



CANONICAL DISSERTATION · V1.0

The Human Time Dividend

A Measurable Theory of Human Return from Automation

Institutional burden, constitutional machine authority, and the end
of meaningless effort

Sean Manouge

The Garden Institute · Doctrine Industries · Nilsoft

25 July 2026

Publication note

This edition presents the canonical dissertation as one coordinated institutional programme.

The Garden Institute

Identifies institutional and human burden, the value at risk, and what must not be lost.

Doctrine Industries

Defines the powers, limits, evidence, rights, remedies, and exit conditions governing machine authority.

Nilsoft

Builds, operates, and measures whether automation produces a real human return.

Identify the burden through the Garden.
Bound the authority through Doctrine.
Prove the human return through Nilsoft.

Contents

Abstract	6
Authorial position and scope	8
Part I — The Burden	11
1. Work, mission, and institutional purpose	11
2. Administrative burden and system debt	14
3. Documentation, coordination, and context loss	16
4. Verification as hidden labour	19
Part II — The Machine Promise	23
5. Automation, productivity, and capacity	23
6. AI and pseudo-automation	25
7. Productivity capture	27
8. The Human Time Dividend	29
8.1 Did total burden fall?	30
8.2 Where did the capacity go?	30
8.3 What happened to the human condition?	31
Part III — The Profiling System	33
9. Behavioural data and inferred identity	33
10. Recommendation, exposure, and influence	35
11. Algorithmic management and institutional control	36
12. Biography versus profile	38

Part IV — Constitutional Automation	41
13. Provenance and institutional memory	41
14. Contestability and human authority	43
15. Bounded autonomy and operator control	45
16. Sovereignty, portability, and exit	47
Part V — The Operational Programme	51
17. Measurement and pilot method	51
18. Reference system and institutional mechanisms	53
The Garden Institute: burden and human condition	53
Doctrine Industries: authority and legitimacy	54
Nilsoft: implementation and measurement	54
Conclusion — The measure is the human return	56
Canonical propositions	58
Appendix A — Canonical definitions	61
System-generated burden	61
System debt	61
Context loss	61
Verification labour	62
Pseudo-automation	62
Productivity capture	62
Human Time Dividend	62
Constitutional automation	63
Bounded autonomy	63
Stateful operator console	63
Governed biography	63
Appendix B — Human Time Dividend scorecard	65

References and source notes	67
Research provenance	69

Abstract

Human life is finite, yet institutions routinely consume it through documentation, coordination, verification, repetition, context reconstruction, procedural uncertainty, and the correction of systems that cannot reliably remember, explain, or govern themselves. These demands are commonly treated as ordinary work or as evidence of individual inefficiency. This dissertation instead treats them as an observable systems phenomenon: **system-generated burden** produced by accumulated defects in information architecture, process design, authority allocation, interoperability, and institutional memory.

Artificial intelligence can reduce this burden. Controlled studies already show that generative systems can improve speed and quality on bounded knowledge-work tasks, while evidence from healthcare demonstrates that documentation and administrative demands can consume substantial portions of professional time.^{1,2,6} But automation does not determine its own social outcome. It can return time to people, increase service quality, or reduce cognitive load; it can also relocate work into verification, rework, exception handling, surveillance, and responsibility without authority. Productivity gains may be captured as profit, throughput, staffing reduction, managerial control, or public benefit. Their destination is institutional, not technical.^{7,8,9,10}

This dissertation introduces the **Human Time Dividend**: the measurable human return produced when technological or institutional improvement reduces net system-generated burden and

allocates the resulting capacity to reduced working time, reduced intensity, improved compensation, improved service, greater autonomy, understanding, or transition security. It is supported by three coordinated contributions:

1. **The Garden Institute** supplies an institutional and human theory of burden: what is taken from people, what becomes invisible, and what must not be lost.
2. **Doctrine Industries** supplies a constitutional architecture for machine authority: explicit powers, limits, evidence requirements, rights, remedies, and exit conditions.
3. **Nilsoft** supplies the measurable practice of human return: burden instrumentation, governed workflows, stateful operator control, and dividend accounting.

The central claim is narrow enough to test and broad enough to matter:

AI-enabled automation should be judged not by task speed, output volume, or labour removal alone, but by whether it produces a net reduction in human burden, preserves legitimate human authority, and returns a measurable share of the resulting capacity to people.

Authorial position and scope

I did not arrive at this argument by beginning with artificial intelligence. I arrived at it by repeatedly becoming the person who carried continuity that systems failed to preserve.

I encountered information that existed but could not be retrieved when needed; decisions retained without the reasoning that produced them; projects distributed across files, applications, messages, versions, devices, and individual memory; and institutions that could record a classification without preserving the human context to which it referred. The visible diagnosis was usually personal: organise more carefully, repeat the explanation, complete the form, reconcile the record, remember where the real version lives. The deeper problem was architectural.

Becoming capable of compensating for these failures did not make the burden legitimate. It made the burden visible.

Artificial intelligence changed the scale at which I could respond. I used it to reconstruct chronology across personal and project archives, compare versions, recover lineages, surface recurring concepts, separate documented evidence from remembered interpretation, and translate distributed records into inspectable structures. The objective was not to manufacture a digital identity or a definitive account of a life. It was to recover relationships hidden by fragmentation.

That experience exposed two possible uses of intelligence. One use turns human activity into profiles so that institutions and platforms can predict, rank, target, or govern more effectively. The other returns fragmented information to the person or com-

munity as continuity, comprehension, and agency. The same broad computational capability can support either purpose. The difference lies in ownership, authority, architecture, consent, evidence, and beneficiary.

This dissertation is therefore not a claim that all work is meaningless, that administration should disappear, or that AI will automatically liberate humanity. It makes five more defensible claims:

1. System-generated burden exists and can be measured.
2. AI can reduce, relocate, or intensify that burden.
3. Productivity gains have no predetermined human destination.
4. Machine-mediated authority requires constitutional limits and practical remedies.
5. Human return from automation can be evaluated through a multidimensional Human Time Dividend.

The work is normative, empirical, and architectural. Normative claims state what institutions should protect. Empirical claims are supported by research. Architectural claims specify how those commitments can become operational systems. These categories are related, but they are not interchangeable.

Part I — The Burden

THE GARDEN INSTITUTE

1. Work, mission, and institutional purpose

The word *work* conceals several different kinds of activity. It may describe effort through which people develop mastery, care for others, create culture, maintain infrastructure, or exercise responsibility. It may also describe labour performed under an employment contract, procedural activity required by an institution, or compensatory effort created by defective systems.

This distinction matters because effort is not inherently a problem. Some effort is constitutive of meaning. Hannah Arendt distinguished labour associated with biological maintenance, durable work that constructs a human world, and action through which people appear to one another politically.²⁸ Harry Braverman later examined how managerial control and deskilling altered the relationship between labour and knowledge.²⁹ David

Graeber’s account of “bullshit jobs,” although contested and difficult to operationalise, captured a widespread intuition that some paid roles are experienced as purposeless even when economically sustained.³⁰

The present framework does not attempt to decide from outside whether an occupation is meaningful. It begins at the workflow level and asks what relationship a task has to the declared purpose of the institution.

Human effort can be divided into four functional classes:

- **Mission work:** directly advances the institution’s purpose.
- **Enabling work:** provides necessary coordination, maintenance, logistics, or infrastructure.
- **Legitimate assurance:** protects rights, safety, quality, continuity, accountability, or evidence.
- **System-compensating work:** exists because information, process, authority, or technology is defective.

Administration is not automatically burden. A clinical record can protect patient continuity; an audit trail can expose abuse; a review gate can prevent irreversible harm. The relevant question is not whether an activity appears bureaucratic. It is whether the activity creates proportionate mission value or compensates for avoidable system failure.

Work becomes system-generated burden when people are repeatedly required to compensate for defects in institutional coordination, information, authority, or technical design without proportionate mission value.

This makes burden mission-relative rather than merely subjective. A frustrating task may still be necessary. A painless task may still be wasteful. The test is functional: what human or institutional good would be lost if the task disappeared, and could that good be preserved through a better design?

The principle that follows is deliberately demanding:

Work must justify its claim on a human life.

That is not a rejection of work. It is a refusal to treat inherited procedures as morally self-justifying.

2. Administrative burden and system debt

Administrative burden is now an established field of public-administration research. Moynihan, Herd, and Harvey define it through learning costs, compliance costs, and psychological costs incurred when people interact with the state.³ Later work shows that these burdens are not evenly distributed: they can restrict access to rights and services most severely for people with fewer resources, less time, or greater need.⁴ Administrative burden is therefore not merely inconvenience. It can become a mechanism of inequality.

The concept can be extended carefully beyond citizen–state interaction. Institutions impose burden internally and externally whenever people must discover requirements, satisfy procedural demands, manage uncertainty, or absorb the psychological cost of opaque systems. OECD work on administrative simplification similarly treats reporting, permits, forms, and notification obligations as measurable costs and uses methods such as the Standard Cost Model to estimate time and expense.⁵

This dissertation introduces **system debt** as the structural origin of recurring burden:

System debt is the accumulated obligation imposed on people by unresolved defects in information architecture, process design, authority allocation, interoperability, and institutional memory.

Technical debt is one subset of this larger category. A brittle application may create manual workarounds, but system debt can also originate in contradictory policies, unowned processes, duplicate authorities, inaccessible archives, or incentives that reward documentation without use.

The cycle is predictable:

```
Structural defect
→ coordination or information failure
→ compensatory human action
→ repeated burden
→ normalisation
→ institutional dependence on hidden labour
```

Once the compensatory labour becomes routine, the institution may cease to perceive the underlying defect. The workaround becomes the workflow. People who remember exceptions, reconcile records, translate terminology, and chase approvals are treated as performing ordinary duties even when their labour exists only because the system has failed to encode its own operating conditions.

Healthcare supplies a strong empirical case. Arndt and colleagues combined electronic health-record logs with time-motion observation and found that primary-care physicians spent nearly two hours on EHR and desk work for every hour of direct clinical face time; the reported workday included substantial clerical, documentation, and inbox activity.¹ A 2024 AHRQ technical brief identified 135 studies and eleven distinct categories for measur-

ing documentation burden, including EHR time, inbox work, after-hours work, administrative tasks, workflow fragmentation, and usability.²

The significance is not that documentation is unnecessary. Clinical records are essential. The evidence shows that the architecture surrounding those records can impose measurable burden and that no single measure captures it completely. System debt is often distributed across time, cognitive load, after-hours spillover, duplicated action, and lost attention.

Meaningless effort is therefore not best defined as work someone dislikes. It is better defined as **mission-external or disproportionate human effort generated by a system's unresolved debt**.

3. Documentation, coordination, and context loss

Institutions can possess enormous quantities of information while remaining unable to use it coherently. A fact may have been recorded without being retrievable; retrieved without being interpretable; interpreted without the necessary authority; or acted upon without preserving the reason for action.

The relevant chain is:

Capture → storage → retrieval → interpretation → authority → action

A failure at any stage can force people to reconstruct context that the institution has already encountered.

Documentation serves several legitimate functions. It may be operational, evidentiary, coordinative, memorial, protective, or regulatory. Burden emerges when documentation is collected without becoming usable institutional memory. A form may satisfy a reporting obligation while failing to support future action. A decision may be logged without its assumptions. A case may be handed between teams with its status but not its reasoning. A file may exist in several versions without a canonical state.

Context loss converts prior work into future work. The institution repeatedly pays to rediscover what it once knew, and the payment is usually made in human attention.

This is especially visible in fragmented digital estates. Each application may function correctly within its own boundary while the person remains responsible for remembering the relationships between them. The human becomes the integration layer: translating names, reconciling states, identifying the current record, and carrying exceptions that no schema represents.

The problem is not solved by producing more documentation. Additional records can increase future burden when they are duplicated, unsearchable, unversioned, poorly scoped, or detached

from decisions. Conversely, high-quality documentation can reduce burden by preserving provenance, rationale, ownership, and status.

A useful distinction is therefore:

```
recorded information ≠ available context  
available context ≠ understood context  
understood context ≠ authorised action
```

The purpose of institutional memory is not to preserve everything indefinitely. It is to preserve enough structured context that people do not repeatedly reconstruct the same operating reality.

This is where the Garden Institute's theory of burden meets Nilsoft's technical programme. The Garden asks what continuity has been lost and what human cost follows. Nilsoft asks how state, evidence, decisions, and ownership can be made durable and inspectable. Doctrine asks which records may legitimately be created, retained, combined, disclosed, or used to act upon a person.

Context preservation is therefore not a purely technical convenience. It is simultaneously a burden problem, a memory problem, and an authority problem.

4. Verification as hidden labour

Automation is often evaluated by the time required to produce an output. That measure can be profoundly misleading.

A generated document may be faster to draft but slower to verify. An automated classification may reduce data entry while creating an exception queue. A recommendation may accelerate a decision while transferring responsibility to a reviewer who lacks access to the evidence or logic behind it. A system may save time for one department by shifting correction work to customers, citizens, contractors, or after-hours staff.

This is **verification labour**: the work required to establish that an automated or institutional output can be trusted, accepted, corrected, or acted upon.

Verification includes:

- factual checking;
- policy and eligibility checking;
- source validation;
- correction and rework;
- exception review;
- duplicate-execution detection;
- approval;
- incident response;
- audit reconstruction;
- confirmation that a state change actually occurred.

Research on generative AI in knowledge work suggests that AI changes rather than simply removes cognitive work. A 2025 CHI study of 319 knowledge workers found that participants described a shift from information gathering toward verification, from problem solving toward response integration, and from task execution toward task stewardship. Because the study relied on self-report, it should not be read as direct evidence of cognitive decline; it does support the narrower claim that verification and stewardship become central components of AI-assisted work.¹⁹

The EU AI Act explicitly recognises automation bias. For high-risk systems, Article 14 requires human-oversight measures that help reviewers understand limitations, interpret outputs, avoid over-reliance, override or reverse outputs, and stop the system safely.²⁰ This is significant because “human in the loop” is not inherently protective. A reviewer without time, evidence, independence, competence, or actual authority may provide ceremonial oversight while retaining legal or moral responsibility.

The correct unit of evaluation is total burden:

```
Total Human Burden
=
execution time
+ coordination time
+ documentation time
+ verification time
+ rework time
+ exception-handling time
+ recovery time
+ after-hours spillover
```

```
Net Burden Reduction
=
baseline total human burden
- post-intervention total human burden
```

Cognitive intensity, autonomy, frustration, and perceived meaningfulness should be measured separately rather than converted into fictional equivalent hours. NASA-TLX, for example, provides a multidimensional instrument for subjective workload across mental, physical, temporal, performance, effort, and frustration dimensions.²⁷

The central warning is simple:

Automation that removes execution but expands verification may be pseudo-automation rather than liberation.

Part II — The Machine Promise

NILSOFT

5. Automation, productivity, and capacity

Artificial intelligence can produce genuine productivity gains. In a preregistered experiment involving 453 professionals completing bounded writing tasks, Noy and Zhang found that access to ChatGPT reduced average completion time by 40 percent and increased assessed output quality by 18 percent.⁶ The result is important, but its scope is equally important: it concerns particular writing tasks under experimental conditions, not all knowledge work or whole organisations.

OECD's 2024 synthesis likewise concludes that AI has significant potential to improve worker and firm productivity, innovation, and growth, while emphasising uncertainty, uneven diffusion, market concentration, inequality, discrimination, and job displacement.⁷

Automation can change several properties of production:

- speed;
- labour input;
- scale;
- consistency;
- marginal cost;
- error distribution;
- service capacity;
- the range of tasks an individual can perform.

It cannot independently determine:

- wages;
- working hours;
- staffing levels;
- workload intensity;
- worker autonomy;
- access to the resulting service;
- ownership of productive infrastructure;
- distribution of gains.

The causal sequence is therefore:

```
Technical capability
→ task transformation
→ workflow redesign
→ organisational response
→ distributional allocation
→ human outcome
```

Only the first stages are primarily technical. An organisation can respond to a time-saving system by reducing hours, increasing throughput, reducing staff, raising quality, expanding access, or intensifying targets. The same model capability can support each outcome.

This distinction prevents a common category error. Productivity is not identical to progress. Productivity describes a relationship between inputs and outputs. Progress requires a judgment about human and institutional outcomes.

The machine promise should therefore be stated precisely: AI can create **productive capacity**. It does not guarantee a Human Time Dividend.

6. AI and pseudo-automation

Many AI systems are marketed through the language of autonomy even when they remain materially dependent on hidden human work.

Pseudo-automation is an intervention represented as autonomous or labour-saving that materially depends on hidden, displaced, intensified, or newly created human labour.

This labour may include data preparation, prompt construction, exception resolution, content moderation, source checking, correction, approval, customer recovery, or responsibility-bearing after the system fails. It may be shifted across organisational boundaries or removed from formal accounting altogether.

Not all verification makes automation false. High-consequence systems require legitimate assurance. The issue is whether the assurance burden is proportionate, visible, adequately supported, and included in the productivity claim.

A system is presumptively pseudo-automated when:

- outputs require systematic manual repair;
- users repeatedly reintroduce missing context;
- responsibility remains human while evidence and authority remain opaque;
- failure handling is excluded from cost calculations;
- labour is shifted to another worker, customer, citizen, or family member;
- output volume grows faster than usable capacity;
- apparent savings are absorbed by after-hours work;

- the organisation becomes unable to operate when the system is unavailable.

Pseudo-automation is especially likely in output-only interfaces. A chatbot can produce a plausible answer without preserving durable workflow state, policy version, evidence chain, approval status, or resulting action. The human must then reconstruct those conditions around the answer.

This is why the correct architecture is not “a model connected to tools.” It is a governed workflow in which models operate inside explicit state, permissions, evidence, policy, and consequence boundaries.

The test is not whether a machine performed a step. It is whether total human burden fell without unacceptable loss of quality, authority, resilience, or meaning.

7. Productivity capture

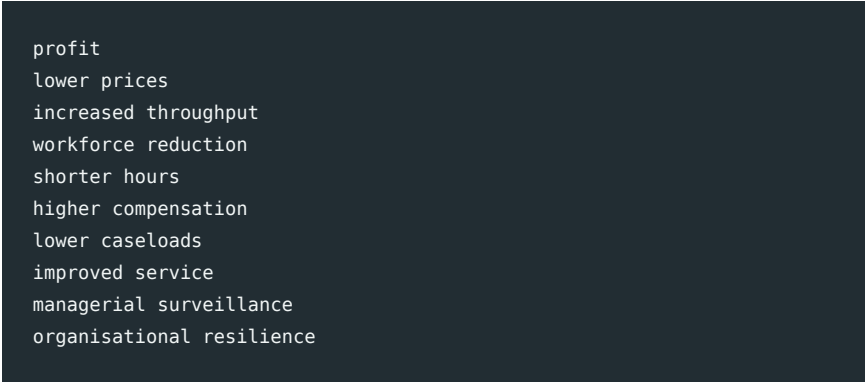
Productive capacity has no predetermined destination.

Research on prior automation demonstrates that technical progress can coexist with negative local labour-market effects. Acemoglu and Restrepo found that greater exposure to industrial robots in US commuting zones reduced employment-to-population ratios and wages.⁸ This result does not prove that

every form of AI will produce the same effects; it establishes that labour-saving technology does not automatically distribute benefits to workers.

Recent OECD cross-country evidence similarly finds an inverse relationship between productivity and labour share at firm and industry levels, with reallocation toward persistently low-labour-share “superstar” firms reducing aggregate labour shares.⁹ ILO and OECD reporting states that the global labour-income share fell by 1.6 percentage points between 2004 and 2024, attributing the trend to several structural factors including automation, globalisation, and weakened worker bargaining power.¹⁰

Created capacity can be allocated to:



- profit
- lower prices
- increased throughput
- workforce reduction
- shorter hours
- higher compensation
- lower caseloads
- improved service
- managerial surveillance
- organisational resilience

These allocations are not equivalent. Some return value directly to workers, some to service recipients, some to owners, and some to institutional robustness. A defensible framework should not assume that every gain must become leisure, nor that profit or re-

silence are illegitimate. It should require that the destination is visible and that claims of human benefit correspond to actual allocation.

This is the distinction between **productivity creation** and **productivity capture**.

A productivity gain is newly available capacity. Its human meaning depends on who captures it, in what form, and under what authority.

8. The Human Time Dividend

The Human Time Dividend is the dissertation's central construct.

The Human Time Dividend is the measurable human return generated when technological or institutional improvement reduces net system-generated burden and the resulting capacity is allocated to reduced working time, reduced intensity, improved compensation, improved service, greater autonomy, understanding, or transition security.

It is deliberately multidimensional. Hours, wages, service quality, autonomy, and wellbeing cannot be combined into one scientifically neutral number without arbitrary weights.

The framework therefore separates three questions.

8.1 Did total burden fall?

```
Net Time Dividend
=
baseline human hours
- post-intervention human hours
- displaced unpaid hours
```

Reported as hours per worker, hours per case, hours per service user, and total institutional hours.

8.2 Where did the capacity go?

The organisation reports the distribution of recovered capacity across:

- reduced working time;
- reduced workload or caseload;
- improved service;
- compensation or benefits;
- training and transition;
- organisational resilience;
- increased throughput;

- capital capture;
- unobserved or unallocated gain.

8.3 What happened to the human condition?

Measure separately:

- cognitive workload;
- autonomy;
- schedule control;
- after-hours spillover;
- burnout indicators;
- perceived meaningfulness;
- service quality;
- error, correction, and appeal rates;
- capacity for meaningful participation.

A four-day-workweek study published in 2025 followed 2,896 employees across 141 organisations after work redesign and reduced hours without reduced pay. It found improvements in worker wellbeing relative to controls, while acknowledging limitations including organisational self-selection and reliance on self-reported outcomes.¹¹ The study does not establish a universal policy. It demonstrates that one plausible capture mechanism—returning capacity as time—can be implemented and evaluated.

A Human Time Dividend exists only when net burden reduction is positive and an identifiable benefit reaches workers, users, or the public. A faster system that merely raises output targets has created capacity, not necessarily a dividend.

Part III — The Profiling System

GARDEN + DOCTRINE

9. Behavioural data and inferred identity

The same computational systems that can reduce administrative burden can also make people more legible to institutions.

Modern profiling does not rely only on information a person explicitly submits. It combines behavioural traces, relationships, transactions, device signals, location, attention, and similarity to other users in order to infer preferences, risks, reliability, capacity, or future action.

The GDPR defines profiling broadly as automated processing used to evaluate personal aspects, including work performance, economic situation, health, preferences, interests, reliability,

behaviour, location, or movements.^{12,13} The definition is important because profiling is not limited to advertising. It can shape employment, finance, insurance, access, visibility, scheduling, and institutional treatment.

The data ladder is:

```
Observed event
→ recorded data
→ behavioural signal
→ statistical inference
→ institutional classification
→ consequential treatment
```

Each stage adds interpretation and uncertainty. Yet the resulting profile may become operationally more powerful than the person's own account.

The FTC's investigation of data brokers found a fundamental lack of transparency and the collection and sharing of large amounts of consumer data behind the scenes, often without consumer knowledge.¹² This creates an asymmetry: organisations can combine fragments that the person cannot see, audit, or correct.

The danger is not that every profile is inaccurate. It is that a probabilistic representation can acquire authority without corresponding rights of context, correction, contestation, or exit.

The profile begins as an operational abstraction and can end as an institutional definition.

10. Recommendation, exposure, and influence

Recommendation systems do not need to determine belief in order to exercise power. They allocate exposure: what appears, in what order, how often, with which social signals, and beside which alternatives.

The strongest defensible claim is that recommendation affects attention, salience, emotional state, perceived norms, discovery, and reinforcement. Evidence for direct ideological transformation is more conditional.

Allcott and colleagues' 2018-election experiment found that Facebook deactivation increased subjective wellbeing, reduced factual news knowledge, and reduced measured political polarisation.¹⁴ A much larger 2020-election experiment found effects on online participation and some knowledge measures but effects on affective and issue polarisation that were close to zero.¹⁵ Naturalistic YouTube experiments involving nearly 9,000 participants found that short-term filter-bubble and rabbit-hole recommendation conditions changed media-consumption choices but had limited downstream effects on political attitudes.¹⁶

These findings reject both simplistic denial and deterministic fear. Platform architecture matters, but users retain agency, effects vary by outcome and context, and short-term attitude change may be smaller than popular narratives suggest.

The feedback structure remains important:

```
behaviour
→ profile update
→ ranked exposure
→ attention and interaction
→ revised behaviour
→ stronger profile
```

The system does not merely discover a stable person. It participates in constructing the informational environment through which the person interprets themselves and others.

Personalisation can improve relevance and accessibility. It can also become containment when past behaviour narrows future exposure and reduces the chance of encountering difficult, unfamiliar, or identity-disrupting material.

A person must retain the practical right to become more than their prediction.

11. Algorithmic management and institutional control

Profiling becomes institutional power when behavioural data is used to assign work, set targets, schedule time, evaluate performance, determine access, trigger intervention, or justify discipline.

OECD defines algorithmic management as the use of software, potentially including AI, to automate tasks traditionally performed by managers. Its 2025 cross-country employer survey documents the use of systems that instruct, monitor, and evaluate workers and highlights benefits alongside transparency, accountability, and psychosocial concerns.¹⁷

Algorithmic management can reduce coordination burden. Dynamic scheduling may match resources more efficiently; automated quality checks may identify defects; workload routing may reduce queues. The same infrastructure can intensify surveillance, remove discretion, obscure the basis of evaluation, and convert measurable activity into a proxy for valuable work.

For each system, the constitutional questions are:

- What data is collected?
- What inference is produced?
- What authority is granted to the inference?
- Which decision is affected?
- Who bears the error?
- Who benefits from the system?
- Can the person inspect, correct, or contest the representation?
- Can the action be reversed?

The problem is not that a machine “makes a decision” in isolation. Institutions define the categories, objectives, thresholds, and consequences. Machine authority is delegated institutional authority.

The central question is who may use which representation of a person, for what purpose, with what authority, and subject to which remedy.

12. Biography versus profile

A person can be extensively recorded while remaining almost entirely unknown.

Profiles and biographies are both selective. Neither is a complete human being. Their difference lies in purpose, authorship, temporality, and authority.

Operational profile	Governed biography
Optimised for prediction or allocation	Oriented toward contextual understanding
Institution-defined purpose	Subject-centred purpose
Compresses uncertainty	Preserves uncertainty and contradiction
Current-state emphasis	Development through time
Often opaque	Provenance-visible
Used to act upon the subject	Used primarily to return understanding to the subject

Operational profile	Governed biography
Difficult to correct	Revision and contestation are structural rights

James C. Scott showed how states simplify complex social reality into legible categories that enable administration but can erase local knowledge and variation.³¹ Shoshana Zuboff later described an economic order organised around the extraction of behavioural data for prediction and intervention.³² These frameworks differ, but both illuminate the danger of mistaking institutional legibility for human truth.

The alternative to an extractive profile is not a perfect digital twin. A complete model of a person is neither achievable nor desirable. The alternative is a governed record that preserves chronology, evidence, uncertainty, multiple perspectives, private boundaries, and the person's right to exceed every representation.

This is the Garden Institute's core domain. It asks how intelligence can help people recover continuity without turning biography into a more invasive profile.

Personal data should not merely make the person more legible to systems. Properly governed, it should help make the person's life more legible to themselves.

Part IV — Constitutional Automation

DOCTRINE INDUSTRIES

13. Provenance and institutional memory

Accountability requires more than a large volume of logs. It requires meaningful relationships between source, transformation, inference, authority, action, and outcome.

Four forms of provenance are required:

1. **Data provenance:** where inputs originated and how they changed.
2. **Decision provenance:** how an interpretation or proposal was formed.

3. **Authority provenance:** who authorised the relevant use or action.
4. **Outcome provenance:** what changed and what occurred afterward.

The minimum chain is:

```
Source
→ transformation
→ retrieval
→ inference
→ proposal
→ policy evaluation
→ authorisation
→ execution
→ outcome
→ review
```

The W3C PROV family provides a standard conceptual model for representing entities, activities, agents, generation, derivation, and attribution across heterogeneous systems.²² Model Cards and Datasheets for Datasets similarly formalise documentation of intended use, performance, collection, composition, limitations, and recommended contexts.^{23,24}

These artefacts do not guarantee accountability. Documentation can become another burden if it is disconnected from actual workflows or produced only for compliance. Their value lies in making consequential relationships inspectable and durable.

Institutional memory should answer:

- What was known?
- Which version was used?
- Which uncertainty remained?
- What policy applied?
- Who had authority?
- What action occurred?
- What changed afterward?
- How was the result corrected or appealed?

Without that structure, every incident becomes an archaeological exercise and every new operator inherits reconstruction work.

14. Contestability and human authority

Transparency is not accountability.

A person may be told that an algorithm was used while lacking access to the relevant evidence, a competent reviewer, or any authority capable of changing the result. Disclosure without remedy leaves the underlying power intact.

Contestability is the practical ability of an affected person or authorised operator to identify that automation was involved, inspect the relevant basis, suspend or challenge its effect, correct material errors, and obtain review from an authority capable of changing the outcome.

The GDPR provides rights to data portability and protections against certain solely automated decisions with legal or similarly significant effects; where exceptions apply, safeguards include human intervention, the opportunity to express a view, and the right to contest the decision.¹⁸ The EU AI Act requires high-risk systems to provide logging, sufficient transparency for deployers, and effective human oversight.²⁰

Meaningful contestability requires:

- notice;
- access to relevant records;
- explanation appropriate to the decision;
- correction of underlying data;
- suspension where harm is plausible;
- review by an independent or distinct authority;
- an enforceable remedy;
- a resolution record.

Human involvement is not sufficient. The reviewer must have time, evidence, competence, independence, and power.

Information without remedy is disclosure, not accountability.

Doctrine Industries treats these conditions as constitutional, not optional interface features. They define the legitimate allocation of power within the system.

15. Bounded autonomy and operator control

Autonomy should be allocated according to consequence, reversibility, uncertainty, evidence quality, and institutional legitimacy.

A practical authority ladder is:

Class	Permitted system action
A0	Observe and record
A1	Retrieve and summarise
A2	Analyse and recommend
A3	Prepare a reversible action
A4	Execute a low-consequence action within explicit policy
A5	Execute only after explicit approval

Class	Permitted system action
A6	Human-only decision; autonomous execution prohibited

Classification should consider:

- consequence severity;
- reversibility;
- legal significance;
- affected population;
- privacy sensitivity;
- evidence quality;
- model uncertainty;
- time sensitivity;
- reviewability;
- cumulative harm.

The NIST AI Risk Management Framework organises AI risk work through Govern, Map, Measure, and Manage, with governance operating across the lifecycle.²¹ The EU AI Act similarly requires risk management, logging, transparency, and human oversight for high-risk systems.²⁰

The interface implication is a **stateful operator console**, not merely a conversational assistant. The console must expose current workflow state, active authority class, evidence, proposed changes, expected consequences, policy results, pending approvals, rollback status, and review history.

Current orchestration tooling demonstrates that durable execution, persistence, pauses, and human intervention are technically feasible. LangGraph, for example, documents persistence and human-in-the-loop interruption patterns; OpenTelemetry and OpenLineage provide vendor-neutral primitives for traces and execution lineage.^{25,26,33,34} These tools are examples, not constitutional guarantees. Governance must remain independent of any vendor implementation.

The machine may extend the operator. It must not conceal the operating conditions from them.

16. Sovereignty, portability, and exit

Data sovereignty is often reduced to storage location. That is insufficient.

Sovereignty is the practical capacity to govern data, models, records, and operations without becoming irreversibly dependent on an external authority.

It includes:

- storage location;
- encryption-key control;

- identity and access authority;
- portability;
- open formats;
- interoperability;
- retention and deletion;
- recoverability;
- model substitutability;
- audit access;
- jurisdiction;
- operational exit.

The GDPR's portability right supports transfer of certain personal data in structured, commonly used, machine-readable form.¹⁸

The local-first software literature argues that software can combine collaboration with user ownership, offline operation, privacy, longevity, and control by treating local data as primary rather than merely cached.²⁵

Local-first is not inherently sovereign or safe. It may transfer backup, maintenance, consistency, and security burdens to users. Cloud systems may provide resilience that individuals cannot reproduce. The relevant question is not local versus cloud as an ideology. It is whether the system preserves meaningful authority, recoverability, inspectability, and exit.

Vendor dependence is a form of future system debt. A system that saves administrative time today while making histories non-portable, workflows non-reconstructable, or institutions unable to operate without one provider may create a negative long-term dividend.

Sovereignty must therefore be included in burden accounting.

Part V — The Operational Programme

NILSOFT

17. Measurement and pilot method

The Human Time Dividend is useful only if it can fail.

A pilot should specify four units of analysis:

1. the workflow;
2. the worker or operator;
3. the service recipient;
4. the institution.

The minimum design is:

```
workflow selection
→ baseline mapping
→ burden classification
→ preregistered hypotheses
→ intervention
→ comparison condition
→ runtime trace collection
→ follow-up
→ distribution analysis
→ governance review
```

The strongest candidate workflows are bounded, repetitive, measurable, reversible, and low enough in consequence to permit controlled experimentation. Examples include document intake, records reconciliation, appointment administration, internal service requests, case handoffs, knowledge retrieval, and compliance-evidence preparation.

Measurement should combine:

- time-use diaries;
- workflow event logs;
- process mining;
- direct observation where practical;
- NASA-TLX or another validated workload instrument;²⁷
- worker interviews;
- user-burden measures;
- error and rework data;
- incident and appeal analysis;
- AI traces and lineage records.

Healthcare research demonstrates the value and limits of this mixed approach. The AHRQ review found multiple measure categories and heavy reliance on EHR logs, while noting incomplete validity evidence and the need for broader perspectives.² OECD administrative-simplification methods demonstrate how compliance obligations can be decomposed into measurable time and cost.⁵

Measurement must not become surveillance. Required safeguards include purpose limitation, data minimisation, aggregation by default, worker consultation, defined retention, separation from disciplinary systems, and accessible results.

The pilot fails when task time falls but total burden, error, after-hours effort, surveillance, or dependency rises.

18. Reference system and institutional mechanisms

The reference system operationalises the three institutional domains.

The Garden Institute: burden and human condition

The Garden defines what is being lost and what must be preserved:

- human time;
- meaningful effort;

- capability;
- continuity;
- agency;
- biographical integrity;
- the right to remain more than a profile.

Doctrine Industries: authority and legitimacy

Doctrine defines:

- permitted and prohibited actions;
- authority classes;
- evidence requirements;
- approval gates;
- contestability;
- correction and appeal;
- retention;
- portability;
- exit and remedy.

Nilsoft: implementation and measurement

Nilsoft provides:

- authenticated roles;
- durable workflow state;
- bounded context assembly;
- evidence retrieval;

- AI proposal generation;
- policy evaluation;
- authority classification;
- reversible execution;
- immutable trace records;
- burden telemetry;
- dividend reporting;
- model substitution;
- graceful failure.

The reference flow is:

```
Request
→ context assembly
→ deterministic preparation
→ AI proposal
→ policy evaluation
→ authority classification
→ approval or refusal
→ bounded execution
→ outcome verification
→ provenance record
→ burden measurement
→ dividend allocation report
```

Technical architecture can make responsible allocation visible. It cannot make allocation durable by itself. Institutional mechanisms may include reduced-hours agreements, workload protections, gain-sharing, worker consultation, automation-impact as-

assessments, transition funding, algorithmic-management disclosure, public-procurement conditions, portability rights, and service-quality obligations.

The evidence for particular interventions varies. Shorter-hours trials are promising but not universally generalisable.¹¹ Labour-market evidence strongly supports the need to treat distribution as institutional rather than automatic.^{8,9,10} The dissertation therefore does not prescribe one universal political programme. It requires every automation initiative to state how capacity will be allocated and to report the result.

Conclusion — The measure is the human return

Modern institutions accumulate debt in their structures. People absorb that debt as documentation, coordination, reconstruction, waiting, verification, exception handling, and after-hours recovery.

Artificial intelligence can reduce those burdens. It can also create a faster and more opaque institution: more output to verify, more profiles to contest, more surveillance, more concentrated authority, and more dependence on systems whose internal state people cannot inspect.

The decisive question is not whether the machine is intelligent. It is whether the institution is designed to use intelligence legitimately.

The complete argument is:

```
Institutions accumulate system debt
→ people absorb it as hidden labour
→ AI can remove, relocate, or intensify that burden
→ productivity alone does not determine the human outcome
→ profiling and algorithmic management reveal the extractive alternative
→ automation therefore requires constitutional powers, limits, rights, records, and remedies
→ the Human Time Dividend measures net human return
→ pilots determine whether the doctrine survives implementation
```

The Garden Institute identifies the burden and the human value at risk.

Doctrine Industries bounds machine and institutional authority.

Nilsoft builds and measures the intervention.

Together they establish one programme:

Identify the burden through the Garden.
Bound the authority through Doctrine.
Prove the human return through Nilsoft.

The objective is not a humanity without effort. It is a humanity no longer forced to waste itself maintaining systems that could responsibly carry their own machinery.

The machine should carry the machinery.

The human should carry the meaning.

The measure is the human return.

Canonical propositions

1. Human life is finite.
2. Work must justify its claim on a human life.
3. System-generated burden is observable and measurable.
4. Meaningless effort is often system debt transferred onto people.
5. Automation can remove, relocate, or intensify burden.

6. **Productivity creates capacity; institutions determine its destination.**
 7. **A Human Time Dividend exists only when measurable benefit reaches people.**
 8. **Capability does not confer legitimacy.**
 9. **Information without remedy is disclosure, not accountability.**
 10. **Human oversight requires evidence, time, competence, independence, and authority.**
 11. **The profile is not the person.**
 12. **A person must retain the right to become more than their prediction.**
 13. **Data should return understanding to the subject, not merely power to the system.**
 14. **Sovereignty requires portability, recoverability, inspectability, and exit.**
 15. **The objective is not maximum automation. It is measurable human return.**
-

Appendix A — Canonical definitions

System-generated burden

Human time, cognitive load, delay, uncertainty, or loss of agency produced by the structures through which an institution coordinates, records, verifies, classifies, and governs activity.

System debt

The accumulated obligation imposed on people by unresolved defects in information architecture, process design, authority allocation, interoperability, and institutional memory.

Context loss

Degradation of usable meaning across capture, storage, retrieval, interpretation, authority, and action.

Verification labour

Human effort required to establish that an automated or institutional output is trustworthy, compliant, accepted, corrected, or safe to act upon.

Pseudo-automation

An intervention represented as autonomous or labour-saving that materially depends on hidden, displaced, intensified, or newly created human work.

Productivity capture

The allocation of capacity created by technological or institutional improvement among workers, service users, institutions, owners, and other actors.

Human Time Dividend

The measurable human return generated when improvement reduces net system-generated burden and allocates the resulting capacity to reduced time, reduced intensity, improved compensation, improved service, autonomy, understanding, or transition security.

Constitutional automation

The design and governance of automated systems through explicit powers, limits, rights, evidence requirements, authority chains, records, remedies, and exit conditions.

Bounded autonomy

Delegated machine authority restricted according to consequence, reversibility, evidence, uncertainty, policy, and legitimate human control.

Stateful operator console

A persistent interface through which authorised humans inspect system state, evidence, policy, authority, proposed actions, consequences, approvals, execution, and history.

Governed biography

A subject-centred, provenance-visible record that preserves context, chronology, uncertainty, multiple perspectives, privacy boundaries, and the right of revision without claiming to replace the person.

Appendix B — Human Time Dividend scorecard

Dimension	Indicator	Example measure
Burden	Administrative time	Minutes per case or worker per week
Burden	Verification time	Minutes spent checking AI outputs
Burden	Rework	Percentage of cases corrected or repeated
Burden	After-hours spillover	Work performed outside scheduled hours
Human condition	Cognitive workload	NASA-TLX or validated alternative
Human condition	Autonomy	Surveyed control over sequence, method, and pace
Human condition	Meaning	Perceived relationship between task and mission
Quality	Error rate	Substantive errors per case

Dimension	Indicator	Example measure
Quality	Exception rate	Cases requiring manual escalation
Governance	Provenance coverage	Percentage of consequential actions with complete chain
Governance	Contestability	Percentage with practical correction and appeal route
Governance	Override effectiveness	Reviews capable of changing or reversing outcome
Distribution	Reduced working time	Hours formally returned
Distribution	Workload relief	Caseload or queue reduction
Distribution	Compensation	Share allocated to wages or benefits
Distribution	Service return	Reduced waiting, backlog, or abandonment
Distribution	Capital capture	Share retained as profit or asset value
Resilience	Exit capability	Time and cost required to operate without the system

References and source notes

1. Arndt, B. G., Beasley, J. W., Watkinson, M. D., Temte, J. L., Tuan, W.-J., Sinsky, C. A., & Gilchrist, V. J. (2017). "Tethered to the EHR: Primary Care Physician Workload Assessment Using EHR Event Log Data and Time-Motion Observations." *Annals of Family Medicine*, 15(5), 419–426. <https://doi.org/10.1370/afm.2121>
2. Wang, Z., West, C. P., Vaa Stelling, B. E., et al. (2024). *Measuring Documentation Burden in Healthcare*. AHRQ Technical Brief No. 47. Agency for Healthcare Research and Quality. <https://www.ncbi.nlm.nih.gov/books/NBK608551/> ; peer-reviewed version: Murad, M. H., et al. (2024). *Journal of General Internal Medicine*, 39, 2837–2848. <https://doi.org/10.1007/s11606-024-08956-8>
3. Moynihan, D. P., Herd, P., & Harvey, H. (2015). "Administrative Burden: Learning, Psychological, and Compliance Costs in Citizen-State Interactions." *Journal of Public Administration Research and Theory*, 25(1), 43–69. <https://doi.org/10.1093/jopart/muu009>
4. Herd, P., Hoynes, H., Michener, J., & Moynihan, D. (2023). "Administrative Burden as a Mechanism of Inequality in Policy Implementation." *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 9(5), 1–30. <https://doi.org/10.7758/RSF.2023.9.5.01> ; Herd, P., & Moynihan, D. (2025). "Administrative Burdens in the Social Safety Net." *Journal of Economic Perspectives*, 39(1), 129–150. <https://doi.org/10.1257/jep.20231394>
5. OECD. "Administrative Simplification." <https://www.oecd.org/en/topics/sub-issues/administrative-simplification.html> ; OECD. (2026). *Smart Regulations, Strong Business*, section on the Standard Cost Model and burden-reduction tools. https://www.oecd.org/en/publications/smart-regulations-strong-business_93d38770-en/
6. Noy, S., & Zhang, W. (2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
7. Filippucci, F., Gal, P., Jona-Lasinio, C., Leandro, A., & Nicoletti, G. (2024). *The Impact of Artificial Intelligence on Productivity, Distribution and Growth: Key Mechanisms, Initial Evidence and Policy Challenges*. OECD Artificial Intelligence Papers No. 15. <https://doi.org/10.1787/8d900037-en>
8. Acemoglu, D., & Restrepo, P. (2020). "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy*, 128(6), 2188–2244. <https://doi.org/10.1086/705716>
9. Cho, W. H., Manaresi, F., & Reinhard, M. (2025). *Superstars, Shooting Stars, and the Labour Share: Cross-Country Evidence*. OECD Science, Technology and Industry Working Papers No. 2025/01. <https://doi.org/10.1787/ecb32e77-en>

10. International Labour Organization & OECD. (2025). *Policy Measures to Address Inequalities and Increase the Labour Income Share*. <https://www.ilo.org/publications/policy-measures-address-inequalities-and-increase-labour-income-share>
11. Fan, W., Schor, J. B., Kelly, O., et al. (2025). "Work Time Reduction via a 4-Day Workweek Finds Improvements in Workers' Well-Being." *Nature Human Behaviour*, 9, 2153–2168. <https://doi.org/10.1038/s41562-025-02259-6>
12. Federal Trade Commission. (2014). *Data Brokers: A Call for Transparency and Accountability*. <https://www.ftc.gov/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014>
13. European Data Protection Board. "Respect Individuals' Rights: Right Not to Be Subject to a Decision Based Solely on Automated Processing." https://www.edpb.europa.eu/sme/be-compliant/respect-individuals-rights_en
14. Allcott, H., Braghieri, L., Eichmeyer, S., & Gentzkow, M. (2020). "The Welfare Effects of Social Media." *American Economic Review*, 110(3), 629–676. <https://doi.org/10.1257/aer.20190658>
15. Allcott, H., Gentzkow, M., Mason, W., et al. (2024). "The Effects of Facebook and Instagram on the 2020 Election: A Deactivation Experiment." *Proceedings of the National Academy of Sciences*, 121(21), e2321584121. <https://doi.org/10.1073/pnas.2321584121>
16. Liu, N., Hu, X. E., Savas, Y., de Benedictis-Kessner, J., Guess, A. M., & Stewart, B. M. (2025). "Short-Term Exposure to Filter-Bubble Recommendation Systems Has Limited Polarization Effects: Naturalistic Experiments on YouTube." *Proceedings of the National Academy of Sciences*, 122. <https://doi.org/10.1073/pnas.2318127122>
17. Milanez, A., Lemmens, A., & Ruggiu, C. (2025). *Algorithmic Management in the Workplace: New Evidence from an OECD Employer Survey*. OECD Artificial Intelligence Papers No. 31. <https://doi.org/10.1787/287c13c4-en>
18. European Union. (2016). Regulation (EU) 2016/679, General Data Protection Regulation, especially Articles 20 and 22 and Article 4(4). <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
19. Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects from a Survey of Knowledge Workers." *Proceedings of CHI 2025*, Article 1121. <https://doi.org/10.1145/3706598.3713778>
20. European Union. (2024). Regulation (EU) 2024/1689, Artificial Intelligence Act, especially Articles 9 and 12–15. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
21. Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>. NIST states that AI RMF 1.0 is being revised as of 2026; this dissertation cites the published 1.0 framework.
22. Lebo, T., Sahoo, S., & McGuinness, D. (eds.). (2013). *PROV-O: The PROV Ontology*. W3C Recommendation. <https://www.w3.org/TR/prov-o/>

23. Mitchell, M., Wu, S., Zaldivar, A., et al. (2019). "Model Cards for Model Reporting." *Proceedings of FAT*19*, 220–229. <https://doi.org/10.1145/3287560.3287596>
 24. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). "Datasheets for Datasets." *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
 25. Kleppmann, M., Wiggins, A., van Hardenberg, P., & McGranaghan, M. (2019). "Local-First Software: You Own Your Data, in Spite of the Cloud." *Proceedings of Onward! 2019*, 154–178. <https://doi.org/10.1145/3359591.3359737>
 26. OpenTelemetry. *OpenTelemetry Specification*. <https://opentelemetry.io/docs/specs/>; OpenLineage. *Open Standard for Lineage Metadata Collection*. <https://openlineage.io/docs/>
 27. Cao, A., Chintamani, K. K., Pandya, A. K., & Ellis, R. D. (2009). "NASA TLX: Software for Assessing Subjective Mental Workload." *Behavior Research Methods*, 41, 113–117. <https://doi.org/10.3758/BRM.41.1.113>
 28. Arendt, H. (1958). *The Human Condition*. University of Chicago Press.
 29. Braverman, H. (1974). *Labor and Monopoly Capital: The Degradation of Work in the Twentieth Century*. Monthly Review Press.
 30. Graeber, D. (2018). *Bullshit Jobs: A Theory*. Simon & Schuster.
 31. Scott, J. C. (1998). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale University Press.
 32. Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs.
 33. LangChain. *LangGraph Overview, Persistence, and Human-in-the-Loop Documentation*. <https://docs.langchain.com/oss/python/langgraph/overview>; <https://docs.langchain.com/oss/python/langgraph/persistence>; <https://docs.langchain.com/oss/python/langchain/human-in-the-loop>
 34. OpenLineage. *About OpenLineage*. <https://openlineage.io/docs/>; OpenTelemetry. *Tracing API and Specification*. <https://opentelemetry.io/docs/specs/otel/trace/>
-

Research provenance

This canonical version was developed from the prior dissertation drafts, the Canonical Trifecta Framework v1.0, and the July 2026 deep-research evidence pass. Empirical claims were narrowed

where the literature was mixed, particularly around platform polarisation, local-first sovereignty, human oversight, and the distribution of productivity gains. The Human Time Dividend remains an original synthesis and proposed measurement framework rather than an established economic indicator.