A basic local contrast is

$$\Delta_i = E\left[Y_i(u) - Y_i(u_0) \mid R_i\right]. \qquad (2)$$

Designs may include randomized or stepped allocation, regression discontinuity, difference-in-differences, interrupted time series, synthetic controls, sequential monitoring and shadow-mode comparison (Athey and Imbens, 2017; Abadie, 2021; Ben-Michael et al., 2021; Rambachan and Roth, 2023; Howard et al., 2021). Subgroup discovery re- quires caution because it can inflate false positives or conceal distributional harm (Wager and Athey, 2018; Ben- jamini and Hochberg, 1995).

Verification is decision-relative. Evidence sufficient to refine a dashboard is not sufficient to impose a sanction, allocate scarce benefits, reconfigure a program or scale a policy. The evidentiary burden increases with consequence, irreversibility, contestability and distributive impact.

4.3. KAM-EISES: a knowledge-audit maturity model

KAM-EISES evaluates whether a governance consequence is supported by evidence. It uses five support axes: local observability (O), falsifiability (F), statistical verifiability (V), legitimacy (L) and auditability (Q). Governance use (G) captures the decision consequence attached to evidence. Each axis is scored from 0 to 4.

Define support capacity as:

$$S = \min\{O, F, V, L, Q\}. \qquad (3)$$

Table 7. Minimum ALKB schema and audit questions.

Component	Required content	Audit question
Regime identifier	Population, institutional context, trace schema, cadence and validity interval	What local regime does the evidence describe?
Governance claim	Falsifiable claim, target indicators, expected contrast, practical effect and failure criteria	What proposition could the evidence weaken or reject?
Eligibility and comparator	Inclusion/exclusion, subgroup definitions, assignment logic, status quo, threshold, randomization, synthetic comparison or registered baseline	Against what reference condition was the action evaluated?
Intervention record	Action class, timing, intensity, authority, override and escalation events	What was done, by whom and under which authority?
Baseline and outcome records	Pre-intervention proxies, missingness, drift status, target outcomes, adverse events and delayed outcomes	What was known before action and what happened after?
Statistical verification	Estimand, estimator, uncertainty interval, sensitivity checks, multiplicity handling and validity threats	How strong is the local contrast?
Constraint checks	Legal, procedural, privacy, fairness, safety, resource and service-level checks	Did the action remain admissible while producing evidence?
Provenance and decision	Data lineage, rule or model version, reviewer, timestamp, reproducibility pointer and decision menu	Can the decision path be reconstructed?

Table 8. Maturity logic in KAM-EISES.

Level	Governance use	Interpretation
M0	$G = 0$	Descriptive setting. Evidence may support exploration but not governance action.
M1	$G = 1$, support ≥ 1	Evidence informs deliberation, dashboard interpretation or hypothesis refinement.
M2	$G = 2$, support ≥ 2	Evidence supports bounded operational adjustment under authorization, basic verification and reproducible records.
M3	$G = 3$, support ≥ 3	Evidence supports targeted rollout, rollback, resource allocation or policy-family adaptation with safeguards and rich provenance.
M4	$G = 4$, support $= 4$	Evidence supports bounded semi-autonomous governance only with drift-aware observability, registered claims, high-grade verification, contestability, complete audit chain, assurance and escalation.

A configuration is knowledge-admissible if and only if $G \leq S$. (4)

The overreach gap is $\Delta^+ = \max\{0, G - S\}$. If $\Delta^+ > 0$, the decision consequence exceeds at least one support condition and should be de-escalated, redesigned or routed to accountable review.

Proposition 2 (Weakest-link principle). In an EISES, governance use of control-generated knowledge cannot be more credible than the weakest support condition among observability, falsifiability, statistical verification, legitimacy and auditability.

Proof sketch. If $G > O$, the decision outruns trace recoverability. If $G > F$, the claim is insufficiently vulnerable to evidence. If $G > V$, the contrast is statistically inadequate for the consequence. If $G > L$, evidence is used beyond institutional authority. If $G > Q$, the decision cannot be re-constructed. Any one failure converts scientific governance into unsupported assertion.

Remark 1 (Rollback as mature governance). Reducing governance use when support capacity degrades is not failure. Drift, schema migration, legal reform, contested evidence, audit failure or subgroup harm should trigger de- escalation. In EISES, rollback is a normal control action that preserves scientific validity.

5. DISCUSSION

5.1. What the framework changes

The framework changes the unit of reasoning. KPI- centered governance asks whether an indicator moved. Auditable scientific governance asks whether a registered claim survived a local contrast under admissible conditions. This distinction matters because KPI targets can induce gaming, strategic adaptation and dashboard theatre (Thomas and Uminsky, 2022; Manheim, 2023). Metrics remain useful, but only when embedded in claims, comparators, uncertainty and reviewable records.

The framework also clarifies the role of AI. Machine learning is valuable for classification, prediction, anomaly detection, ranking and pattern extraction, but governance is not only prediction. A model can predict delay without establishing whether an intervention is authorized, causal, fair or auditable. Responsible AI research therefore emphasizes structural and procedural governance, not accuracy alone (Janssen et al., 2020; Papagiannidis et al., 2025; Agbabi- aka et al., 2025; Miedema et al., 2026). In the proposed framework, AI outputs become acceptable only when they map to bundle components: claim, comparator, uncertainty, constraints, provenance and decision use.

Table 9. KPI governance versus falsifiable auditable governance.

Issue	KPI/intuitive governance	Falsifiable auditable governance
Unit of reasoning	Indicator movement or expert impression	Local claim with comparator, horizon, uncertainty and failure condition
Main evidence	Dashboard value, ranking or trend	ALKB linking regime, action, contrast, verification, constraints and provenance
Role of AI	Prediction or recommendation treated as sufficient decision support	Model output treated as one artifact inside a reviewable evidence chain
Risk	Metric gaming, confounding and automation bias	Explicit validity threats, drift checks, constraint evidence and rollback conditions
Decision consequence	Often escalates by managerial confidence	Bounded by weakest support condition in KAM-EISES
Learning	Success stories and informal memory	Versioned repository including negative evidence and retired claims

5.2. Why supervisory restraint is the conservative option

The central design implication is restraint. Supervisory governance separates estimation, admissibility checking, intervention selection, evidence updating and consequence assignment. It can recommend or execute only within registered bounds, and it can also choose inaction, escalation or rollback when uncertainty, risk or contestability is too high. This structure is compatible with control theory with- out requiring plant-like assumptions. Supervisory control contributes admissibility and forbidden-transition logic (Ramadge and Wonham, 1987); adaptive control con- tributes bounded retuning under drift (Åström, 1983; Ku- mar, 1985); stochastic predictive control contributes finite-horizon reasoning under uncertainty (Mesbah, 2016); event- triggered control contributes cadence discipline (Tabuada, 2007); runtime assurance contributes evidence that action remained inside an authorized envelope (Hobbs et al., 2023; Hsu et al., 2024). The integrated result is a governance posture: monitor, adjust, test, scale, roll back or escalate only at the level that local evidence can support.

5.3. Implementation-independent transfer of control concepts

The framework is intentionally implementation-independent. It does not depend on a particular software stack, measurement pipeline or simulation environment. It requires every technical artifact to contribute reviewable evidence.

5.4. What counts as accumulated knowledge

Accumulated knowledge is not a growing collection of dash- boards. It is a versioned archive of local claims, contrasts, failures, uncertainty intervals, drift episodes, constraint vi- olations, appeals, rollbacks and successful transfers. An institution learns when it can say what worked, where, under which regime, against which comparator, with what uncertainty and under which constraints.

The ALKB creates that memory. Its value is highest when it preserves negative evidence: a failed intervention with clear assumptions and audit trail is scientifically more valuable than a successful-looking KPI without provenance. Over time, bundle repositories can support meta-learning about which claims are portable, which regimes are fragile, which indicators are unstable and which governance consequences tend to exceed evidence.

5.5. Ethical and political boundaries

No formal design removes normative judgement. Legitimacy, contestability and redress remain constitutive. Control is therefore treated as a discipline for bounded evidence generation, not as a substitute for public authority. The weakest-link rule is intentionally strict because it prevents technical performance from bypassing institutional account- ability. In contested socio-economic regimes, a decision is mature when its evidentiary basis can be challenged, recon- structed and revised.

Table 10. Control concepts as evidence contributors.

Control concept	Evidence contribution	Required reinterpretation in EISES
Feedback	Links action, observation, updating and consequence	Feedback is claim testing, not only performance correction.
Supervision	Filters admissible actions and forbidden transitions	Safety includes legality, fairness, privacy, procedure and escalation rights.
Stochasticity	Represents uncertain states, effects and outcomes	Uncertainty is a governance constraint.
Adaptation	Updates parameters or policy families under drift	Adaptation must remain authorized and leave provenance.
Event triggering	Determines when evidence warrants action or review	Cadence must respect latency, risk and audit burden.
Runtime assurance	Preserves bounded operation under uncertainty	Assurance must include reconstructable evidence and contestability.
Analytical evidence tools	Clarify uncertainty reduction, sensitivity and collective patterns	Technical outputs are admissible only when interpreted inside a local validity regime.

6. APPLICATION CLASS: INNOVATION ECOSYSTEMS AND INSTITUTIONAL SUPPORT PROGRAMS

Innovation ecosystems are a useful stress test because they combine heterogeneous actors, complementary assets, path dependence, institutional layering and contested definitions of value (Cooke et al., 1997; Adner, 2017; Jacobides et al., 2018; Schot and Steinmueller, 2018; Granstrand and Hol- gersson, 2020). They are event-rich: calls, applications, eligibility checks, triage, evaluation, contracting, milestone review, partner matching, reporting, payments and follow- up all generate records. They are also risky: the presence of operational traces can create the illusion that innovation itself is controllable.

The framework avoids that overclaim. An innovation ecosystem as a whole is not an EISES. Novelty, complementarities, welfare, territorial justice and entrepreneurial discovery cannot be reduced to short-run operational indicators without normative loss, especially where innovation measurement depends on conventions about novelty, implementation and diffusion (OECD and Eurostat, 2018). Several subdomains can nevertheless qualify as local regimes: portfolio intake and screening, evaluation and contracting, milestone monitoring, intermediation and support routing, compliance review and operational resilience. These sub- domains are bounded processes through which institutions can learn and make better-governed decisions about timing, allocation, escalation, support routing and rollback.

Consider a reviewer-assignment protocol for an innovation-support call. The regime is defined by the call text, reviewer pool, conflict-of-interest rule, trace schema and decision cadence. Suppose the claim, comparator, eligibility, outcome horizon and failure criteria are registered; authority, override, escalation and review procedures are codified; and rule versions, traces, constraint checks and escalation reasons are reconstructable. If the statistical design is still a historical

comparator rather than a stepped rollout or threshold-based design, a plausible KAM-EISES score is

$$O = 3,\ F = 3,\ V = 2,\ L = 3,\ Q = 3, \quad (5)$$

so $S = 2$. A targeted adaptive rollout at $G = 3$ would be overreach, whereas bounded operational adjustment under human authorization at $G = 2$ would be admissible. The example shows that stronger analytics cannot compensate for weak verification.

Evidence accumulation then proceeds through a bundle repository. It should store successful interventions, null effects, failed rollouts, constraint violations, contested evidence, drift episodes and retired claims. A time-to-contract KPI may improve because the protocol helped, because ineligible cases were excluded faster, because appeals were sup- pressed, because documentation changed or because the call attracted easier applications. Only a bundle can distinguish among these possibilities. Rigorous innovation governance therefore requires evidence structure, not only innovation dashboards.

Table 11. Illustrative claims for innovation-support EISES subdomains.

Claim type	Evidence-compatible formulation	Why it works
Intake workflow	A revised completeness-check protocol reduces median time-to-contract by at least 15% without increasing appeal or exclusion-error rates above the registered bound.	Local regime, intervention, outcome, comparator, constraint and failure condition are explicit.
Reviewer assignment	For eligible proposals in segment s, assignment rule u reduces conflict bottlenecks relative to the registered comparator without increasing unresolved conflicts.	It tests a bounded coordination mechanism rather than claiming to control innovation.
Milestone monitoring	Early risk flag r predicts delayed milestone completion with calibrated uncertainty sufficient for human escalation but not automatic sanction.	Prediction quality is linked to a bounded governance use.
Support routing	For eligible firms in segment s, support-routing protocol u reduces time-to-first-service by δ days without increasing referral loops.	It produces a local contrast on a process mechanism.
Overstretched claim	The platform increases regional innovation.	No local regime, comparator, horizon, failure criterion or validity envelope is specified.
Overdelegated claim	The system should autonomously choose the most valuable projects.	It collapses contested public value into ranking unless legitimacy, verification, auditability and appeal are high.

Table 12. Worked ALKB example for a reviewer-assignment protocol.

ALKB field	Example entry
Local regime	2026 innovation-support call, reviewer-assignment stage, fixed call text, defined reviewer pool and documented conflict-of-interest rule.
Claim	A revised assignment rule reduces median time-to-contract by at least 15% without increasing appeals or exclusion-error rates above a registered bound.
Comparator	Historical rule, stepped rollout, threshold-neighbour comparison or another registered baseline selected before consequence assignment.
Intervention	Assignment rule u with conflict-of-interest screening, documented overrides and escalation paths.
Outcomes	Time-to-contract, unresolved conflicts, appeals, exclusion-error rate, reviewer workload and delayed contracting events.
Verification	Local contrast with uncertainty interval, missingness assessment, sensitivity check and multiplicity handling for subgroups.
Constraints	Conflict-of-interest compliance, procedural fairness, appeal rights, privacy, workload limits and service-level requirements.
Governance use	$G = 2$ for bounded operational adjustment unless verification, legitimacy and auditability jointly justify a higher consequence.
Rollback trigger	Appeals, unresolved conflicts, exclusion-error rate or subgroup harm exceed the registered bound; trace schema or authority changes invalidate the local evidence.

7. LIMITATIONS AND RESEARCH AGENDA

7.1. Scope and non-claims

The paper does not claim that socio-economic systems are globally controllable, that event data reconstruct a full la- tent state, or that any single implementation can solve public governance. Its narrower claim is that, in bounded local regimes, non-classical control concepts can convert interventions into auditable knowledge bundles with explicit falsification conditions, statistically defensible uncertainty and decision consequences constrained by local support capacity.

The framework is conceptual and requires

empirical calibration. The maturity scores are ordinal and should not be treated as precise measurements. Inter-rater reliability, domain-specific scoring rubrics and validation studies are needed before KAM-EISES is used for high-consequence certification. The ALKB schema is a minimum evidence unit, not a complete institutional policy.

7.2. Research agenda

Four priorities follow. First, scoring and validation work should develop maturity-axis rubrics, example libraries and inter-rater reliability protocols for observability, falsifiability, verification, legitimacy and auditability. Second, causal designs for constrained governance learning should be adapted to domains where randomized trials are difficult, including stepped rollout, threshold comparison, interrupted time series, synthetic controls, sequential monitoring and shadow-mode evaluation (Abadie, 2021; Ben-Michael et al., 2021; Rambachan and Roth, 2023; Howard et al., 2021). Third, trace and drift governance should model how process data change after interventions, how strategic response affects recoverability and how schema migration changes the validity interval. Fourth, assurance and rollback research should connect stability, legality, privacy, fairness, security, accountability and statistical validity into review- able cases, with de-escalation protocols when support capacity degrades (Hobbs et al., 2023; Liu et al., 2024; Shahandashti et al., 2024; Paterson et al., 2025).

Progress will come less from declaring more autonomy and more from producing better local evidence. Control contributes when it disciplines intervention as knowledge- generating action under explicit constraints and when analytical tools are tied to falsifiable decisions rather than used as decorative analytics.

8. CONCLUSIONS

Digital instrumentation has changed the evidentiary status of many socio-economic processes. Public and institutional systems now generate traces that enable feed- back, prioritization, prediction and assurance. The scientific mistake is to confuse this visibility with governability. Event-instrumented socio-economic systems remain com- plex adaptive regimes in which traces, categories, objectives, incentives and authority co-evolve.

This article argued that governance in such regimes must move beyond isolated KPIs, unaudited expert judgement and AI prediction used as implicit authority. The defensible standard is auditable scientific governance: interventions should be formulated as local, falsifiable and statistically checkable claims; evidence should include comparators, uncertainty, constraints, provenance and decision con- sequences; and governance use should be limited by observability, verification, legitimacy and auditability.

Supervisory restraint was identified as a plausible and conservative form for this setting. Its value is not that it removes human judgement or automates sovereignty. Its value is that it screens admissible action, represents uncertainty, adapts inside authorized policy families, preserves evidence and routes unsupported or contested cases toward escalation or rollback. The strongest contribution of control is therefore not classical command of a fixed plant, but disciplined production of local knowledge under complex adaptive conditions.

The paper introduced the ALKB as the minimum evidence unit for this program and KAM-EISES as the maturity model that limits governance consequence by the weak- est support condition: $G \leq \min\{O, F, V, L, Q\}$. This rule is intentionally conservative. A decision cannot be more credible than what can be observed, falsified, statistically verified, authorized and audited.

The broader contribution is a bridge between control and evidence culture. Control provides feedback, intervention and assurance; complex-systems reasoning provides collective-regime interpretation; information theory pro- vides uncertainty and discrimination among claims; sensitivity and identifiability analysis clarifies local learnability; causal inference provides comparators and verification; re- sponsible AI provides legitimacy and contestability. Institutions become more scientific when they accumulate auditable local knowledge, preserve failures as evidence and refuse to let authority outrun support.

FUNDING

This research did not receive any specific grant from funding agencies in the public, commercial or not-for-profit sectors.

AUTHOR CONTRIBUTIONS

Jorge Omar Portilla Jaimes: Conceptualization, methodology, literature curation, formal analysis, writing–original draft, project administration. Oscar J. Suarez: Method- ology, control-systems interpretation, validation, writing, review and editing. José Aguilar Castro: Conceptualization, complex-systems and AI-governance framing, supervision, writing–review and editing. Oliver López-Corona: Methodology, statistical and complex-systems interpretation, vali- dation, writing–review and editing. All authors have

read and agreed to the submitted version of the manuscript.

DECLARATION OF COMPETING INTEREST

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

DATA AVAILABILITY

No primary experimental dataset was generated for this conceptual review. The project package includes the manuscript source, a search-strategy summary, an editable corpus-sheet template, an auditable local knowledge bundle schema and a KAM-EISES maturity-scoring rubric intended to support replication of the conceptual assessment.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used OpenAI ChatGPT for language refinement and consistency checking. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

ACKNOWLEDGEMENTS

The authors thank their respective academic institutions and research groups for the scholarly environment that sup- ported the development of this review.

REFERENCES

1. Abadie, Alberto. (2021). Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects. *Journal of Economic Literature*, 59(2), 391–425. doi:10.1257/jel.20191450.
2. Adams, Jan Niklas, van Zelst, Sebastiaan J., Rose, Thomas and van der Aalst, Wil M. P. (2023). Explainable Concept Drift in Process Mining. *Information Systems*, 114, 102177. doi:10.1016/j.is.2023.102177.
3. Adner, Ron. (2017). Ecosystem as Structure: An Actionable Construct for Strategy. *Journal of Management*, 43(1), 39–58. doi:10.1177/0149206316678451.
4. Agbabiaka, Olusegun, Ojo, Adegboyega and Connolly, Niall. (2025). Requirements for Trustworthy AI-Enabled Automated Decision-Making in the Public Sector: A Systematic Review. *Technological Forecasting and Social Change*, 215, 124076. doi:10.1016/j.techfore.2025.124076.
5. Amari, Shun-ichi and Nagaoka, Hiroshi. (2000). *Methods of Information Geometry*. American Mathematical Society, Providence, RI. Anderson, Philip W. (1972). More Is Different. *Science*, 177(4047), 393–396. doi:10.1126/science.177.4047.393.
6. Anderies, John M. and Mathias, Jean-Denis. (2025). Leveraging Control Theory Tools to Enable Real-Time Policy Action for Sustainable Social-Ecological-Technical Systems. *Communications Earth and Environment*, 6, 806. doi:10.1038/s43247-025-02767-3.
7. Aoki, Naomi, Tatsumi, Tomohiko, Naruse, Go and Maeda, Kentaro. (2024). Explainable AI for Government: Does the Type of Explanation Matter to the Accuracy, Fairness, and Trustworthiness of an Algorithmic Decision as Perceived by Those Who Are Affected? *Government Information Quarterly*, 41(4), 101965. doi:10.1016/j.giq.2024.101965.
8. Åström, Karl J. (1983). Theory and Applications of Adaptive Control–A Survey. *Automatica*, 19(5), 471–486. doi:10.1016/0005-1098(83)90002-X.
9. Athey, Susan and Imbens, Guido W. (2017). The State of Applied Econometrics: Causality and Policy Evaluation. *Journal of Economic Perspectives*, 31(2), 3–32. doi:10.1257/jep.31.2.3.
10. Battiston, Federico, Amico, Enrico, Barrat, Alain, Bianconi, Ginestra, Ferraz de Arruda, Guilherme, Franceschiello, Benedetta, Iacopini, Iacopo, Kéfi, Sonia, Latora, Vito, Moreno, Yamir, Murray, Micah M., Peixoto, Tiago P., Vaccarino, Francesco and Petri, Giovanni. (2021). The Physics of Higher-Order Interactions in Complex Systems. *Nature Physics*, 17(10), 1093–1098. doi:10.1038/s41567-021- 01371-4.
11. Ben-Michael, Eli, Feller, Avi and Rothstein, Jesse. (2021). The Augmented Synthetic Control Method. *Journal of the American Statistical Association*, 116(536), 1789–1803. doi:10.1080/01621459.2021.1929245.
12. Benjamini, Yoav and Hochberg, Yosef. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society: Series B*, 57(1), 289–300. doi:10.1111/j.2517-6161.1995.tb02031.x.
13. Cooke, Philip, Gomez Uranga, Mikel and Etxebarria, Goio. (1997). Regional Innovation Systems: Institutional and Organisational Dimensions. *Research Policy*, 26(4–5), 475–491. doi:10.1016/S0048-7333(97)00025-5.
14. de Fine Licht, Karl. (2025). Resolving Value Conflicts in Public AI Governance: A Procedural Justice Framework. *Government Informa- tion Quarterly*, 42(2), 102033. doi:10.1016/j.giq.2025.102033.

15. Eulerich, Marc, Huang, Qing, Pawlowski, Justin and Vasarhelyi, Miklos A. (2025). Using Process Mining as an Assurance Tool in the Three-Lines-Model. *International Journal of Accounting Information Systems*, 56, 100731. doi:10.1016/j.accinf.2025.100731.
16. Fo̊hr, Tassilo L., Reichelt, Valentin, Marten, Kai-Uwe and Eulerich, Marc. (2025). A Framework for the Structured Implementation of Pro- cess Mining for Audit Tasks. *International Journal of Accounting Information Systems*, 56, 100727. doi:10.1016/j.accinf.2025.100727. Galanti, Riccardo, de Leoni, Massimiliano, Navarin, Nicolò and Marazzi, Alan. (2023). Object-Centric Process Predictive Analytics.
17. *Expert Systems with Applications*, 213, 119173. doi:10.1016/j.eswa.2022.119173.
18. Gama, João, Žliobaitė, Indrė, Bifet, Albert, Pechenizkiy, Mykola and Bouchachia, Abdelhamid. (2014). A Survey on Concept Drift Adaptation. *ACM Computing Surveys*, 46(4), 44:1–44:37. doi:10.1145/2523813.
19. Garmestani, Ahjond S., Allen, Craig R., Angeler, David G., Sundstrom, Sarah M. and Chaffin, Brian C. (2023). Multi-Scale Adaptive Management of Social-Ecological Systems. *BioScience*, 73(11), 800–807. doi:10.1093/biosci/biad096.
20. Ghavasieh, Arsham, Nicolini, Carlo and De Domenico, Manlio. (2020). Statistical Physics of Complex Information Dynamics. *Physical Review E*, 102(5), 052304. doi:10.1103/PhysRevE.102.052304.
21. Granstrand, Ove and Holgersson, Marcus. (2020). Innovation Ecosystems: A Conceptual Review and a New Definition. *Technovation*, 90–91, 102098. doi:10.1016/j.technovation.2019.102098.
22. Gritsenko, Daria and Wood, Matthew. (2022). Algorithmic Governance: A Modes of Governance Approach. *Regulation and Governance*, 16(1), 45–62. doi:10.1111/rego.12367.
23. Hobbs, Kerianne L., Mote, Mark L., Abate, Matthew C. L., Coogan, Samuel D. and Feron, Eric M. (2023). Runtime Assurance for Safety-Critical Systems: An Introduction to Safety Filtering Approaches for Complex Control Systems. *IEEE Control Systems Magazine*, 43(2), 28–65. doi:10.1109/MCS.2023.3234380.
24. Howard, Steven R., Ramdas, Aaditya, McAuliffe, Jon and Sekhon, Jasjeet. (2021). Time-Uniform, Nonparametric, Nonasymptotic Confidence Sequences. *The Annals of Statistics*, 49(2), 1055–1080. doi:10.1214/20-AOS1991.
25. Hsu, Kai-Chieh, Hu, Haimin and Fisac, Jaime F. (2024). The Safety Filter: A Unified View of Safety-Critical Control in Autonomous Systems. *Annual Review of Control, Robotics, and Autonomous Systems*, 7, 47–72. doi:10.1146/annurev-control-071723-102940.
26. Jacobides, Michael G., Cennamo, Carmelo and Gawer, Annabelle. (2018). Towards a Theory of Ecosystems. *Strategic Management Journal*, 39(8), 2255–2276. doi:10.1002/smj.2904.
27. Janssen, Marijn, Brous, Paul, Estevez, Elsa, Barbosa, Luís S. and Janowski, Tomasz. (2020). Data Governance: Organizing Data for Trustworthy Artificial Intelligence. *Government Information Quarterly*, 37(3), 101493. doi:10.1016/j.giq.2020.101493.
28. Jaynes, Edwin T. (1957). Information Theory and Statistical Mechanics. *Physical Review*, 106(4), 620–630. doi:10.1103/PhysRev.106.620.
29. Jia, Qing-Shan, Panetto, Hervé, Macchi, Marco, Siri, Silvia, Weichhart, Georg and Xu, Zhanbo. (2022). Control for Smart Systems: Challenges and Trends in Smart Cities. *Annual Reviews in Control*, 53, 358–369. doi:10.1016/j.arcontrol.2022.04.010.
30. Kaniadakis, Giorgio, Di Matteo, Tiziana, Scarfone, Antonio Maria and Gervino, Giampiero. (2024). New Trends in Statistical Physics of Complex Systems: Theoretical and Experimental Approaches. *The European Physical Journal B*, 97, 203. doi:10.1140/epjb/s10051- 024-00832-y.
31. Kolt, Noam, Shur-Ofry, Michal and Cohen, Reuven. (2025). Lessons from Complex Systems Science for AI Governance. *Patterns*, 6, 101341. doi:10.1016/j.patter.2025.101341.
32. Kumar, Panganamala R. (1985). A Survey of Some Results in Stochastic Adaptive Control. *SIAM Journal on Control and Optimization*, 23(3), 329–380. doi:10.1137/0323023.
33. Lamnabhi-Lagarrigue, Francoise, Annaswamy, Anuradha M., Engell, Sebastian, Isaksson, Alf J., Khargonekar, Pramod P., Murray, Richard M., Nijmeijer, Henk, Samad, Tariq, Tilbury, Dawn M. and Van den Hof, Paul M. J. (2017). Systems and Control for the Future of Humanity, Research Agenda: Current and Future Roles, Impact and Grand Challenges. *Annual Reviews in Control*, 43, 1–64. doi:10.1016/j.arcontrol.2017.04.001.
34. Laux, Johann. (2024). Institutionalised Distrust and Human Oversight of Artificial Intelligence: Towards a Democratic Design of AI Governance under the European Union AI Act. *AI and Society*, 39(5), 2853–2866. doi:10.1007/s00146-023-01777-z.
35. Liu, Yang-Yu, Slotine, Jean-Jacques and Barabási, Albert-László. (2011).

Controllability of Complex Networks. *Nature*, 473(7346), 167–173. doi:10.1038/nature10011.

36. Liu, Yang-Yu, Slotine, Jean-Jacques and Barabási, Albert-László. (2013). Observability of Complex Systems. *Proceedings of the National Academy of Sciences*, 110(7), 2460–2465. doi:10.1073/pnas.1215508110.
37. Liu, Hezhen, Huang, Chengqiang, Sun, Ke, Yin, Jiacheng, Wu, Xiaoyu, Wang, Jin, Zhang, Qunli, Zheng, Yang, Nigam, Vivek, Liu, Feng and Sifakis, Joseph. (2024). Design for Dependability–State of the Art and Trends. *Journal of Systems and Software*, 211, 111989. doi:10.1016/j.jss.2024.111989.
38. Longpre, Shayne, Mahari, Robert, Chen, Anthony, Obeng-Marnu, Naana, Sileo, Damien, Brannon, William, Muennighoff, Niklas, Khazam, Nathan, Kabbara, Jad, Perisetla, Kartik, Wu, Xinyi (Alexis), Shippole, Enrico, Bollacker, Kurt, Wu, Tongshuang, Villa, Luis, Pentland, Sandy and Hooker, Sara. (2024). A Large-Scale Audit of Dataset Licensing and Attribution in AI. *Nature Machine Intelligence*, 6(9), 975–987. doi:10.1038/s42256-024-00878-8.
39. Manheim, David. (2023). Building Less-Flawed Metrics: Understanding and Creating Better Measurement and Incentive Systems. *Patterns*, 4(10), 100842. doi:10.1016/j.patter.2023.100842.
40. Mesbah, Ali. (2016). Stochastic Model Predictive Control: An Overview and Perspectives for Future Research. *IEEE Control Systems Magazine*, 36(6), 30–44. doi:10.1109/MCS.2016.2602087.
41. Miedema, Emiel, Waschull, Sabine and Emmanouilidis, Christos. (2026). Towards Trustworthy Artificial Intelligence for Decision- Making: A Lifecycle Perspective on Knowledge- and Data-Driven Artificial Intelligence Systems. *Computers in Industry*, 174, 104409. doi:10.1016/j.compind.2025.104409.
42. Montanari, Arthur N., Duan, Chao, Aguirre, Luis A. and Motter, Adilson E. (2022). Functional Observability and Target State Estimation in Large-Scale Networks. *Proceedings of the National Academy of Sciences*, 119(1), e2113750119. doi:10.1073/pnas.2113750119.
43. Motter, Adilson E. (2015). Networkcontrology. *Chaos*, 25(9), 097621. doi:10.1063/1.4931570.
44. OECD and Eurostat. (2018). *Oslo Manual 2018: Guidelines for Collecting, Reporting and Using Data on Innovation*. OECD Publishing, Paris. doi:10.1787/9789264304604-en.
45. Papagiannidis, Emmanouil, Mikalef, Patrick and Conboy, Kieran. (2025). Responsible Artificial Intelligence Governance: A Review and Research Framework. *The Journal of Strategic Information Systems*, 34(2), 101885. doi:10.1016/j.jsis.2024.101885.
46. Paterson, Colin, Hawkins, Richard D., Picardi, Chiara, Jia, Yijun, Calinescu, Radu and Habli, Ibrahim. (2025). Safety Assurance of Machine Learning for Autonomous Systems. *Reliability Engineering and System Safety*, 264, 111311. doi:10.1016/j.ress.2025.111311.
47. Pearl, Judea. (2009). *Causality: Models, Reasoning, and Inference*. 2nd ed., Cambridge University Press, Cambridge.
48. Polojärvi, Dana, Palmer, Erika and Dunford, Charlotte. (2023). A Systematic Literature Review of Sociotechnical Systems in Systems Engineering. *Systems Engineering*, 26. doi:10.1002/sys.21664.
49. Preiser, Rika, Biggs, Reinette, De Vos, Alta and Folke, Carl. (2018). Social-Ecological Systems as Complex Adaptive Systems: Organizing Principles for Advancing Research Methods and Approaches. *Ecology and Society*, 23(4), 46. doi:10.5751/ES-10558-230446.
50. Quinn, Katherine N., Abbott, Michael C., Transtrum, Mark K., Machta, Benjamin B. and Sethna, James P. (2023). Information Geometry for Multiparameter Models: New Perspectives on the Origin of Simplicity. *Reports on Progress in Physics*, 86(3), 035901. doi:10.1088/1361-6633/aca6f8.
51. Ramadge, Peter J. and Wonham, W. Murray. (1987). Supervisory Control of a Class of Discrete Event Processes. *SIAM Journal on Control and Optimization*, 25(1), 206–230. doi:10.1137/0325013.
52. Rambachan, Ashesh and Roth, Jonathan. (2023). A More Credible Approach to Parallel Trends. *The Review of Economic Studies*, 90(5), 2555–2591. doi:10.1093/restud/rdad018.
53. Rinderle-Ma, Stefanie, Winter, Karolin and Benzin, Janik-Vasily. (2023). Predictive Compliance Monitoring in Process-Aware Informa- tion Systems: State of the Art, Functionalities, Research Directions. *Information Systems*, 115, 102210. doi:10.1016/j.is.2023.102210. Rossiter, John Anthony, Cassandras, Christos G., Hespanha, João P., Dormido, Sebastián, de la Torre, Luis, Ranade, Gireeja, Vi- sioli, Antonio, Hedengren, John, Murray, Richard M., Antsaklis, Panos, Lamnabhi-Lagarrigue, Francoise and Parisini, Thomas. (2023). Control Education for Societal-Scale Challenges: A Community Roadmap. *Annual Reviews in Control*, 55, 1–17. doi:10.1016/j.arcontrol.2023.03.007.

54. Schot, Johan and Steinmueller, W. Edward. (2018). Three Frames for Innovation Policy: R and D, Systems of Innovation and Transformative Change. *Research Policy*, 47(9), 1554–1567. doi:10.1016/j.respol.2018.08.011.
55. Shahandashti, Kimya Khakzad, Belle, Alvine B., Lethbridge, Timothy C., Odu, Oluwafemi and Sivakumar, Mithila. (2024). A PRISMA-Driven Systematic Mapping Study on System Assurance Weakeners. *Information and Software Technology*, 175, 107526. doi:10.1016/j.infsof.2024.107526.
56. Stapleton, Larry, Wang, Fei-Yue, Netto, Mariana, Jia, Qing-Shan, Visioli, Antonio and Kopacek, Peter. (2024). Shaping the Future of Advanced Automation and Control Systems for Society: Strategic Directions and Multidisciplinary Collaborations of IFAC's Social Systems Coordinating Committee. *Annual Reviews in Control*, 58, 100967. doi:10.1016/j.arcontrol.2024.100967.
57. Tabuada, Paulo. (2007). Event-Triggered Real-Time Scheduling of Stabilizing Control Tasks. *IEEE Transactions on Automatic Control*, 52(9), 1680–1685. doi:10.1109/TAC.2007.904277.
58. Teinemaa, Irene, Dumas, Marlon, La Rosa, Marcello and Maggi, Fabrizio Maria. (2019). Outcome-Oriented Predictive Process Monitoring: Review and Benchmark. *ACM Transactions on Knowledge Discovery from Data*, 13(2), 17:1–17:57. doi:10.1145/3301300.
59. Thomas, Rachel L. and Uminsky, David. (2022). Reliance on Metrics Is a Fundamental Challenge for AI. *Patterns*, 3(4), 100476. doi:10.1016/j.patter.2022.100476.
60. Thurner, Stefan, Hanel, Rudolf and Klimek, Peter. (2018). *Introduction to the Theory of Complex Systems*. Oxford University Press, Oxford. doi:10.1093/oso/9780198821930.001.0001.
61. Ulbricht, Lena and Yeung, Karen. (2022). Algorithmic Regulation: A Maturing Concept for Investigating Regulation of and Through Algorithms. *Regulation and Governance*, 16(1), 3–22. doi:10.1111/rego.12437.
62. van der Aalst, Wil M. P., Reijers, Hajo A. and Maruster, Laura. (2024). Process Mining Beyond Workflows. *Computers in Industry*, 161, 104126. doi:10.1016/j.compind.2024.104126.
63. Wada, Tatsuaki and Scarfone, Antonio M. (2024). A Hamiltonian Approach to the Gradient-Flow Equations in Information Geometry. *The European Physical Journal B*, 97, 103. doi:10.1140/epjb/s10051-024-00747-8.
64. Wager, Stefan and Athey, Susan. (2018). Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *Journal of the American Statistical Association*, 113(523), 1228–1242. doi:10.1080/01621459.2017.1319839.