Wellbeing Futures: Innovation, Policy and Research

wfipr.com

Research Article

Dissolve What Extracts: AI, Wellbeing, and the Composition of Economic Life (A Compositional Theory of the AI Wellbeing Dividend, with a Constructive Policy Triad)

John Ricketts[1,*] and Chris D. Beaumont[2]

[1]*AI+Wellbeing Institute, International College of Liberal Arts (iCLA), Yamanashi Gakuin University, Japan; E-mail: ricketts.john@c2c.ac.jp*

[2]*University of Tokyo, Tokyo, Japan*

(*Corresponding Author)

Cite This: https://doi.org/10.65638/2978-882X.2026.2.05

ABSTRACT: AI's most important wellbeing dividend may not come from producing more output, but from changing the composition of economic life. By reducing the administrative, cognitive, financial, and procedural burdens that absorb human capacity without building capability, AI can dissolve forms of **extractive friction**. But unless the resulting gains are protected from rent migration and redirected into formative infrastructure, extraction will not disappear; it will reappear in AI-native form. This paper develops a **compositional theory of AI and wellbeing**, distinguishing the philosophical hinge between extractive friction (burdens that drain capacity without building it) and **formative friction** (difficulty that develops judgment, skill, relationship, and responsibility). The wellbeing dividend is mapped across four BANI channels, with Brittleness treated explicitly as a hypothesis-generating channel rather than an established mechanism. The paper distinguishes theory from prescription. The prescription is a triadic policy mechanism: *dissolution* compresses extractive friction; *anti-recapture* prevents migrated rents from forming a new layer of extraction (taxation is its principal fiscal instance, alongside interoperability mandates, fiduciary duties, antitrust scrutiny, and public AI provision); and *formation funding* directs released capacity into relational and developmental infrastructure. The triad is operationalised through a five-test evaluative toolkit (Capability, Proportionality, Distribution, Migration, Formation) designed to assess whether any AI deployment or proposed policy actually advances compositional change rather than merely automating the existing equilibrium.

Keywords: AI, Economic composition, Extractive dissipation, Extractive friction, Formative friction, BANI, Administrative burden, Rent migration, Wellbeing, Adaptive infrastructure.

1. INTRODUCTION: THE GOVERNING MODEL

Standard AI debates ask how much more output AI will produce; this paper asks what kinds of economic activity AI allows societies to stop reproducing.

This paper uses a compositional framework that decomposes national economic overhead into three categories — Institutional Infrastructure (capability-building), Remediation Load (repair of preventable harm), and Extractive Dissipation (value transfer without proportionate capability gain) — on the premise that what economies do with overhead matters more than how much they have.

Received: July 13, 2026
Accepted: October 22, 2026
Published: November 02, 2026

1.1. The Governing Model

The paper organises around a single governing model. AI can create a wellbeing dividend only if three conditions are jointly met:

Condition 1 — Dissolution. Extractive friction must be dissolved. The AI-enabled mechanism: AI is deployed to compress burdens that absorb human capacity without building capability. Section 3 develops the capability shift; Section 5 sets out the evaluative tests.

Condition 2 — Anti-recapture. Concentrated AI rents must be prevented or captured. The rent-prevention mechanism: without it, dissolution gains migrate to whoever controls the new substrate. Section 4 develops the historical pattern, the broader toolkit, and the fiscal instance.

Condition 3 — Formation. Released capacity must be reinvested in formative infrastructure. The wellbeing-conversion mechanism: time, attention, and revenue freed by dissolution and protected by anti-recapture must flow into the relational and developmental institutions that convert saved capacity into resilience. Section 5 develops what counts as formation.

The three conditions form a single causal loop: dissolution releases capacity, anti-recapture protects it from migration, formation converts it into wellbeing. Two compressed formulations stand for the architecture as a whole — one diagnostic, one prescriptive:

AI changes the substrate of extraction; policy determines whether extraction dissolves or migrates.

Dissolve what extracts. Tax what concentrates. Fund what forms.

In the compressed triad, “tax what concentrates” names the fiscal core of a broader anti-recapture agenda: prevent rents where possible, and capture them where they form. Section 4 develops the broader toolkit; the slogan is the fiscal centre of gravity, not the whole machinery.

1.2. Theory and Prescription

The argument proceeds in two registers. The **theoretical register** identifies extractive friction as a compositional target and explains why AI makes its dissolution newly actionable. The **prescriptive register** asks what institutional design is required to make that dissolution durable. The registers are separable: a reader sceptical of the prescriptive case need not reject the theoretical one; a reader convinced by the diagnostic case still has further work to do before the policy programme follows.

1.3. Roadmap: Four Questions

The remainder of the paper answers four questions, each tracking one piece of the governing model.

1. **Section 2.** What is extractive friction? Develops the philosophical hinge between extractive and formative friction, the proportionality test, and the four BANI channels through which extractive friction harms wellbeing.

2. **Section 3.** Why can AI dissolve it now? Describes the capability shift and the three structural conditions — fiduciary alignment, procedural recognition, effective access — under which the capability translates into actual dissolution.

3. **Section 4.** Why will extraction migrate without anti-recapture? Establishes the migration claim as a historical regularity and develops the broader anti-recapture toolkit, with taxation treated in detail as the fiscal instance most directly tied to formation funding.

4. **Section 5.** What counts as formation, and how do we evaluate? Defines formative infrastructure with empirical anchoring in social-support evidence, then operationalises the entire framework through a five-test evaluative toolkit.

Section 6 concludes.

1.4. Approach

This is a conceptual, theory-building paper. It does not collect or analyse new data, so sampling and data-collection procedures do not apply. The argument is built in three steps. First, it takes the compositional categories set out in Section 1 and the findings of a companion paper (Ricketts & Beaumont 2026) as its starting point. Second, it draws on literatures on administrative burden, sludge, rent extraction, platform concentration, capability, and recent evidence on generative AI, selected for their direct bearing on each condition of the governing model. Third, it applies the resulting framework to one worked case, healthcare claims administration (Section 5.3), chosen because it is large in scale, well documented across countries, and already subject to AI deployment on both sides.

2. WHAT IS EXTRACTIVE FRICTION?

This section answers the first question: what counts as extractive friction, and why does it matter for wellbeing?

Modern life is increasingly spent maintaining access to systems that were meant to support life: healthcare, taxation, law, housing, education, finance, welfare, employment, and public administration. Citizens decode bills, contest denials, complete forms, interpret eligibility rules, manage compliance obligations, compare contracts, challenge errors, wait on phone lines, and navigate procedures whose complexity often serves no human purpose proportionate to the burden imposed. These activities generate income, employment, fees, billable hours, software contracts, consultancy work, and measurable economic activity. Yet much of this activity does not heal, teach, house, protect, or form anyone. It converts human attention into institutional maintenance.

We call this **extractive dissipation**: the loss of human capacity into activities that transfer value, preserve asymmetry, or absorb effort without producing proportionate gains in human flourishing (cf. Mazzucato 2018; Graeber 2018). The term is deliberately stronger than inefficiency. Inefficiency suggests waste. Extractive dissipation suggests captivity. It names the condition in which people are forced to spend life-energy servicing systems whose burdens are not justified by the capabilities they create.

Two related terms recur throughout the paper and are worth distinguishing here. **Extractive friction** is the micro-level mechanism: the everyday burdens experienced as forms, denials, opaque contracts, procedural asymmetry, compliance load, and administrative exhaustion. **Extractive dissipation** is the macro-level compositional result: the economic activity sustained by those frictions — billable hours, fees, navigation industries, duplicated administrative labour, avoidable disputes, and defensive expenditure. Not all dissipation is experienced as friction, and not all friction creates measurable dissipation, but many major forms of extractive dissipation are produced through extractive friction. The paper uses *extractive dissipation* when discussing the compositional accounting framework set out in Section 1, and *extractive friction* when discussing the lived burdens AI is positioned to dissolve.

2.1. Forms of Extractive Friction and Dissipation

Having distinguished the micro-level mechanism (extractive friction) from the macro-level result (extractive dissipation), this section sets out the recurring forms in which they appear together.

Rent-preserving complexity. A system becomes difficult enough to navigate that intermediaries can charge for access, even when the underlying task could be simpler. Tax preparation is the clearest example. The citizen's need is straightforward: comply with the law. But the procedural environment creates demand for paid navigation.

Administrative exhaustion. The system does not openly deny access; it makes access costly enough that people give up. Healthcare claims, welfare benefits, visa procedures, university bureaucracy, insurance appeals, and housing disputes often work this way. The burden is shifted from institution to individual (Herd & Moynihan 2018; Halling & Baekgaard 2024).

Informational asymmetry. Contracts, bills, financial products, legal notices, or eligibility criteria are written in ways that ordinary people cannot confidently interpret. The person becomes dependent on the specialist not because the matter is inherently beyond them, but because the system has made comprehension scarce.

Strategic opacity. A process is designed so that individuals cannot easily know whether they have been treated fairly. Pricing, denial, ranking, recommendation, risk scoring, and enforcement decisions may all be insulated from meaningful challenge.

Defensive participation. People are forced to spend time and money protecting themselves from avoidable harms produced by other actors: scams, predatory contracts, manipulative advertising, hidden fees, exploitative platform design, or excessive legal risk (Sunstein 2022).

The empirical scale is large. A major eight-country comparison using 2010–2011 data found that administrative costs accounted for 25.3% of total US hospital expenditures, the highest share among the countries studied; Scotland and Canada had the lowest administrative costs (Himmelstein *et al*. 2014). The differential is composed almost entirely of activities that produce no clinical benefit: billing, eligibility verification, claims adjudication, prior authorisation, denial appeals. More recent US-specific work estimates that administrative simplification could save about $265 billion, or roughly 28% of annual administrative spending, without compromising quality or access (Sahni, Carrus & Cutler 2021).

2.2. The Philosophical Hinge: Extractive vs Formative Friction

The aim of this paper is not a frictionless society. That would be a poor and dangerous ideal. Some friction extracts; some friction forms. The distinction is the philosophical hinge of the entire argument.

Contemporary technology often treats friction as the enemy. Platforms reduce clicks, shorten pathways, remove waiting, personalise interfaces, anticipate desire, and optimise for seamlessness. In many domains, this is useful. But a general war on friction is a mistake. A frictionless life is not necessarily a flourishing life. It may be passive, dependent, shallow, and easily manipulated.

Human beings need some difficulty. We are formed through tasks that resist us and through people who do not immediately conform to us. Conversation, education, care, craft, democratic participation, friendship, parenting, scholarship, artistic practice, and moral responsibility all involve forms of friction. They require patience, attention, interpretation, humility, negotiation, discipline, and repair. Sennett (2008) and Crawford (2009) develop versions of this point with respect to skilled practice and material work; the broader claim extends well beyond craft.

Extractive friction is difficulty that drains capacity without forming it. It consumes time, cognition, money, trust, and emotional bandwidth while producing little or no durable gain in agency, skill, relationship, or understanding.

Formative friction is difficulty that builds capacity. It may be uncomfortable, demanding, or slow, but it develops judgment, courage, care, skill, responsibility, attachment, and resilience.

Note on the term “capability”. In capability terms (cf. Sen 1999; Nussbaum 2011), extractive friction consumes resources without expanding substantive freedom, while formative friction can enlarge the practical ability to choose, act, relate, and participate. The paper's use of “capability” is therefore not merely productivity-oriented; it refers to the conversion of economic activity into agency, relationship, and practical freedom.

The same life domain can contain both. A medical appointment may be formative in the broad sense: a patient is heard, interpreted, cared for, advised, and supported. But the phone tree to book it, the insurance codes attached to it, the denial letter that follows it, and the billing dispute that must be fought afterwards may be extractive. A school may be formative: a place where students encounter others, develop attention, accept correction, and acquire orientation. But admissions bureaucracy, assessment overload, compliance paperwork, and parent-facing procedural confusion may become extractive. A legal system may be essential to justice, yet routine legal intimidation and procedural opacity may be extractive.

The operational principle. *AI should compress the administrative surface of formative encounters while preserving the encounter itself.*

The aim is not to replace the doctor, teacher, mentor, counsellor, neighbour, colleague, or civic association. The aim is to dissolve the procedural crust that prevents people from reaching them, trusting them, or having enough bandwidth to participate in them.

This distinction also answers a likely objection. Critics may worry that AI-enabled convenience will produce weaker citizens: less patient, less skilled, less socially embedded, less capable of dealing with difficulty. That objection has force if AI is used to remove formative friction. But it loses force when AI is directed against extractive friction. The question is not whether people should face difficulty. The question is which difficulties are worth facing.

2.3. The Proportionality Test

Distinguishing extractive from formative friction requires a working test. Not every instance of complexity is extractive. A tax system may be complex because a democracy has chosen to subsidise childcare, housing, education, disability, research, or regional development through the tax code. A legal process may be complex because rights must be protected. Healthcare review may prevent fraud or unsafe treatment. Financial regulation may reduce systemic risk. Complexity becomes extractive not because it exists, but because the capability gains it produces are disproportionate to the burden it imposes.

The proportionality test. *A burden becomes extractive when the human cost of navigating it exceeds the capability, justice, safety, or coordination it plausibly creates — especially where the same legitimate function could be achieved through a less burdensome design. The test has four parts: What legitimate function does the burden serve? Is there evidence that it actually serves that function? Could the same function be achieved with less human cost? Who bears the burden, and does it fall most heavily on those with least capacity?*

The test must be calibrated. Some burdens were never necessary from the standpoint of flourishing. Others may once have served a legitimate function but became extractive as procedural safeguards hardened into rent-preserving systems. Others remain mixed: partly necessary, partly extractive. AI's role is not to abolish complexity indiscriminately, but to make these distinctions practically actionable.

2.4. How Extractive Friction Harms Wellbeing: Four BANI Channels

Extractive friction does not produce a single, undifferentiated wellbeing harm. It harms people through distinct channels. The BANI framework — Brittle, Anxious, Nonlinear, Incomprehensible (Cascio 2020) — is useful because it treats vulnerability as a formative vector rather than a single reflective condition. Ricketts and Beaumont (2026) argue that the four letters describe distinct failure modes (structural, affective, dynamic, and epistemic respectively) that should not be expected to co-vary as interchangeable indicators of one underlying state.

The four channels do not have identical evidential status in this paper. Incomprehensibility and Anxiety are supported by adjacent literatures on administrative burden, sludge, scarcity, and cognitive bandwidth (Herd & Moynihan 2018; Sunstein 2022; Mullainathan & Shafir 2013). Nonlinearity is supported by work on the regressive incidence of administrative burdens. Brittleness is treated more conceptually, drawing on cascading-failure and fragility logics (Taleb 2012). The purpose of distinguishing the channels is not to claim that all four mechanisms are already established, but to specify the channels through which the AI wellbeing dividend should be tested.

Incomprehensibility

Extractive systems often begin with incomprehensibility. The bill is unreadable. The contract is opaque. The eligibility rule is conditional. The denial letter is coded. The appeal process is hidden. The risk score is unavailable. The tax rule is buried. The platform decision is unexplained. Incomprehensibility is not merely inconvenience. It is a form of power. It converts the person from participant into subject.

Anxiety

Anxiety is often treated as a private psychological state. In administered life, it is frequently a rational response to unresolved institutional threat. The unread bill, the unexplained charge, the ambiguous visa status, the missed deadline, the unresolved claim — each occupies cognitive and emotional space. Extractive friction plausibly contributes to anxiety because it creates uncertain obligations under asymmetric conditions. The mechanism is consistent with administrative-burden and scarcity accounts of cognitive load: unresolved administrative threats consume the same executive resources required for planning, care, learning, self-regulation, and participation.

Brittleness

Administrative systems often become harmful not through single catastrophic decisions but through cascades. A form is misunderstood. A deadline is missed. A claim is denied. A payment is delayed. A debt is triggered. A credit score falls. A housing application fails. A health condition worsens. This is brittleness: the tendency of small failures to propagate through linked systems. The person experiences life as a chain of institutional dominoes. This paper treats Brittleness as a **hypothesis-generating channel** rather than an established mechanism. A testable form of the hypothesis: *in domains where AI authorised representatives gain procedural standing, the rate at which small administrative errors propagate into downstream consequences (debt referral, eviction filing, benefit termination, credit-score events) should decline measurably under an appropriate comparison or synthetic-control design.* If it does not, the channel as specified here would need revision.

Nonlinearity

The harms of extractive friction are not evenly distributed. They are nonlinear. For affluent households, administrative friction is irritating but often manageable: they can buy expertise, delegate tasks, absorb delays, hire lawyers, recover from mistakes. The wealthy already purchase human interposition. For low-income households, the same friction can be catastrophic. A small fee, delay, denial, mistake, or missed form can trigger cascading consequences. The burden is not merely larger relative to income. It is structurally different because there is less buffer (Herd & Moynihan 2018; Mullainathan & Shafir 2013).

3. WHY CAN AI DISSOLVE IT NOW?

This section answers the second question: what specifically has changed about AI's capability that makes the dissolution of extractive friction newly actionable, and what conditions must hold for the capability to translate into actual dissolution?

3.1. The Capability Shift

Extractive friction has long been recognised. People know when systems are needlessly difficult. They know when a process is designed to exhaust them. They know when a fee exists because complexity has been preserved. They know when a contract is unreadable by design. They know when a denial letter is not the end of a decision but the beginning of a battle they lack the strength to fight. The novelty is not diagnosis. The novelty is capacity.

Before AI, contesting extractive friction required expertise, time, memory, drafting skill, confidence, and the endurance to survive procedural delay. These were unevenly distributed. Wealthy households purchased them through lawyers, accountants, financial advisors, and other professional intermediaries. Less wealthy households absorbed the burden directly — or gave up. AI changes the balance by making several previously scarce capabilities widely available at low marginal cost: it can render language legible, make procedure navigable, draft inexpensively, hold memory persistently, compare options, surface anomalies, prepare representation-ready materials, and support representation where procedural recognition exists. Experimental and field evidence suggests that these gains are largest for less experienced and lower-skilled users (Noy & Zhang 2023; Brynjolfsson, Li & Raymond 2025).

The compositional consequence is that AI does not simply accelerate existing activity; it changes which activities remain necessary. If citizens can file taxes directly, the routine tax-preparation industry shrinks. If patients can contest denials effectively, denial-based administrative strategies become less profitable. If tenants can understand leases and challenge abusive clauses, some legal intimidation loses force. If small businesses can manage compliance without expensive intermediaries, regulatory burden becomes less regressive.

This does not mean professional expertise disappears. Many legal, medical, financial, and administrative situations require human judgment, ethical responsibility, and institutional accountability. But AI can change the boundary between genuinely expert work and work that exists because non-experts have been artificially disempowered. The aim is not to dissolve expertise but to dissolve dependency on expertise where that dependency is produced by avoidable opacity. AI's distinctive value therefore lies less in replacement than in interposition: it stands between individuals and systems, not to remove institutions, but to rebalance the relationship between institutions and those subject to them.

One clarification matters here. **AI interposition should be treated as a transitional dissolution mechanism, not the final institutional ideal.** The highest form of dissolution is not an AI assistant that helps people survive an unnecessary burden, but a redesigned institution in which the burden no longer exists. Where systems can be made simpler at source, that is the deeper compositional goal. Interposition is the response to systems that cannot or will not be redesigned in the short run; it should not be allowed to become a permanent substitute for the redesign itself.

3.2. Direction of Agency: When AI Dissolves Extraction and When It Intensifies It

The capability shift is not directionally neutral. The same AI capability can dissolve extraction or intensify it. A patient-side AI can help contest denied claims; an insurer-side AI can deny claims faster, more selectively, and more opaquely. A citizen-side AI can make tax filing easier; a tax-preparation firm can use AI to reduce its own costs while lobbying to preserve the complexity that sustains demand. A tenant-side AI can identify abusive lease terms; a landlord-side AI can optimise rent increases, screen tenants, and automate pressure. A worker-side AI can explain rights and draft grievances; an employer-side AI can monitor performance, personalise incentives, and weaken bargaining (Zuboff 2019). A consumer-side AI can compare products and resist manipulation; a platform-side AI can personalise persuasion and make manipulation more adaptive. Acemoglu and Johnson (2023) make the general historical case: whether the gains from a technology are widely shared depends on the institutions that direct it.

AI is therefore not intrinsically compositional in the direction we want. It becomes compositional for wellbeing only when deployed to reduce extractive friction rather than to deepen it. Three structural conditions follow.

Loyalty (fiduciary alignment)

AI systems that act on behalf of individuals must be loyal to those individuals. The legal-policy term for this is *fiduciary alignment*: a duty to act in the principal's interest rather than the agent provider's. If an agent is funded by advertising, affiliate commissions, platform incentives, data capture, or seller-side integration, fiduciary alignment is compromised. The agent may appear to help the user while subtly routing them through

monetised pathways. Loyalty is not a property of model capability alone. It is a property of the deployment contract. Who pays? Who benefits? Who sees the data? Who can override the recommendation? Who is the principal? Without fiduciary alignment, AI interposition becomes another interface layer through which extraction is personalised.

Standing (procedural recognition)

Even AI systems with fiduciary alignment may fail if they lack *procedural recognition*: the formal admission of AI-mediated representation by the institutions a user must engage with. Many systems recognise only certain kinds of representatives, signatures, filings, and communications. If AI can explain a denial but cannot help submit the appeal, the burden remains partly intact. If it can draft a letter but the institution refuses to accept AI-assisted representation, the asymmetry persists. The question is not only whether AI can help people understand systems, but whether systems will permit AI-mediated action within them. This is especially important in healthcare, welfare, legal, tax, and immigration contexts, where procedural recognition determines whether assistance becomes power.

Distribution (effective access)

Finally, AI capability must not be monopolised by the same actors who profit from complexity. The relevant condition is *effective access*: the practical availability of capable, affordable, trustworthy, and legally usable AI on the user side of every transaction. If high-quality AI remains concentrated among platforms, insurers, employers, financial institutions, and large firms, then AI may increase asymmetry rather than reduce it (Khan 2017). The institution becomes more capable while the individual remains exposed. Effective access does not necessarily require that every model be open or public. But it does require that user-side capability exist at a quality and cost that genuinely rebalances the institutional encounter.

The compositional case for AI therefore cannot be separated from institutional design. AI dissolves extraction only when fiduciary alignment, procedural recognition, and effective access are all present. The compressed shorthand — loyal enough to serve, recognised enough to act, distributed enough to rebalance — remains a useful summary, but the operative concepts are legal and infrastructural rather than dispositional.

4. WHY WILL EXTRACTION MIGRATE WITHOUT ANTI-RECAPTURE?

This section answers the third question. Even if dissolution succeeds, the gains do not automatically translate into wellbeing. Extraction migrates. Without an anti-recapture mechanism, the friction economy reappears as the AI-agent economy. Capture must therefore be designed in.

4.1. The Migration Claim

The migration of extraction across technological transitions is a regularity, not a speculative worry. Older intermediaries collapse while new platform intermediaries emerge; concentrated rents in information industries reform around whichever layer of the new stack remains controllable (Wu 2010; Khan 2017). Music streaming is the most studied recent case: digitalisation lowered the barriers to distribution while creating new forms of concentration at the platform and recommendation layer, with rent and value-capture migrating from labels and retailers to the platforms that aggregate listening (Hesmondhalgh 2021). The substrate changes; the rent persists. There is no reason to assume that AI will be the first transition in which rents simply disappear.

Extractive dissipation is not a fixed stock. It is an equilibrium outcome of actors seeking rents under existing conditions: capital and talent will search for whichever rent structure the new substrate makes available. The likely AI-native forms include subscription dependency, platform tolls, data rents, personalised persuasion, agent-to-agent transaction fees, and proprietary access to procedural capability (Zuboff 2019). Without a mechanism to capture concentrated gains and either redirect them or prevent their formation, the friction economy reappears as the AI-agent economy.

The anti-recapture principle. *The general principle is capture what concentrates. Taxation is the fiscal version of that principle — the focus of this section — but it is one instrument among several.*

4.2. The Broader Anti-Recapture Toolkit

Anti-recapture is not synonymous with taxation. The broader toolkit includes at least the following mechanisms, each addressing a different point of potential rent formation in the AI stack.

Interoperability mandates for AI agents reduce the risk that the platform layer locks users into proprietary agent ecosystems. *Data portability and user-controlled data*

vaults limit the ability of the data-advantage layer to become a permanent rent source. *Public or civic AI agents* for high-burden domains (healthcare appeals, benefits navigation, legal aid) ensure that user-side capability exists where private provision would underproduce it. *Fiduciary duties for AI agents* acting on behalf of individuals codify the loyalty condition of Section 3 into enforceable legal form. *Antitrust scrutiny* of vertical integration across cloud, model, and platform layers makes it harder for the most concentrated rent structures to form in the first place. *Procurement rules* for public-sector AI use can require interoperability, fiduciary alignment, and audit access as conditions of contract. *Open standards* for appeals, forms, claims, benefits, and machine-readable rights lower the cost of building user-side tools that interface with administered systems. *Rights to contest* and to human review in high-stakes automated decisions preserve the procedural surface on which Section 3.2's standing condition operates.

These mechanisms are not alternatives to taxation; they are complements. Some prevent rent formation; taxation captures rents that form despite prevention. A serious anti-recapture strategy uses both.

The remainder of this section develops the tax instrument in detail because it is the mechanism most directly tied to formation funding (Section 5): without revenue capture or savings capture, there is no reliable fiscal flow into the institutional infrastructure that converts dissolution gains into wellbeing. The other instruments in the toolkit are essential to preventing recapture, but they do not by themselves generate the revenue that the third leg of the triad requires.

4.3. Tax What Concentrates

The tax target is not AI deployment as such. Taxing productive use would slow beneficial diffusion. The target is AI rent: supernormal returns arising from scarcity, concentration, platform control, data advantage, or infrastructural chokepoints (cf. Mazzucato 2018).

The tax-design principle. *Tax AI rents, not AI use.*

The AI-stack structure makes this distinction more tractable than a generic tax on AI use. Rents tend to concentrate at three layers. *Compute and infrastructure*: a small number of cloud and chip providers control the capacity on which downstream AI deployment depends, with margins reflecting that scarcity. *Frontier model access*: a small number of frontier-model developers earn returns reflecting capability concentration rather than ordinary returns to capital and labour. *Platform-mediated agent ecosystems*: where AI agents become the dominant interface between consumers and services, the platform that controls the agent layer captures a share of every downstream transaction. Each layer exhibits the structural features — concentration, scarcity, infrastructural position — that distinguish rents from ordinary returns. Each is therefore a candidate tax base. None of this makes implementation easy: rent identification, transfer-pricing defence, and jurisdictional coordination remain hard problems on which the existing tax-design literature has only partial answers.

This concern is already visible in competition policy. The US Federal Trade Commission has scrutinised partnerships among major cloud providers and generative-AI firms, including their effects on competition and access to key AI inputs (FTC 2024, 2025). The UK Competition and Markets Authority has similarly treated foundation-model markets as a live competition and consumer-protection concern (CMA 2024). These developments do not prove the migration thesis, but they show that the relevant layers of the AI stack — cloud, model access, and platform integration — are already recognised by regulators as potential sites of concentration.

This is not unprecedented. Tax systems have long distinguished ordinary returns from rents in domains where scarcity, concentration, or public resources generate supernormal profits. Norway's petroleum tax system is the closest large-scale analogue: the petroleum tax regime is designed to capture resource rent while maintaining investment neutrality, with official sources describing a combined marginal tax rate of 78% (Norwegian Petroleum 2025). Norway also demonstrates the second half of the political economy: petroleum revenues were channelled into the Government Pension Fund Global, established to manage them with a long-term perspective for the benefit of future generations (NBIM n.d.). The relevance of the analogy is not the rate but the architecture — rent identification, ring-fenced capture, long-horizon redirection — which is exactly the architecture the compositional argument requires.

The analogy is of course imperfect. AI rents are less territorially fixed than petroleum rents, more intangible, and more vulnerable to transfer-pricing strategies. But the Norwegian case demonstrates that concentrated scarcity rents can be taxed without treating ordinary productive activity as the tax base.

Windfall and excess-profits taxes offer further design precedent. The UK Energy Profits Levy was introduced

as a temporary tax on exceptional profits of oil and gas companies, with threshold mechanisms linked to historically normal price ranges (HM Treasury 2024). These instruments show that tax systems can target exceptional or rent-like returns without taxing ordinary use, and that the threshold-and-exception architecture used for energy could be adapted for AI capability rents.

Digital services taxes provide a cautionary analogue. The OECD Pillar One process aims to reallocate taxing rights over a share of profits from the largest and most profitable multinational enterprises and, in doing so, to remove unilateral digital services taxes (OECD 2023). As of March/April 2026, the Pillar One Multilateral Convention remained not yet open for signature, and the European Parliament's Legislative Train recorded the Commission's September 2025 acknowledgement that discussions were “on hold” while it did not intend to table a new DST proposal at that stage (European Commission 2025; Tridico & Lalucq 2025; Baert 2026). The trajectory illustrates the jurisdictional difficulty AI taxation will face: even with broad design agreement among more than 140 jurisdictions, opening the convention for signature has proven politically intractable. National or regional instruments in jurisdictions that host AI infrastructure, control market access, or possess strong tax-enforcement capacity are therefore likely to precede any coordinated international architecture.

The design task is not to produce a perfect AI rent tax in the abstract. It is to adapt existing rent-identification instruments to the distinctive structure of AI: cloud concentration, model access, data advantage, platform control, market power, and agent-mediated dependency.

4.4. The GDP-Measurement Caveat

One implication of the migration argument is that the wellbeing dividend from dissolution is partly invisible to standard economic measurement. In some sectors, dissolving extractive activity may reduce measured GDP. This should be understood as a measurement artefact rather than a loss of productive capacity or human flourishing. When fewer people pay unnecessary fees, fight avoidable disputes, hire intermediaries to navigate artificial complexity, or repair preventable administrative failures, recorded activity may fall even as wellbeing rises. The compositional frame is necessary precisely because GDP does not distinguish activity that builds capability from activity generated by avoidable friction (cf. Mazzucato 2018).

Measured properly, the question is not whether the economy grows, but whether more of what it produces becomes human flourishing.

5. WHAT COUNTS AS FORMATION, AND HOW DO WE EVALUATE?

This section answers the fourth question and provides the evaluative closure to the framework, because the formation argument and the evaluative toolkit form a single operational endpoint.

5.1. Fund What Forms

The third leg of the triad is to redirect the dividend toward formation. This is where the paper departs most sharply from standard productivity narratives. The value of AI is not simply that it frees time for consumption, leisure, or more work. Released capacity should be directed toward the institutions and relationships that build durable human capability.

Formation funding matters because released capacity does not automatically become wellbeing. It may be absorbed into consumption, longer work intensity, platform rents, or private convenience. It becomes adaptive capacity only when channelled into institutions and relationships that improve people's ability to interpret change, rely on others, recover from shocks, and participate meaningfully in common life. In this sense, formation is the conversion mechanism between saved time and durable resilience — the leg of the triad that closes the BANI loop set out in Section 2.4 by strengthening the relational and developmental systems that reduce anxiety, prevent cascades, and improve recovery after shocks.

Formation infrastructure includes schools that cultivate orientation rather than only credential throughput; libraries as civic commons; community health workers who convert abstract systems into trusted relationships; youth clubs and mentoring systems; public-interest legal support; local arts and cultural institutions; care networks; mutual-aid scaffolding; and public spaces where weak ties can form before crisis (Putnam 2000).

Formation funding should therefore be ring-fenced through transparent local or sectoral funds, with priority given to communities and domains where the Distribution test shows the highest extractive-friction burden. This makes the wellbeing-conversion mechanism institutional rather than aspirational: the dividend reaches formation by design, not by hope.

These are not soft extras. They are adaptive infrastructure. In the Ricketts and Beaumont (2026) analysis, the relational form of capacity (social support) appears to do more of the descriptive work in the moderation patterns observed than the biological form (healthy life expectancy). The strong causal interpretation is qualified by their robustness analysis and remains provisional pending the country-specific identification strategy of their planned follow-up; but even on the descriptive evidence, the pattern suggests that what allows societies to absorb disruption may be at least as much relational as material.

A society's ability to absorb disruption depends not only on income, health, or formal institutions, but on relational capacity: whether people have someone to call, whether institutions are trusted, whether citizens can interpret change together, whether help arrives before crisis, whether young people are formed in contexts of responsibility and belonging.

The compositional AI dividend should therefore be measured not only in productivity gains or fiscal savings, but in capacity returned to relational life. That is the loop.

5.2. The Evaluative Toolkit

The four questions of Sections 2–5 also describe what a proposed AI deployment must be evaluated against. The proportionality test introduced in Section 2 identifies whether a burden is extractive in the first place. For application to AI policy, however, four further tests are needed. The full toolkit is summarised below.

The tests are not intended to produce a single composite score. They are intended to discipline judgment: a proposed AI deployment may pass one test while failing another, and the analytic value of the toolkit lies in surfacing which test it fails.

The five tests are designed to be applied in sequence. The Capability and Proportionality tests together identify whether a burden is a candidate for dissolution at all. The Distribution test specifies the wellbeing rationale and incidence. The Migration test disciplines the institutional design (without it, dissolution is a one-period claim and the gains return as new rents). The Formation test asks whether the dividend reaches the relational and developmental infrastructure on which long-run resilience depends.

The triad and the toolkit map onto each other directly. *Dissolution* is informed by the Capability and Proportionality tests; *anti-recapture* (in any of its forms — taxation, interoperability mandates, fiduciary duties, antitrust action, public AI provision) is the response to the Migration test; *Formation funding* is the institutional answer to the Formation test. The Distribution test cuts across all three. A policy or deployment that passes only the first two tests is descriptively interesting but structurally insufficient: it identifies friction worth dissolving but does not protect against rent migration or convert the dividend into resilience.

5.3. Applying the Toolkit: Healthcare Claims Administration

The toolkit is most legible when worked through a single case. Healthcare claims administration is a useful case because it is large in scale, a strong candidate for extractive-friction analysis, and increasingly subject to AI deployment on both sides of the patient–insurer relationship.

The **Capability test** asks whether the denial-and-appeal burden builds patient, clinical, or institutional capability. Patients may learn fragments of procedural survival, but the burden does not build health, clinical judgment, institutional trust, or transferable agency proportionate to its cost. Clinicians spend time on documentation that

Table 1: The evaluative toolkit: five tests for AI deployments

Test	Question	Function
Capability test	Does the burden build human or institutional capability?	Distinguishes formative from extractive friction.
Proportionality test	Is the burden justified by the value it creates?	Identifies where the human cost exceeds the capability, justice, safety, or coordination value produced.
Distribution test	Who bears the burden, and who benefits if it is dissolved?	Surfaces nonlinear and regressive incidence; distinguishes progressive from neutral or regressive AI deployment.
Migration test	If AI removes this burden, where does the rent go?	Forces explicit accounting of whether dissolution is durable or merely relocates extraction.
Formation test	Does the saved capacity flow into human development, or is it absorbed elsewhere?	Closes the loop: tracks whether the wellbeing-conversion mechanism is actually operating.

does not improve care; the system as a whole does not become more responsive over time. The burden generates billing-code expertise, administrative employment, and intermediary fees, but these accrue to the navigation of the maze rather than to the production of health.

The **Proportionality test** asks whether fraud prevention, cost control, or coordination justify the cognitive and financial burden placed on patients. Some prior-authorisation activity does prevent fraud or unsafe over-treatment, and some claims adjudication does enforce coverage rules that reflect democratic decisions about what insurance pays for. But cross-country comparison (Himmelstein *et al*. 2014; Sahni, Carrus & Cutler 2021) shows that the burden imposed in the US system substantially exceeds the cost-control or fraud-prevention value generated, by margins that comparable single-payer systems achieve at a fraction of the administrative cost. The scale of the burden strongly suggests that much of the US claims-administration burden exceeds what fraud prevention, cost control, or coordination can justify.

The **Distribution test** asks who bears the burden. Patients with fewer resources bear it disproportionately: they have less time, less expertise, less procedural confidence, and less capacity to absorb the consequences of cascading administrative failure. Wealthy patients buy human interposition (advocates, billing specialists, attorneys); low-income patients absorb the burden directly or give up. The incidence is regressive and the dissolution dividend would be progressive.

The **Migration test** asks where the rent goes if AI removes the friction. This is the test that determines whether dissolution is durable. If patient-side AI compresses the appeal burden but insurers deploy AI to deny faster and more selectively, the rent has not migrated — it has remained in the same hands. If user-side AI is provided through proprietary platforms that take a share of every successful appeal, the rent has migrated to AI deployers. If AI capability remains concentrated among insurers and absent on the patient side, dissolution does not occur at all. The test forces explicit accounting of which scenario is being designed for. Anti-recapture instruments addressing this case include fiduciary duties for patient-side AI agents, interoperability mandates that prevent appeal data lock-in, public or civic AI provision in healthcare appeals, and procedural recognition of AI-mediated representation in claims processes.

The **Formation test** asks whether the saved capacity flows into care relationships, community health workers, patient support, and the relational infrastructure that produces health, or whether it is absorbed into lower insurer costs and AI-platform margins. This is the closing test. A reform that compresses administrative cost without redirecting any portion of the savings into formation would pass the first four tests and fail the fifth. The Formation test is what distinguishes a productivity story from a wellbeing story.

Healthcare claims administration thus illustrates the toolkit's analytical structure: the first three tests identify the case for dissolution, the fourth disciplines the institutional design, and the fifth closes the wellbeing loop. A policy design that addresses only the first three is incomplete; one that addresses all five operationalises the full triad.

5.4. Further Domains

Applying the Capability and Proportionality tests identifies further candidate domains for AI-driven dissolution: those where burdens are sustained by informational asymmetry, procedural complexity, routinised drafting, memory load, eligibility confusion, or strategic opacity.

Tax preparation. Routine tax compliance should not require a large navigation industry. AI could support direct filing, explain obligations, check errors, and reduce dependency on paid intermediaries. The deeper goal is not faster tax preparation, but a system in which routine citizens no longer need tax preparation as an industry.

Routine legal services. Wills, leases, small claims, employment rights, immigration forms, debt disputes, and consumer complaints often require basic legal understanding that is priced beyond ordinary reach. AI can expand access to legal comprehension and procedural action while preserving the role of qualified human lawyers for complex, high-stakes, or discretionary matters.

Small-business compliance. Compliance burdens are often regressive because fixed costs do not scale down. AI can help small organisations understand obligations, prepare documentation, avoid penalties, and compete without unnecessary dependence on consultants.

Housing, finance, education administration, welfare systems, and consumer protection all provide further domains. In each case, the test is not whether AI can automate a task, but whether it can reduce the share of human life spent navigating avoidable asymmetry — and

whether the released capacity, protected from migration, ultimately reaches formation.

5.5. Failure Modes

The framework also identifies failure modes. AI may preserve extractive friction by making bad systems easier to navigate rather than simpler to redesign. It may generate false confidence through erroneous procedural, legal, or medical advice. It may trigger adversarial escalation, with institution-side AI denying, delaying, or pricing more adaptively than user-side AI can contest. It may create platform dependency if the user-side agent becomes a new tollgate. It may convert assistance into surveillance if sensitive user data are monetised. And it may deepen inequality if high-quality AI interposition remains available only to those already able to buy procedural help. These failure modes do not defeat the compositional argument, but they define the institutional conditions under which it can fail.

5.6. Relation to Existing Literature

The framework builds on several literatures. Like Mazzucato (2018) and Graeber (2018), it separates value creation from value extraction, but it locates extraction in the everyday burdens of administered life rather than only in sectors or occupations. It extends the administrative-burden and sludge literatures (Herd & Moynihan 2018; Sunstein 2022; Halling & Baekgaard 2024), which document these burdens and their regressive incidence, by asking what happens to them when user-side AI makes them cheap to contest. It agrees with Acemoglu and Johnson (2023) that the gains from technology depend on institutional direction, and specifies that direction as three conditions and five tests. It is consistent with early evidence that generative AI helps less experienced users most (Noy & Zhang 2023; Brynjolfsson, Li & Raymond 2025), while arguing that such gains reach wellbeing only if they are protected from rent migration and reinvested in formation. Its contribution is to connect these literatures in one compositional model with a toolkit that can be applied to specific deployments and policies.

6. CONCLUSION

The paper's governing model identifies three conditions under which AI delivers a wellbeing dividend rather than a productivity story. Dissolution is the AI-enabled mechanism; anti-recapture (of which taxation is the principal fiscal instance) is the rent-prevention mechanism; formation funding is the wellbeing-conversion mechanism. The five-test toolkit operationalises all three. The headline claim follows: *AI changes the substrate of extraction; policy determines whether extraction dissolves or migrates.* Without institutional design, the friction economy becomes the agent economy — the billing maze becomes the subscription maze, the professional gatekeeper becomes the platform gatekeeper.

Dissolve what extracts. Tax what concentrates. Fund what forms.

The goal is not a frictionless society. It is a society capable of distinguishing between the friction that extracts from us and the friction that forms us — and dissolving the first so that the second can matter again.

Three practical implications follow. Governments deploying or regulating AI in high-burden domains should apply the five tests before deployment, not after. Anti-recapture instruments, especially fiduciary duties for user-side agents, interoperability mandates, and procedural recognition of AI-assisted representation, should be designed alongside dissolution rather than added later. And part of any savings from dissolution should be ring-fenced for formation. Future research should test the framework empirically: measuring changes in administrative burden and its incidence where user-side AI is introduced, testing the Brittleness hypothesis in Section 2.4, and tracking whether released capacity reaches formative infrastructure.

The central question for AI policy is therefore not only what AI makes possible, but what forms of economic life it allows societies to stop reproducing.

FUNDING

This research received no external funding.

INSTITUTIONAL REVIEW BOARD STATEMENT

Not applicable.

INFORMED CONSENT STATEMENT

Not applicable.

DATA AVAILABILITY STATEMENT

No new data were created or analyzed in this study. Data sharing is not applicable to this article.

COMPANION PAPER

The present paper is self-standing, but it extends a companion paper by the same authors, Ricketts and

Beaumont (2026), *Wellbeing, VUCA(t), and BANI: The Unequal Human Burden of Global Disruption*, which develops the formative-vector reading of BANI used in Section 2.4 of the present paper, and the social-support evidence drawn on in Section 5.1. That paper operationalises BANI as a formative vector and shows that capacity moderates the descriptive VUCA → anxiety pattern; the present paper extends it into the question of how AI changes the substrate of extraction and what institutional design is required to make the change durable.

ACKNOWLEDGMENTS

We thank colleagues at the AI+Wellbeing Institute for comments on earlier drafts. The paper benefited from iterative methodological feedback that shaped successive draft architectures, particularly the integration of the migration argument, the calibration of the BANI evidential register, the addition of the broader anti-recapture toolkit, and the restructuring around the four-question architecture used in v2.0.

CONFLICTS OF INTEREST

The authors declare no competing interests.

DISCLAIMER/PUBLISHER'S NOTE

The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of the publisher and/or the editor(s). The publisher and/or the editor(s) disclaim responsibility (or use the journal's standard note; "MDPI" is left over from a template) for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

REFERENCES

Acemoglu, D., & Johnson, S. (2023). Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity. New York: PublicAffairs.

Baert, P. (2026). The implementation of the OECD global agreement on re-allocation of taxing rights (Pillar One). European Parliamentary Research Service, Legislative Train Schedule. As of 20 March 2026. Accessed 30 April 2026. europarl.europa.eu/legislative-train

Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. Quarterly Journal of Economics, 140(2), 889–942. https://doi.org/10.1093/qje/qjae044

Cascio, J. (2020). Facing the Age of Chaos. Medium, 29 April.

CMA. (2024). AI Foundation Models: Update Paper. Competition and Markets Authority, 11 April. gov.uk/cma-cases/ai-foundation-models-initial-report

Crawford, M. B. (2009). Shop Class as Soulcraft: An Inquiry into the Value of Work. New York: Penguin Press.

European Commission. (2025). Response to European Parliament question O-000022/2025 on taxation of large digital platforms in the light of international developments. Plenary debate, September 2025. Summarised in Baert (2026).

FTC. (2024). FTC launches inquiry into generative AI investments and partnerships. Federal Trade Commission, 25 January. ftc.gov/news-events/news/press-releases/2024/01/ftc-launches-inquiry-generative-ai-investments-partnerships

FTC. (2025). FTC issues staff report on AI partnerships and investments study. Federal Trade Commission, 17 January.

Graeber, D. (2018). Bullshit Jobs: A Theory. New York: Simon & Schuster.

Halling, A., & Baekgaard, M. (2024). Administrative burden in citizen-state interactions: A systematic literature review. Journal of Public Administration Research and Theory, 34(2), 180–195. https://doi.org/10.1093/jopart/muad023

Herd, P., & Moynihan, D. P. (2018). Administrative Burden: Policymaking by Other Means. New York: Russell Sage Foundation. https://doi.org/10.7758/9781610448789

Hesmondhalgh, D. (2021). Is music streaming bad for musicians? Problems of evidence and argument. New Media & Society, 23(12), 3593–3615. https://doi.org/10.1177/1461444820953541

Himmelstein, D. U., Jun, M., Busse, R., Chevreul, K., Geissler, A., Jeurissen, P., Thomson, S., Vinet, M.-A., & Woolhandler, S. (2014). A comparison of hospital administrative costs in eight nations: US costs exceed all others by far. Health Affairs, 33(9), 1586–1594. https://doi.org/10.1377/hlthaff.2013.1327

HM Treasury. (2024). Energy Profits Levy: Energy Security Investment Mechanism legislation. London: HM Government. gov.uk/government/publications/energy-profits-levy-energy-security-investment-mechanism-legislation

Khan, L. M. (2017). Amazon's antitrust paradox. Yale Law Journal, 126(3), 710–805.

Mazzucato, M. (2018). The Value of Everything: Making and Taking in the Global Economy. London: Allen Lane.

Mullainathan, S., & Shafir, E. (2013). Scarcity: Why Having Too Little Means So Much. New York: Times Books.

NBIM. (n.d.). About the fund: Government Pension Fund Global. Norges Bank Investment Management. Accessed 30 April 2026. nbim.no/en/about-us/about-the-fund

Norwegian Petroleum. (2025). The petroleum tax system. Updated 15 October 2025. Accessed 30 April 2026. norskpetroleum.no/en/economy/petroleum-tax

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

Nussbaum, M. C. (2011). Creating Capabilities: The Human Development Approach. Cambridge, MA: Harvard University Press. https://doi.org/10.4159/harvard.9780674061200

OECD. (2023). Multilateral Convention to Implement Amount A of Pillar One. Paris: OECD Publishing. Released October 2023; not yet open for signature as of early 2026. oecd.org/en/topics/sub-issues/reallocation-of-taxing-rights-to-market-jurisdictions/multilateral-convention-to-implement-amount-a-of-pillar-one

Putnam, R. D. (2000). Bowling Alone: The Collapse and Revival of American Community. New York: Simon & Schuster.

Ricketts, J., & Beaumont, C. D. (2026). Wellbeing, VUCA(t), and BANI: The Unequal Human Burden of Global Disruption. Wellbeing Futures: Innovation, Policy and Research, (2): 61–75. https://doi.org/10.65638/2978-882X.2026.2.06

Sahni, N. R., Carrus, B., & Cutler, D. M. (2021). Administrative simplification and the potential for saving a quarter-trillion dollars in health care. JAMA, 326(17), 1677–1678. https://doi.org/10.1001/jama.2021.17315

Sen, A. (1999). Development as Freedom. Oxford: Oxford University Press.

Sennett, R. (2008). The Craftsman. New Haven: Yale University Press.

Sunstein, C. R. (2022). Sludge: What Stops Us from Getting Things Done and What to Do About It. Cambridge, MA: MIT Press. https://doi.org/10.7551/mitpress/13859.001.0001

Taleb, N. N. (2012). Antifragile: Things That Gain from Disorder. New York: Random House.

Tridico, P., & Lalucq, A. (2025). Taxation of large digital platforms in the light of international developments. Question for oral answer O-000022/2025 to the Commission, on behalf of the Committee on Economic and Monetary Affairs. European Parliament, September 2025. europarl.europa.eu/doceo/document/O-10-2025-000022_EN.html

Wu, T. (2010). The Master Switch: The Rise and Fall of Information Empires. New York: Knopf.

Zuboff, S. (2019). The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. New York: PublicAffairs.

©2026 Ricketts and Beaumont. Published by Wellbeing Futures: Innovation, Policy and Research. This is an open access article licensed under the terms of the Creative Commons Attribution Non-Commercial License which permits unrestricted, non-commercial use, distribution and reproduction in any medium, provided the work is properly cited. (http://creativecommons.org/licenses/by-nc/4.0/)