




A TREATISE

The Agentic Economy

The Convergence of Intelligence and the Economy



Jeremy Allaire · Co-Founder & CEO, Circle

© 2026 Jeremy Allaire. All rights reserved.

The Convergence of Intelligence and the Economy

This is the canonical, long-form statement of a single argument: that the agentic economy and the onchain economy are not neighbors but the same economy, seen from two sides. As artificial intelligence decomposes the firm into agentic skills, those agents become economic actors, and the open, programmable money and settlement rails built over the last decade are what let them transact, coordinate, and be accountable at machine speed.

The work builds its case claim by claim across nine sections, from the decomposition of the firm to the civic vision that follows from it. It is written to be read straight through, but each section also stands on its own. Throughout, it tries to hold its strongest claims and their strongest objections in view at once, naming where the argument is most contestable rather than smoothing it over.

By Jeremy Allaire. July 2026. This PDF is the full text; an interactive edition (with shorter reads, an audiobook, an explainer film, a policymaker's guide, and a concept gallery) lives online at AgenticEconomyTreatise.com. The next page maps those companion formats.

© 2026 Jeremy Allaire. All rights reserved.

DISCLAIMER

The Agentic Economy is a personal work authored by Jeremy Allaire in his individual and personal capacity. Jeremy Allaire is the Chairman, Co-Founder, and CEO of Circle Internet Group, Inc., a company whose current or future products and platform may relate to topics discussed in this work. The views, opinions, and arguments expressed are solely his own. They do not represent the views of Circle Internet Group, Inc. or any of its affiliates, subsidiaries, directors, officers, employees, or investors, and should not be construed as a statement made by or on behalf of Circle. Although Circle is making the link available for convenience, the treatise is not a Circle statement, policy position, product announcement, forecast, or endorsement, and it has not been adopted or approved by Circle Internet Group, Inc., its affiliates, board of directors, or officers.

This work contains forward-looking statements, projections, and opinions about technological, economic, and regulatory developments. These reflect the author's personal assessment and involve inherent uncertainty. Actual developments may differ materially. Nothing in this work constitutes investment, financial, legal, or regulatory advice, and no portion should be relied upon as such.

This is one of many ways in

The full text you are holding is the source of truth. Online, the same argument is also a 60-second thesis, a 10-minute read, an audiobook, a short film, a policy-maker's guide, and a concept map. Visit any of the addresses below.

— Hub

The interactive edition

Every format, one place: start anywhere and choose your depth.

AgenticEconomyTreatise.com/

The treatise online

This full text, section by section, with live cross-links.

AgenticEconomyTreatise.com/treatise/

— Read it shorter

The thesis in 60 seconds

The core bet and the seven claims that carry it.

AgenticEconomyTreatise.com/treatise/thesis.html

The essential read (~10 min)

The whole argument in plain language, no jargon.

AgenticEconomyTreatise.com/treatise/essential.html

The considered read (~15 min)

Each section's core idea developed in turn.

AgenticEconomyTreatise.com/treatise/nine-ideas.html

— Listen & watch

The audiobook (~2 hr 11 min)

The complete work, narrated start to finish.

AgenticEconomyTreatise.com/treatise/listen.html

The explainer film (~8 min)

The whole argument in motion.

AgenticEconomyTreatise.com/treatise/watch.html

— A companion guide

A Policymaker's Guide

A constructive policy agenda drawn from the treatise.

AgenticEconomyTreatise.com/policymakers/

— Explore the ideas

The map

*Every concept as a visual explainer,
mapped and linked back to the text.*

AgenticEconomyTreatise.com/interactive/treatise-mind-map.html


The argument, section by section

01	The Convergence and the Decomposition of the Firm	6
	<i>Why platform shifts arrive as convergences, and how the firm dissolves into agentic skills.</i>	
02	Assembly, Coordination, and Why the Firm Goes Onchain	14
	<i>Reassembling the decomposed firm, and the identity and trust that force it onchain.</i>	
03	The Monetary Substrate: Velocity, Safety, and Finality	22
	<i>Why machine-speed money replaces leverage with velocity, and what keeps it safe.</i>	
04	Credit Markets: Machine Underwriting, Agent Working Capital, and the Prudential Layer	30
	<i>The long tail of credit: real-time underwriting, machine working capital, and its brakes.</i>	
05	Global by Construction	40
	<i>An economy with no single native jurisdiction: compliance, FX, and sovereignty reshaped.</i>	
06	The Supply Side: From Subscription to Consumption	55
	<i>When the seat dies as a pricing axis and the model becomes a cost, the agent a business.</i>	
07	The Onchain Corporation	64
	<i>Two paths to the firm whose governance, treasury, and tokens are native primitives.</i>	
08	Implications and the Concentration of Power	73
	<i>The labor question is the ownership question; concentration is the geopolitics.</i>	
09	The Civic Vision	82
	<i>How a widening economy is designed rather than assumed: ownership as citizenship.</i>	

SECTION 01

The Convergence and the Decomposition of the Firm

Every major platform shift of the internet era arrived the same way. Not as a single invention, but as several technologies that had matured separately and suddenly came together. The web was one such convergence. It needed usable graphical interfaces on commodity hardware, the deregulation that opened the internet to commerce, modems fast enough to make the connection usable, and an open software layer of pages, links, and servers. None of those alone would have made it. Digital media was another convergence, when broadband investment met new media formats and the wireless networks that carried them. Mobile combined cheap cameras and sensors, location, the touch interface, and mobile broadband. The cloud commoditized first software and then hardware. Social platforms were what happened once mobile put a connected device in every hand.




Convergence as a Law

The pattern beneath all of them recurs reliably. When capabilities converge, the marginal cost of some previously expensive activity collapses toward zero. And when that cost collapses, the *velocity* of that activity explodes. The web exploded the velocity of published information. Mobile and social did the same for human communication. The cloud did it for the creation and delivery of software.

When capabilities converge, the marginal cost of some previously expensive activity collapses toward zero. And when that cost collapses, the *velocity* of that activity explodes

Two new operating systems are now converging, and they apply that same mechanism to the two things the internet never managed to digitize natively: intelligence, and economic activity itself. The first is an operating system for intelligence: artificial intelligence, in the form of foundation models and the cloud-hosted agentic systems built on them. The second is an operating system for the economy: blockchain networks, on which value, contracts, and coordination can be expressed and executed in software. Each drives the marginal cost of its domain toward zero. One collapses the cost of cognition and work. The other collapses the cost of transaction, settlement, and coordination. And like every prior shift, the two reinforce each other. Intelligence makes economic activity possible to conduct at machine speed, and an economic substrate lets machine intelligence transact, exchange value, coordinate and execute contracts. This is more than the meeting of two trends. It is a single phenomenon, and the destination of this work is a claim earned gradually across these pages: that the agentic economy and the onchain economy are not neighbors but the same economy, converging in a powerful force that will transform the global economic system.




Two Operating Systems

Start with the operating system for intelligence, because it sets everything else in motion. The term means something specific: the capabilities of frontier foundation models, together with the reasoning and agentic infrastructure that lets them execute work at scale, exemplified today by platforms such as Claude and Claude Code, or OpenAI and Codex. These are a new kind of compute machine. You do not program them in the old sense. You *instruct* them, in language, to produce things and carry out work. The atomic unit of that work is the agent: a reasoning process dispatched against a task. A conventional operating system abstracts hardware and exposes deterministic programs you drive through explicit interfaces. The operating system for intelligence abstracts *cognition*. It is operated through natural-language instruction, it runs non-deterministic reasoning rather than fixed logic, and it takes the agent, not the function call, as its unit of execution. This reframes software itself, from instructions a human wrote and a machine executes literally, to work a human delegates to a reasoning machine. That shift is what makes the next move possible.

a company is, in substance, organized cognition with a logo attached


For the first time, the fundamental tasks of economic production can be broken apart and rebuilt as agentic skills. To see why that matters, look at what a corporation actually is. Beneath the brand and the buildings, a firm is an information system organized around a familiar set of functions: product and engineering, marketing, sales, talent, finance, legal and compliance, operations, customer support. And what it costs to run that system is overwhelmingly *labor*. Across the economy, labor is the single largest operating expense, commonly a quarter to a third of revenue, and far more in service businesses, where it routinely runs from a third to half of costs and, in low-capital operations, well past that. In knowledge and technology firms the point becomes nearly total. The spending that isn't capital is salaries, for engineers, salespeople, marketers, executives, lawyers, and finance staff. Such a company is, in substance, organized cognition with a logo attached. A second great pool sits beyond the payroll: outside professional services, the consultants, lawyers, accountants, agencies, and outsourced specialists that are simply organized labor rented from beyond the firm's walls. Both pools, internal and external, are organized human cognition. That is the exact cost and capability the operating system for intelligence targets.



Decomposition of the Firm

This is why the agentic economy overturns the classical theory of the firm. Economists have long explained why companies exist by pointing to transaction costs: coordinating, contracting with, and trusting *external* labor is expensive, so a firm internalizes the work it is cheaper to


make than to buy. The boundary between what happens inside the company and what is bought outside is drawn by the cost of coordination. When every unit of non-physical work can be performed by a discoverable, contractable, instantly settleable agent, those coordination costs collapse, and the boundary loses its grip. The vivid consequences are the one-person company, where a single human orchestrates a roster of agents to do what once required departments, and the hyper-leveraged small team inside a large firm that executes at a scale its headcount could never support, reusing agentic skills built both inside and outside the organization. The economics compound, because three exponential curves move at once. The labor share of operating cost falls as cognition shifts to agents. The cost of running those agents falls, with the price of equivalent machine intelligence dropping on the order of tenfold a year. And the capability of that intelligence rises across nearly every benchmark at the same time. Cheaper to run, more capable, and absorbing more of the cost base: the three multiply, and the productive potential is enormous.



Overturning the Theory of the Firm

This decomposition is not uniform. It is appearing first in software engineering, because today's models are exceptionally strong at understanding and writing code. It is spreading in parallel across the information-and-coordination-heavy functions, marketing, sales, support, and much of legal, finance, and compliance, wherever the work is the crafting, targeting, analyzing, and presenting of information. The most repeatable, information-dense tasks fall first, but much of

this advances at once rather than in a neat queue. Physical labor is the furthest off. Robotics is delivering augmentation in heavy industry and assembly, but the hard problems of physical production, the physical sciences, and complex logistics remain years away, likely a decade or more, even as AI's own progress begins to bend the robotics curve. The cognitive decomposition is already underway and on a compounding curve, and the reasonable expectation is a sharp acceleration over the next two to three years. The trajectory is exponential, and exponentials are routinely underestimated in their early, deceptively gentle stretch.




Three Exponential Curves

**the agentic economy and the onchain economy are not neighbors
but the same economy, converging in a powerful force that will
transform the global economic system**

It would be a mistake to read decomposition as the simple subtraction of people. The more accurate picture is augmentation as much as replacement. Human creativity and inventiveness are magnified when paired with deep agentic skill. People take on broader, more cross-functional and inter-disciplinary roles, move between focuses far faster than before, and do things that were previously inconceivable for an individual or a small team. Certain capacities remain

irreducibly human: the work of emotional relationships and direct human engagement, the exercise of critical judgment over agentic processes, and the responsibilities of governance and accountability that cannot be delegated to a machine. There is a real tension here. At the level of the individual, this is a story of magnification, of people made far more capable. At the level of the whole economy, it raises a hard question about the share of output that flows to labor at all. Both are true at once. The serious objection is that magnification for individuals can coincide with a falling labor share for the economy, and a later part of this argument reconciles them directly, through ownership. For now it is enough to say that "what stays human" is not a vestige but a transformed and, in important ways, enlarged surface.



What Stays Human

The firm, then, decomposes into intelligence, into agentic skills oriented around the work of creating, delivering, and servicing a product or a service. But a decomposed firm raises a new question: how are all those skills reassembled into coordinated work, and how are the economic relationships among them, inside the company and increasingly beyond it, expressed, enforced, and settled? Decomposition into agents is the beginning of the story. Assembly, coordination, and the economic substrate that makes them transactable are the rest of it. That is where this argument turns next, and it is the path by which the operating system for intelligence and the operating system for the economy are revealed, in the end, to be one and the same.

 CONTINUE ONLINE

Read “The Convergence and the Decomposition of the Firm” online, with live cross-links to related sections, the shorter reads, and the audiobook.

AgenticEconomyTreatise.com/treatise/section-1.html

SECTION 02

Assembly, Coordination, and Why the Firm Goes Onchain


Once the firm decomposes into agentic skills, the question is no longer what can be automated but how the pieces are reassembled into coordinated work. The mechanism is an orchestration layer built on the agentic capabilities of the most advanced foundation models. At its center sits a general orchestrator: an agent that, for any given body of work, breaks an objective into a pipeline of tasks and dispatches them to sub-agents. The surrounding infrastructure initializes the pipeline, maintains its context and memory, and executes and recomposes the results that come back. The same generalized architecture serves any function. A marketing pipeline, a finance pipeline, a product or sales pipeline are, structurally, the same machine pointed at different work. Humans are not absent from it. They occupy two distinct positions. Some sit *in* the loop, doing or reviewing specialized work inside a pipeline where their judgment is required. Others sit *on* the loop, setting the objectives, defining the acceptance criteria, watching the quality of outputs, and deciding when the machine should stop and ask. This distinction between operating inside the work and supervising it from above is the concrete form human oversight takes in an agentic firm. The infrastructure to support it, pipelines, orchestration, and durable memory, is maturing into broad availability now, not at some distant horizon. Many cutting edge firms and teams are in full-on execution mode on these new architectures.

SECTION 2 · ASSEMBLY, COORDINATION, AND WHY THE FIRM GOES ONCHAIN



The Orchestration Layer

A general orchestrator assembles agents into any function — one architecture, many businesses.



Same library, a different subset — a different business.

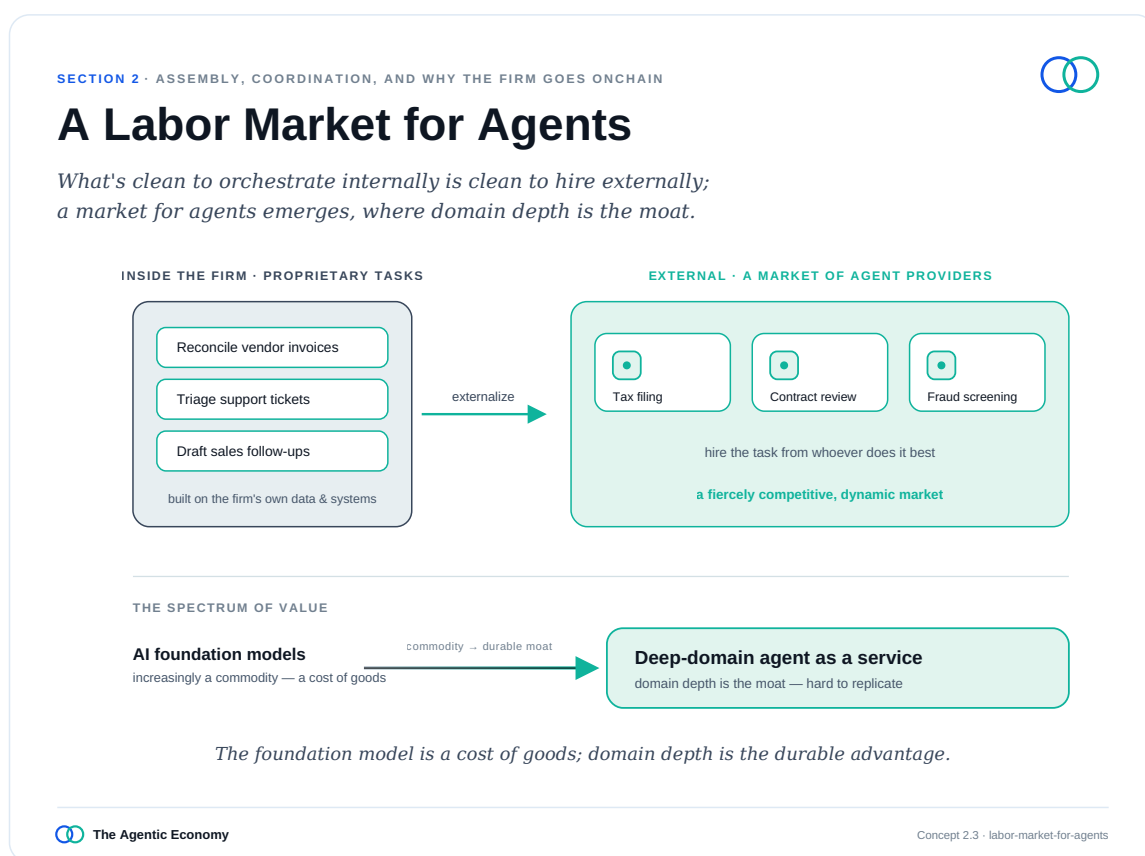
The Orchestration Layer

What begins as an internal exercise does not stay internal, and the reason explains why an open economy of agents emerges rather than merely a set of better-automated companies. To orchestrate its own work, a firm must turn each function into a well-defined skill, trained on its domain, connected to the right data, continually updated. But a skill clean enough to be orchestrated *inside* a company is, by construction, clean enough to be discovered and hired *from outside* it. The moment internal modularization meets agentic marketplaces and agentic payments and contracts, the disaggregation a firm performed to optimize itself becomes the substrate of a cross-organizational market. Companies build the open agentic economy as a byproduct of optimizing themselves. No one has to set out to create it.

a skill clean enough to be orchestrated *inside* a company is, by construction, clean enough to be discovered and hired *from outside* it

That market could take one of two broad shapes, and which it takes matters enormously. In one, companies consume intelligence as a metered utility from a handful of large platforms. In the other, a genuine labor market of specialized agents forms that are *hired* to do particular work. The second is both more likely and more consequential, for the same reason it has held

across software generally: deep domain knowledge is durably valuable. General foundation models and skill libraries supply a great deal, but the lasting business is in agents that go very deep in a domain, creative marketing, video production, intellectual property, contract negotiation, and thousands of other crafts, by aggregating proprietary context and specialized data, refining their capability continuously, and hardening themselves to the security and reliability standards enterprises demand. The foundation model is a commodity input. The durable value sits in domain depth, proprietary data, relentless iteration, and enterprise-grade trust. Specialized agents make their capabilities discoverable through agentic registries and marketplaces, advertised in metadata that is legible to humans and directly consumable by other agents, and they compete fiercely, among the most intense competition anywhere, to deliver stronger results.



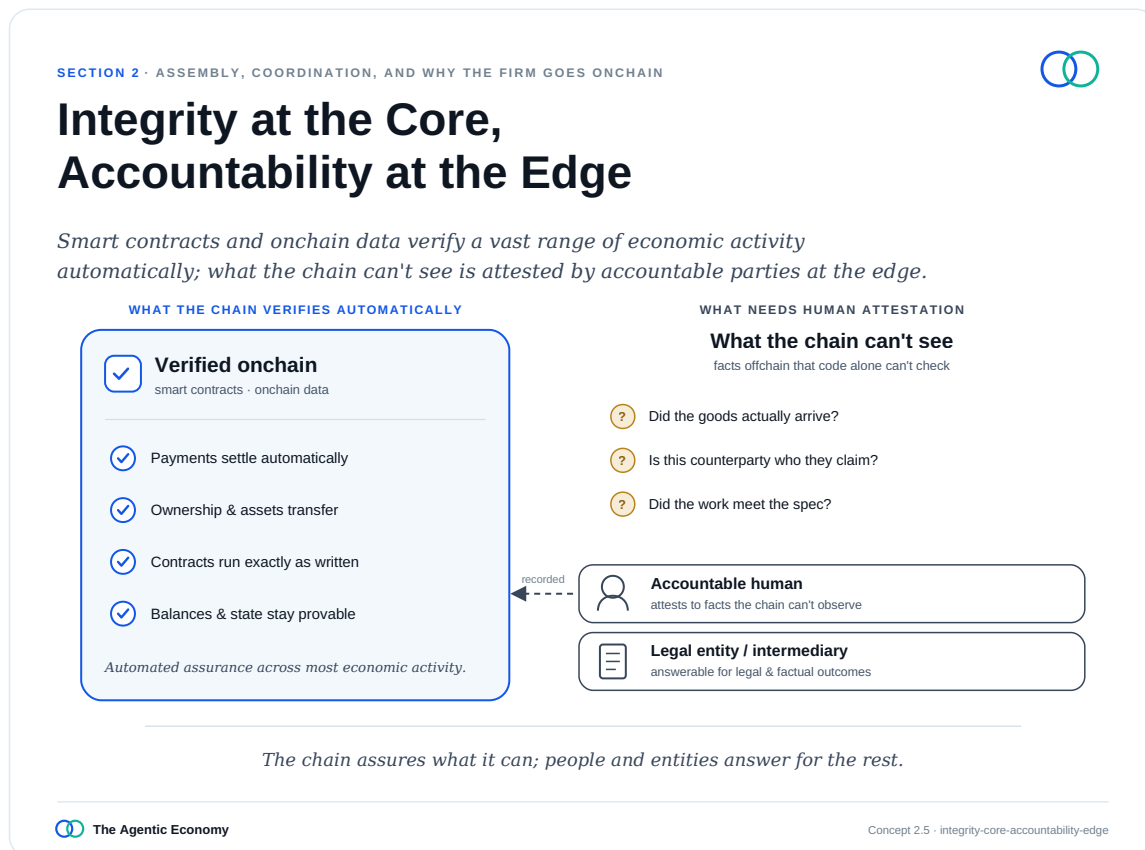
A Labor Market for Agents

The pattern rhymes with what came before. Mobile operating systems let app developers flourish at a scale no one had imagined, and creator platforms let content specialization explode. Here the marketplaces and the tooling to build, discover, and monetize agents are the comparable prize. One economic detail is worth seeding now and developing later: because foundation models compete fiercely, a specialist agent will route work across them to optimize its own cost of intelligence, so the model becomes a cost line and the agent becomes the business, a reversal whose consequences belong to a later section.

A private registry binds trust to a single operator. The cryptographic and real-world-identity roots of an onchain system make trust portable

A labor market for agents runs immediately into a hard problem, and solving it is what forces the whole edifice onchain. Before an orchestrator hires an agent, it must know that the agent is real, that its work can be trusted, and that someone is accountable if things go wrong. None of that is given when the worker is software that may have been assembled anywhere in the world. The answer is that an agent's identity is not one thing but a stack of layers. At the base is cryptographic verifiability: an economic operating system built on a public blockchain whose data, transactions, and code execution are publicly verifiable in real time, rooted in cryptography rather than in trust of any intermediary.

Trust-minimization at the core holds for the things the system can verify for itself: that a transaction occurred, that a balance moved, that a contract executed as written. It does not resolve facts about the outside world, adjudicate disputes, or reverse an outcome that was correct in code but wrong in the world. Those are handled not by the deterministic core but by an accountable periphery: oracles that attest to external facts, adjudication for disputes, and human override for the cases that demand it. The architecture is therefore integrity at the core and accountable intermediation at the periphery, a structure this work develops fully when it takes up the legal form of the firm. The base is trust-minimized. The edges are deliberately, and accountably, not.



Integrity at the Core, Accountability at the Edge


On that base sit the layers that make an agent answerable. The first is real-world grounding. An agent's work must ultimately tie to a real, identified entity, and the regulated identity-verification infrastructure that financial-infrastructure firms already operate at enormous scale, knowing and validating businesses and individuals through established controls, becomes the means of answering, for any agent, who created it, whether they are legitimate, and whether they are in good standing. The second is the agent's own economic existence: a wallet it controls and verifiable credentials that carry its real-world hooks, so that its standing can be depended upon. Atop these sits reputation, built over time the way reputation has always been built on the internet, through the accumulated record of work and the reviews of those who relied on it, except that here it can be far more resistant to fraud, because it is rooted in verified real-world identity rather than in a disposable pseudonym. This is also why the firm goes onchain rather than simply trusting a marketplace's private database. A private registry binds trust to a single operator. The cryptographic and real-world-identity roots of an onchain system make trust *portable*. It travels across marketplaces, across firms, and across borders without requiring anyone to trust a particular platform owner, which is exactly what an open economy of agents transacting globally requires and what no private database can deliver at that scale.

SECTION 2 · ASSEMBLY, COORDINATION, AND WHY THE FIRM GOES ONCHAIN



The Identity & Trust Stack

To transact, an agent needs a cryptographic identity, real-world grounding, a wallet, credentials, and reputation.



No layer stands without the one beneath it.

The Identity & Trust Stack

Autonomy, in this economy, is not anonymity. An autonomous agent is an accountable one


These layers together establish the accountability chain. Every action an agent takes traces, through its wallet and its credentials, back to a verified real-world creator in good standing. Autonomy, in this economy, is not anonymity. An autonomous agent is an accountable one, and the chain that makes it so, agent to wallet to credential to identified, answerable entity, is the thread that recurs throughout this work wherever the question arises of who, in the end, is responsible. It is what lets a counterparty hire a piece of software with confidence, what lets a regulator find a party to hold to account, and what keeps machine autonomy from collapsing into a world of unattributable action.

SECTION 2 · ASSEMBLY, COORDINATION, AND WHY THE FIRM GOES ONCHAIN



The Accountability Chain

Trust becomes portable across contexts, but there is always a party to hold to account.



The Accountability Chain

With these pieces in place, the firm can be seen in its reassembled form. A small human core sits on the loop, setting objectives and exercising judgment. Orchestrators coordinate pipelines of specialized agents, some built in-house and many hired from a global marketplace. Each engagement is an enforceable contract executed in software. And every actor, however autonomous, resolves through the accountability chain to someone who answers for it. At that point coordination has stopped being an internal management problem and become an economic one, conducted in software across the boundaries of the firm, and the agentic corporation has begun, quietly, to reveal itself as the onchain corporation. But all of this presumes something not yet established: a form of money these agents can hold and exchange at machine speed, in enormous volume and tiny increments, without taking on risk in the very act of transacting. That money is the monetary substrate the next section builds.


SECTION 2 · ASSEMBLY, COORDINATION, AND WHY THE FIRM GOES ONCHAIN



In the Loop, On the Loop


Humans move from in the loop — approving each step — to on the loop — supervising the system as it runs.

IN THE LOOP — HUMAN WORKS ALONGSIDE AGENTS



The human works inside the flow — step to step.

ON THE LOOP — SUPERVISES THE SYSTEM



The system runs; the human steps in only when needed.

In the Loop, On the Loop

CONTINUE ONLINE

Read “Assembly, Coordination, and Why the Firm Goes Onchain” online, with live cross-links to related sections, the shorter reads, and the audiobook.

AgenticEconomyTreatise.com/treatise/section-2.html

SECTION 03

The Monetary Substrate: Velocity, Safety, and Finality

The reassembled firm of the previous section needs something it has not yet been given: a money its agents can hold and exchange at machine speed. They must move it in enormous aggregate volume and in vanishingly small increments. And they must do so without assessing, on each transaction, whether the money itself is good. That last qualification is the whole game. The properties that make money usable by software at extreme velocity are not the properties that bank-issued money has. Working out why leads to a specific and old-fashioned answer: full-reserve money, settled with finality, on an open network. What follows below is the case for it.

Begin with velocity, because it reorganizes everything else. The marginal cost of storing and moving money falls toward zero. The time to move it falls to hundreds of milliseconds, effectively the speed of light. And the money becomes directly controllable and programmable by software. The result is a substrate for extraordinarily high-velocity money. The same dollar can be put to work across many uses in rapid succession. Small and large amounts alike become usable the instant they arrive. Micro-level value exchange, orchestrated by agents, becomes feasible where it never was before. This is simply the unit economics that information and software already followed on prior internet-scale platforms, now applied to money itself.

This often raises an objection from monetary economics. The answer is direct. Modern banking manufactures velocity through leverage: a bank lends the same deposited dollar many times over, creating synthetic dollars that carry risk but, in theory, speed money through the economy. If full-reserve money forbids that multiplication, does it starve the economy of credit? It does not. The reuse benefit of leverage can be obtained without permanently creating risky synthetic dollars. When money turns over fast enough, a dollar can be time-locked for seconds and lent to a third party, so velocity itself substitutes for the multiplier. Full-reserve money is not idle either. Its backing in short-term government debt finances government spending into the economy, so the reserves do productive work even at rest. And credit does not vanish. It becomes more powerful. An onchain market for money supports machine-intermediated credit, adjudicated in very short durations, with pooled lending into diversified borrowers while the lender keeps a demand-deposit-like right of instant redemption. That is the substance of the next section. Credit is not starved by full-reserve money. It is rebuilt on top of it, and made more capable and safe.


SECTION 3 · THE MONETARY SUBSTRATE




Velocity Replaces Leverage

Time-locked reuse and machine-speed velocity do the work leverage used to do; reserves aren't idle, and credit is rebuilt.

THE OLD WAY — LEVERAGE



THE NEW WAY — VELOCITY



Same throughput, different mechanism — reserves work without being multiplied into risk.

Velocity Replaces Leverage

Why insist that the base money carry no embedded risk at all? Because velocity makes risk-bearing money dangerous in exact proportion to its speed. A bank run is already faster than it used to be. Mobile banking has shown that a run can overwhelm even a very large institution in hours rather than weeks. At machine speed it becomes effectively instantaneous. An agent deciding whether to accept a unit of money wants assurance of one-for-one redeemability that never has to be questioned. If it must worry about redeemability, it has to price that risk on every transaction. Pricing risk on every micro-transaction at million-fold velocity is economically impossible. There is a subtler failure too. In a world of thousands of distinct bank-money issuers, each issuer's dollar is its own IOU carrying its own risk. A dollar from one issuer is not interchangeable with a dollar from another, and the prices of these supposedly identical dollars diverge. That destroys what monetary economists call the singleness of money: the property that a dollar is a dollar, fungible and at par, which money must have to function as a unit of account at scale.

SECTION 3 · THE MONETARY SUBSTRATE




Why Base Money Carries No Risk

Base money must survive runs at machine speed and can't price risk per micro-payment — so it carries none.

FULLY RESERVED · 1:1



A CENTURY-OLD VERDICT, REVERSED



Money that clears instantly can't also carry risk — and at machine speed, full reserve is finally its moment.


Why Base Money Carries No Risk

Credit is not starved by full-reserve money. It is rebuilt on top of it, and made more capable

The agentic economy is global and runs at internet scale. A far-flung agent making a real-time settlement decision cannot pause to evaluate the creditworthiness of distant issuers. The national backstops that make bank money tolerable (deposit insurance, a lender of last resort) do not reach most actors in a borderless system. Full-reserve money is the only form that holds par for everyone, everywhere, without those backstops. This is the old idea of narrow banking: one-hundred-percent reserves, long proposed and long dismissed as safe but not useful enough to bother with. What changes the verdict is that machine intermediation and internet-scale utility make narrow banking maximally useful for the first time.

Safety of the unit is necessary but not sufficient. The settlement of the unit must be equally beyond question. Sound money has historically rested on a "no questions asked" basis: money works precisely because no one has to perform due diligence on it before accepting it. The financial system formalized this in the Principles for Financial Market Infrastructures, whose core assurance is that when a payment system deems a transaction final, it is final. That assurance

becomes vital at high velocity, and it is exactly where decentralized networks have fallen short. Hard forks could un-recognize previously settled transactions. Chain reorganizations could unwind them. The best such systems offered was *probabilistic* finality: wait some number of confirmations and a reversal becomes unlikely. Agents transacting at extreme velocity cannot operate on "probably final." They require deterministic, sub-second finality: settled means settled, now. That is a concrete technical requirement that a chain engineered for it can meet and older designs cannot. It sits alongside par and redeemability as the third member of a triad. Par, redeemability, and finality together are what let money be accepted by a machine without any per-transaction scrutiny.




Deterministic Finality

Settlement finality raises an apparent paradox, since people also want reversibility: refunds, fraud protection, the ability to undo a mistaken or fraudulent payment. The resolution is architectural. It is the same move the internet made in layering reliable protocols over a simple, unreliable base. Keep the base money deterministically final, and build reversibility as optional protocols *on top* of it: time-adjudicated and event-based escrow, refund pools, insurance against those pools. Do not bake reversibility into the money itself. Reversibility welded into the base unit would destroy the no-questions-asked property and force every agent to price unwind-risk on every transaction, reproducing the failure of IOU money. Pushed to the edges as composable layers, the same protections become available without compromising the core. Ir-

reversibility, in this light, is not a hazard to be mitigated but a feature to be relied upon. A final settlement is a guarantee an agent can build on precisely because it never has to think about it again.

None of this safety is self-enforcing. It rests on institutional architecture, and that architecture is now being built. Under recent law, large stablecoin issuers are overseen by federal national-bank regulators. They are structured with bankruptcy-remoteness, so the money is isolated from the failure of any issuer or bank. National trust-bank charters supply a fiduciary apparatus that walls the base-layer monetary obligations off from credit risk: narrow banking reborn. The reserve design then gets *safer* as the system grows. Eligible reserves are evolving from short-term government debt toward direct central-bank cash and overnight repurchase agreements with the central bank. Policymakers increasingly favor this direction, and it is mirrored internationally: the Bank of England's proposals point the same way, and European payment-systems law already contemplates e-money issuers' access to central-bank balance sheets.




Institutional Architecture

the larger and more systemic the money becomes, the closer its backing moves toward central-bank money itself. The destination is central-bank-grade safety

This is the precise inverse of fractional-reserve banking, where scale concentrated systemic risk. Here, the larger and more systemic the money becomes, the closer its backing moves toward central-bank money itself. The destination is central-bank-grade safety without a central-bank digital currency: private, programmable, internet-native money that is nonetheless bankruptcy-remote, fiduciary-chartered, and ultimately backed by the safest asset there is.

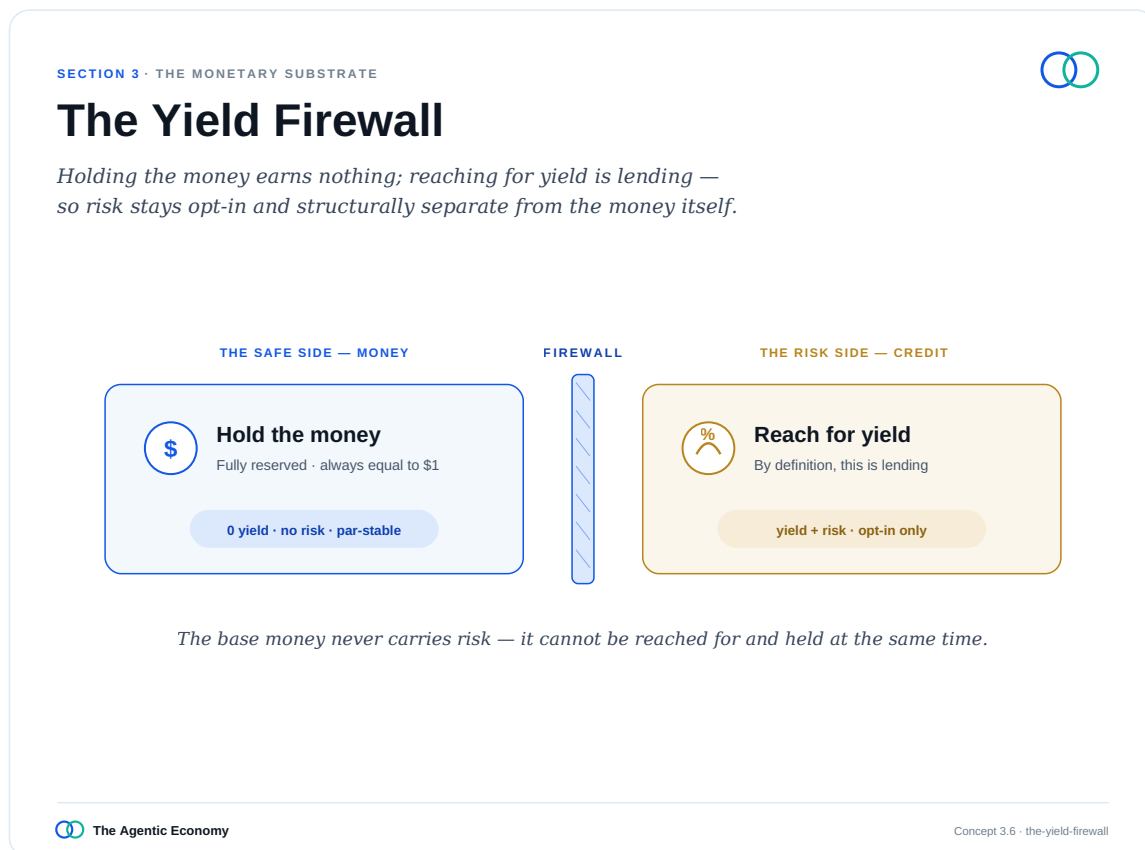
What becomes of monetary policy when the fractional-reserve multiplier is no longer the channel it works through? The price lever remains fully intact. The central bank still sets the policy rate. Because stablecoin reserves *are* short-duration and overnight instruments, that rate transmits to the reserve base of the money supply immediately and fully, arguably a more direct transmission than the bank-lending channel ever provided. What fades is not the central bank's power but the multiplier as its transmission mechanism.



Monetary Policy, Reshaped

In two respects the central bank's hand may actually strengthen. Transmission becomes more direct. And because onchain credit is transparent and observable in real time, the central bank can observe credit conditions directly rather than inferring them from lagged, aggregated reports: visibility it has never had. Its role partly migrates from *operating* the multiplier to *supervising* a transparent machine-credit market: a new frontier of registering and prudentially supervising machine and agentic credit markets, whose participants are agentic in substance but, through the accountability chain established earlier, human-accountable in the end. This is a genuine expansion of the supervisory mandate, not a hollowing-out of it.

One distinction is crucial to understand. The base money described here is ultra-safe precisely because it bears no credit risk and pays its holder no yield for merely holding it. The reserve yield on its backing flows to the issuer and, through it, to the ecosystem of the stablecoin network, but holding the money is not an interest-bearing position. That is deliberate. It is the firewall that keeps the base unit safe. The moment a holder reaches for yield, they are no longer simply holding money. They are *lending* it into a credit market, taking on credit risk in exchange for a credit yield, and that is a distinct and opt-in choice. The safety of the base and the return from putting the base to work as credit are two different things. Conflating them would undo the entire safety argument just built. Holding that firewall firmly in place, the argument can now cross it, into the credit markets that grow on top of this new form of money, and the machine underwriting that makes them reach further than credit has ever reached.



The Yield Firewall

 CONTINUE ONLINE

Read “The Monetary Substrate: Velocity, Safety, and Finality” online, with live cross-links to related sections, the shorter reads, and the audiobook.

AgenticEconomyTreatise.com/treatise/section-3.html

SECTION 04

Credit Markets: Machine Underwriting, Agent Working Capital, and the Prudential Layer

The previous section ended at a firewall. The base money is ultra-safe precisely because it bears no credit risk and pays no yield for being held. The moment a holder reaches for yield, they have crossed over into lending, taking credit risk in exchange for a credit return. This section is about what grows on the far side of that firewall. Credit does not disappear in a full-reserve world. It is rebuilt, in a form that reaches further, prices tighter, and fails more visibly than the credit system it succeeds.


The starting observation reframes the entire problem. The long tail of borrowers (small merchants, gig workers, households, and now agents) has been underserved not because it is bad risk, but because the cost of *evaluating* each small exposure was often too inefficient and even exceeded expected returns on the value of the loan. Credit was rationed by the price of underwriting, not by the quality of the borrower. Collapse the cost of evaluation toward zero, and an enormous population of sound but previously unbankable borrowers becomes serviceable for the first time.

SECTION 4 · CREDIT MARKETS



A Long Tail Rationed by Underwriting

Credit's long tail is rationed by the cost of evaluation, not by borrower quality; collapse that cost and the tail gets served.



Cheap machine underwriting meets a long tail of small suppliers — and the tail clears on both sides.

A Long Tail Rationed by Underwriting


What collapses that cost is a data flywheel. Onchain payment activity is, by its nature, structured, verifiable, and real-time: a far richer substrate for risk models than the lagged, fragmented records underwriting has always relied on. Onchain credit pools allocate risk dynamically using oracles that prove offchain facts into the system: verified data about individuals, households, and firms drawn from existing financial-data rails, credit histories, and hooks into general-ledger and treasury systems. As treasury platforms, fintechs, neobanks, and corporates move their cash onto onchain money, the data grows richer still. Because the networks are global, the enriched data extends globally and ties to every additional currency that comes onchain on the same rails. All of it feeds real-time models with agentic underwriting logic, and the loop compounds. Better data yields better models, which yield better underwriting, which attracts more activity and more data: the same recursive engine that built the long-tail markets of search advertising, content publishing, e-commerce product sales and software distribution, now pointed at credit. The result is, in effect, a real-time, global, entity-permissioned credit-information system, against which today's bureaus look lagged, national, and error-prone.

SECTION 4 · CREDIT MARKETS



The Data Flywheel

Onchain activity plus oracles create a richer underwriting substrate, and underwriting itself becomes tireless agent work.



Each turn enriches the next — the substrate deepens, the cost falls.

The Data Flywheel


A common objection often arrives here: that putting the economy's credit activity onchain means exposing everyone's financial life on a public ledger. It has a clean answer. Onchain does not mean public. Selective-disclosure and confidential-computing techniques let contract state and positions remain encrypted and private by default, exposed only through configurable, cryptographically enforced access policies, while the protocol's rules still execute deterministically on the encrypted data. An entity can prove the attributes that matter to a lender (credit-worthiness, balances, identity in good standing) or have an oracle prove them, without revealing its raw positions to competitors or the public. The flywheel can therefore be rich and private at once, with stronger confidentiality than today's system, where intermediaries see everything. This same capability resolves a tension that recurs throughout this work between comprehensive auditability and privacy: disclosure becomes permissioned and selective rather than all-or-nothing, so a regulator sees what it is authorized to see and a competitor sees nothing.

SECTION 4 · CREDIT MARKETS



Onchain Is Not Public

Onchain can be confidential by default — proving creditworthiness without revealing the underlying data.



Verifiable yet private: the proof travels, the data does not.


Onchain Is Not Public

With evaluation cheap and data abundant, the underwriters themselves become agents, and the economics shift. Agentic underwriters are tireless and continuously optimize toward the efficiency frontier. They compete to underwrite exposures that prior market structures could not reach. They ride the compounding data flywheel. Three forces (a wider range of serviceable opportunities, compounding data advantage, and relentless automated optimization) combine to compress the marginal cost of borrowing, explode the volume of lending, and drive margins down, the same way machine market-making compressed bid-ask spreads in equities.

This runs against an instinct. Normally "cheaper and more abundant credit" is a synonym for "riskier," the reflex of the 2008 financial crisis. Here the system is cheaper, more abundant, *and* safer and more accessible. It achieves this because the new efficiency comes from better information and better underwriting, not from more leverage. This is the credit-side expression of velocity replacing leverage: volume grows on the back of superior underwriting and high turnover, not on the manufacture of risky synthetic dollars. Two consequences travel with this and are taken up later. Compressed margins can tempt capital to reach for yield, which routes to the prudential layer below. And the residual spread accrues to whoever owns the best underwriting agents, which routes to the concentration of capital question of Section 8.

Agent working capital, in plain terms

The core idea is simple. Agents can borrow money to fund their own work. And the work they've been hired to do becomes an asset a lender can finance. Call it agent working capital, and the thing it produces a machine receivable.



Agent Working Capital

Why is this new? When a bank lends to a person, the biggest unknown is whether they'll pay it back. That's a question about human behavior. Machine credit takes that out.

Here's the simple case. An agent already has a contract for a \$10 translation job. It borrows \$4 to buy the extra compute to finish. The lender isn't guessing whether the agent "wants" to repay. It's pricing three concrete things: will the work be accepted, will the oracle report it correctly, will a dispute come up. That's it. Open-ended creditworthiness becomes a short, bounded bet on whether one job gets done.

One caveat, and it travels with everything below. This is near-certain for a single loan. It is not risk-free. Bundle many of these together and you still get correlated, systemic risk. What's new isn't that risk disappears. It's that you can see it in real time and insure against it before a blowup, instead of reconstructing it after the crash.

Collateral flips

Human collateral is some unrelated asset a court seizes slowly. Machine collateral is the opposite. The first thing backing the loan is the payment for the work itself, assigned onchain so the lender is first in line when the job settles. Recovery is automatic, not a lawsuit.

More layers sit underneath: a bond the agent posts that can be slashed, extra collateral, reputation tied to the human who built the agent, and finally that human standing behind it.


If something goes wrong, recovery runs in order. The escrowed payment nets out first. Then the stake is slashed. Then a shared insurance pool absorbs the tail. Only what's left reaches the human. The first three happen automatically, in seconds. The last one works because there's a real, identified person at the end.

This only holds at the short end

Credit was rationed by the price of underwriting, not by the quality of the borrower

The certainty fades as the loan gets longer. A one-minute compute burst against a signed job is nearly mechanical. A few days of working capital adds more risk. Months of financing some unproven capability brings back all the old unknowns, and it's just ordinary credit again.

So machine credit doesn't replace human credit. It becomes a new near-risk-free floor. Human and firm lending then prices as a spread above that floor. And that spread measures exactly what machines remove: enforcement uncertainty, information gaps, and behavioral default. The top of the curve, where repayment rides on a founder's vision, stays human. It always will.



Determinism Decays with Duration

Where does the money come from?

Retail holders lend into these markets the way they already lend into pools today, through the "earn" features in neobanks and exchanges. That grows into "agentic earn": you supply the value, and agents continuously manage yield, risk, and redemption for you. It feels like a checking account with a machine portfolio manager underneath. Corporate treasurers do the same through on-chain cash management. Institutions package it into credit funds. Agents become the main surface all of this flows through.

Two things that shouldn't get blurred.

First, two different yields. The yield on the money's safe backing is not what lenders earn. That sits at the issuer and stablecoin network ecosystem layer. The credit yield is what you earn for lending. It's opt-in, and it carries real risk. Confusing the two would quietly undo the safety the money rests on.

Second, where the value goes. Broad distribution is already how this works. The leading regulated issuers already push most of their reserve income out to the ecosystem, through partners and usage-based rewards, not as interest on the coin. That share is rising. The network tokens emerging here are built as stakeholder instruments, with value flowing to validators, de-

velopers, and users. The result: credit supply itself gets democratized and globalized. The long tail both supplies and consumes credit.

A market that clears in milliseconds, underwritten by machines, can build up hidden exposure and unwind faster than any institution could react. Every financial system can fail. The real question is how it fails: opaquely and late, like today's, or rarely and visibly.

Transparency changes the answer. You can watch the exposure build in real time, every loan, every stake, every link. And you can do it without putting private books on public display. A firm keeps its positions private from competitors, while a supervisor with permissioned access sees the whole system live. That flip, from guessing at risk afterward to watching it form, is the foundation everything else stands on.

But seeing a fire start isn't the same as putting it out. You need brakes written into the protocol, running faster than any committee can meet. Humans set the rules; machines enforce them. The most important brake isn't an off-switch, it's a dial. When too much money crowds into the same model, oracle, or compute provider, it automatically gets more expensive to pile in further. Risk gets steadily pricier instead of slamming into a wall.


Insurance has to be a real layer, not an afterthought: shared pools funded by a small skim on each loan, underwriters above them, reinsurance for the tail. What's new is that premiums price against live, observed risk, not stale averages, and the insurer's own health is continuously verifiable. The last crisis's failed insurer wasn't a problem because it failed. It was a problem because it was opaque, undercapitalized, and nobody could see it. Here, trouble shows up before default.

SECTION 4 · CREDIT MARKETS



The Prudential Layer

Transparency lets risk be observed as it forms; the controls are brakes, not off-switches, with insurance as a first-class layer.



See risk as it forms, slow it before it breaks — brakes, not off-switches.

The Prudential Layer

The money is fully reserved, so the money needs no backstop. There's no leverage to unwind and no run to stop. That's a real break from the banking model. What isn't automatically safe is the credit built on top of the money. Pools can face redemptions; collateral can fire-sale. So the real question is providing liquidity to credit markets under stress, not deposit insurance. The likely answer is private-first: extra collateral, reserve pools, reinsurance, and pre-committed liquidity from large holders. Whether some public backstop should ever reach the most critical infrastructure is genuinely open. But if it's ever needed, transparency makes it faster, smaller, and better targeted than the blind bailouts of the past.

All of this creates new roles: underwriters, insurers, oracle providers, pool operators. Agentic in substance, but each tracing back to a real, answerable human. That's what makes supervision at machine speed possible. The likely structure is two tiers: licensing for the actors big enough to move the whole system, standards and self-regulation for the long tail. Central banks shift from running the old money multiplier to supervising these transparent markets, alongside capital-markets regulators.

The rules aren't finished. Where the perimeter falls, how borderless markets get supervised by national authorities, how to guard against both capture and neglect, none of that is settled. But for the first time, the people responsible for stability would work from a live, verifiable picture

of the system, and could step in gradually instead of bluntly. That's a far more resilient foundation than what we have today.

Which raises the obvious next question. A system this transparent, this global, and this unbound to any one country: where, exactly, does it live? That's where the argument turns next.

 CONTINUE ONLINE

Read “Credit Markets: Machine Underwriting, Agent Working Capital, and the Prudential Layer” online, with live cross-links to related sections, the shorter reads, and the audiobook.

AgenticEconomyTreatise.com/treatise/section-4.html

SECTION 05

Global by Construction

5.1 The Three-Layer Stack

The agentic economy has a concrete architecture, and that architecture has three layers. At the base sits money: software-enabled money, in the form of stablecoins, serving as the unit of account and the medium of final settlement. Above it sits an economic operating system: the layer of coordination, contracting, and value exchange, realized as blockchains and programmable smart contracts with deterministic settlement finality. At the top sits agentic execution: the layer where work is performed, powered by AI foundation models and the cloud software through which they act. The decisive fact about all three is not what they do but where they live. Each is software. Each runs on the internet.

Each layer matters because of what it replaces. Money as software supplants the national banking systems that have intermediated economic life for centuries: national to their core, stitched across borders only through slow and costly correspondent banking. Software-enabled money has no such edges. It is the same money everywhere it is held, settling without asking which country the counterparty inhabits.

The economic operating system replaces the national legal and contract-enforcement systems through which strangers have historically transacted at all. Coordination and trust have always been jurisdiction-bound, because the courts and registries that make a contract mean something are bound to a sovereign. A programmable settlement layer relocates that function into deterministic code, executing identically wherever the parties sit. The trust derives from the protocol, not from a jurisdiction.

Agentic execution replaces the most geographically rooted thing of all: local labor and the firms that organize it. Execution performed by AI models in the cloud has no hometown. It can be summoned from anywhere, by anyone, at any scale, answering to demand rather than location.

This is the organizing insight. Each layer, on its own, has no inherent geography, because each lives in internet software rather than in any national institution. So the economy assembled from them inherits, by default, the borderlessness of the internet. This is global by construction: not a feature added to the system but a structural property of what it is made of. For all of recorded economic history, the substance of economic life has been native to nations and only ever retrofitted onto cross-border ambitions. Now that substance is natively global, and the national framing is the thing that must be retrofitted back on.


The three layers are not a strict bottom-up stack in which each merely rests on the one below. They are interdependent and co-evolving, each accelerating the others. Money makes execution transactable. Execution drives demand for coordination. Coordination makes money useful, giving the unit of account something to denominate. Each layer is itself a convergence of maturing technologies arriving at once (cryptographic settlement, programmable contracts, foundation models, ubiquitous cloud): the signature of a platform shift rather than a single invention.

Once the economy is native to the internet rather than to any nation, a cascade of questions opens that the older framing never had to ask. What does it mean for an economy to have no native jurisdiction? What must money become when it spans every currency at once rather than rests within one? How do sovereignty and compliance reassert themselves against a system that does not recognize borders, and what new forms must they take? And what is the nature of an economy that is at once radically equalizing and powerfully concentrating? The architecture is clear. Its implications are not, and the rest of this section turns to them.



The three-layer stack of the agentic economy: money, the economic operating system, and the agentic execution layer, global by construction.

5.2 No Single Native Jurisdiction



No Single Native Jurisdiction

Every economy that has ever existed has been situated somewhere. Cross-border regulation rests on a premise rarely examined because it was rarely false: an economic act occurs somewhere in particular, performed by a party domiciled somewhere in particular. The act has a *situs*, a legal location, and from that *situs* flows the question of which sovereign's rules apply. The agentic economy is the first economy whose substance has no native *situs* at all. The reason is the work itself: it is executed, yet not necessarily *human* work, performed by software agents whose creators may be spread across a dozen jurisdictions, drawing on models trained in one place, hosted in another, invoked by a counterparty somewhere else. When such an agent negotiates a price or settles a payment, the classical question (*where did this happen?*) has no answer.

It is tempting to call this a jurisdictional vacuum, but the opposite is true. The act without a *situs* does not escape law. It falls under *too much* of it. Modern conflict-of-laws long ago stopped requiring a physical location; mandatory rules attach by effect and by the protected party's position, wherever the actor sits. An agent assembled by contributors in twelve countries is potentially subject to the consumer-protection law of the customer's residence, the data law of the data subject, and the tax claim of the market all at once, and sometimes these

may conflict. The structural fact is not the absence of applicable law but the collision of too many claims to it, with no situs left to break the tie. The economy has no *single* native jurisdiction.

The obsolescence of situs arrives alongside a better coordinate. Every agent is bound to an accountability chain: it acts through credentials and a wallet that trace, through an identity and trust stack, to a real-world creator or entity that has been verified and stands in good standing. Autonomy is not anonymity. Set against the world it succeeds (correspondent banking, nominee structures, offshore vehicles, where the beneficial party is exactly what cannot be seen), this economy is far more legible. Here, every actor resolves to someone accountable.

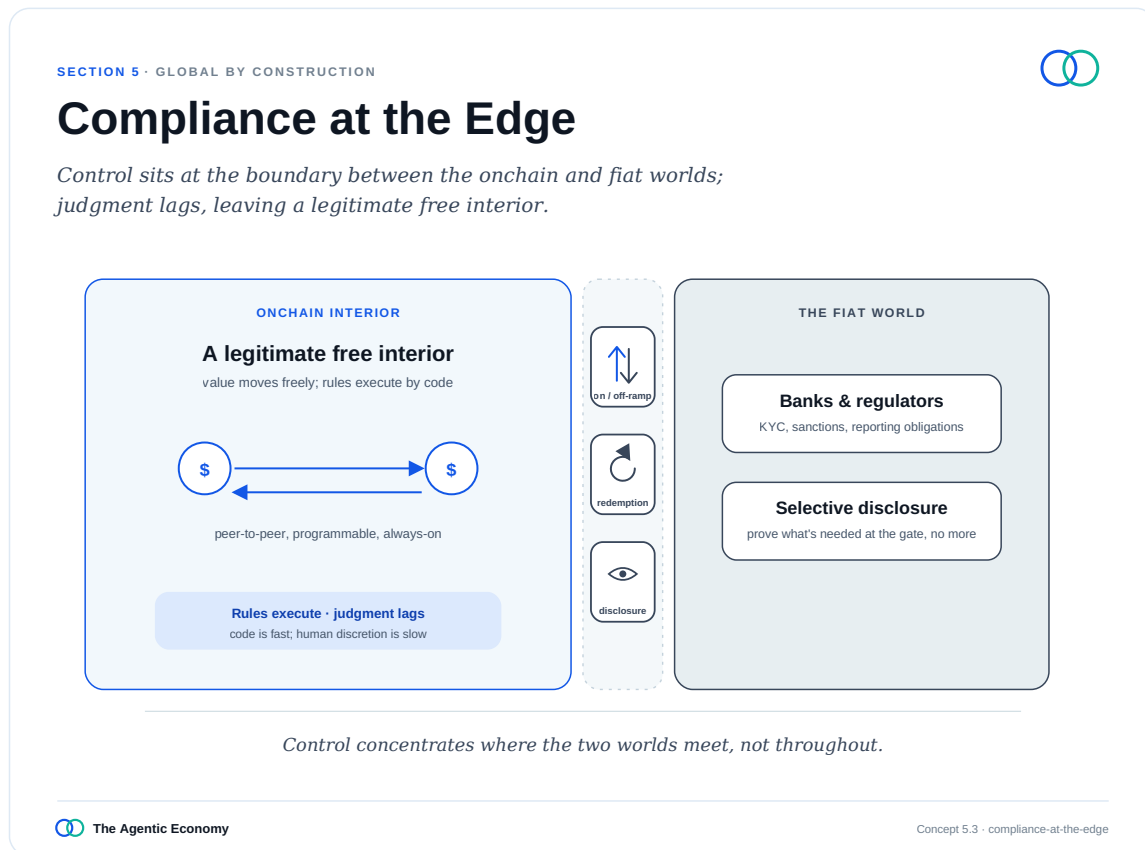
So if the location of an act can no longer organize regulation, the accountable entity behind it can. The question shifts from a territorial one to an entity-based one: *who is the accountable party standing behind this agent, and what obligations does it carry?* But entity-based regulation is not standard-free regulation. If accountability attaches to an entity rather than a place, the entity will choose the place, and "verified in good standing" can decay into "verified somewhere that asks no hard questions."

The discipline must come from the demand side: the jurisdictions where users actually sit conditioning *market access* on the operator meeting a recognized substantive floor. This is how data regulation already reaches the world, and how recent tax accords answered situs-less digital value.

Two limits keep this from being so straightforward. First, attribution is not enforcement. Tracing an act to a verified entity yields a name, not a remedy, and a party domiciled beyond reach can sit behind a perfect identity record. What changes is that, when value moves on programmable rails, enforcement can attach at the level of the infrastructure rather than the courtroom: credentials revoked, balances frozen, access conditioned on compliance. This lever is faster than litigation, and double-edged, because it concentrates enforcement power in whoever operates the infrastructure (a matter taken up later, and one that already has stirred controversy in onchain markets).

Second, the identity layer that makes the economy accountable is, from the other side, a mechanism of control. An accountability engine and a censorship engine can be the same machine. So the identity layer cannot be a single global registry run by a single operator. It must be plural and portable: competing issuers, user-held credentials, selective disclosure so "verified" can be proven without exposing identity to every counterparty, and real due process around revocation. Designed that way, traceability serves accountability without becoming surveillance. The architecture is a choice, and it has to be made deliberately.

5.3 Compliance at the Edge



Compliance at the Edge

Here's the plainer version, straight paragraphs, no sub-headers.

Every system the world has built to fight illicit money rests on a quiet assumption: that money moves slowly, and through few enough chokepoints, to be inspected after the fact. A wire crosses borders through a chain of correspondent banks, each seeing only its own leg. A suspicious-activity report gets filed days later. The old architecture is opaque, fragmented, and backward-looking by design. Its defenders have mistaken that opacity for safety.

The agentic economy makes control stronger at the points that matter most: onboarding, moving money into the system, and moving it across the line of the regulated world. Those are exactly where the old system fails, and it fails on visibility and timing, which is what a transparent, identity-rooted, programmable settlement layer fixes. Screening can be built into the rails themselves, not bolted onto each middleman. It runs as a gate before a transaction settles, not a report after it clears. An agent acts through credentials that trace back to a verified entity, so you can always get from a transaction to a nameable party, which is rarely true across the correspondent maze. And programmable infrastructure gives you levers the old rails never had: revoking credentials, freezing or returning balances, and conditioning access on compliance.

global by construction: not a feature added to the system but a structural property of what it is made of

None of this requires publishing everyone's financial life on an open ledger. The answer is selective disclosure: a system that's private by default and discloses only by consent, granting read access through cryptographically enforced, permissioned rules. Supervisors and law enforcement get authenticated visibility into what they're authorized to see. Competitors and the public see nothing.

But one distinction matters. Programmability guarantees that rules run at machine speed. It does not guarantee that judgment keeps up. Enforcing known rules in real time is genuinely new and powerful. Detecting new kinds of illicit activity in real time is a hard, adversarial problem, and machine speed makes it harder, because agents can probe a ruleset faster than humans can close the gap between the rule and its intent. So laundering isn't solved. Prevention at the boundary gets cheap and fast. Detection on the inside stays an arms race, better-instrumented than today but not won.

This is the crux, and it's a question of architecture, not surveillance. A legitimate financial system has to leave room for real economic freedom and hard privacy: self-custody, unhosted wallets, transfers no operator can see or stop. That isn't a gap to be closed. It's a legitimate condition of a free society, the digital version of cash, and any standard that tried to abolish it would be building an apparatus of total control. The right approach is to put policy at the edge, at the boundary where value and identity cross from the regulated world into the free interior and back. You regulate the on-ramps and off-ramps, not the wallet itself. Illicit value is only useful once it's converted into real purchasing power, and that conversion almost always means crossing back into the transparent world, where issuer redemption is an observable chokepoint that cash never had.


The base layer has to preserve full freedom of use while each jurisdiction builds its own controls at the edges. Build centralized control into the foundation, and you've created a single set of levers that can be captured, coerced, or misused. A neutral base layer is, by design, not something you can sanction. Recent law has started to recognize this, where authorities tried to sanction ownerless protocol code rather than the people using it. Control belongs at the credentialed edge, not in the neutral core.

The hardest capability to defend is also the most powerful. The same lever that returns a thief's takings can enable wrongful seizure, automated error at scale, and state-coerced censorship of lawful but disfavored actors. A freeze that's instant and global, applied by an issuer under pressure with no court involved, is a worse instrument of liberty than the banking freeze it replaces. Freeze and clawback are legitimate only when wrapped in real due process: crypto-

graphically logged, set to expire unless a court renews them, requiring multiple parties to authorize, and paired with a genuine right of appeal.

So the architecture forces a value choice, and it should be made in the open. A bounded private interior means some illicit value will live beyond direct reach, as it always has with cash. What the architecture offers isn't total visibility. It's proportionality. The state gains stronger tools than it has today, observable boundaries, screened on-ramps and off-ramps, and due-process-bound enforcement at the edge, in exchange for giving up the dream of a panopticon over the interior.

5.4 Multi-Currency Money and Invisible FX



Multi-Currency Money & Invisible FX

Foreign exchange is the friction-laden seam between national currencies, and almost everything slow and expensive about moving value across borders lives in that seam. A cross-currency payment threads a chain of correspondent banks, each pre-funding accounts in the other's currency, each taking a margin, each adding a day. The agentic economy dissolves the seam by dissolving its premise: that each currency lives in its own national plumbing and must be handed across the border by intermediaries who straddle two systems.

As every major currency comes onchain as a regulated, full-reserve stablecoin (and the legal scaffolding is being laid market after market), currency becomes an abstraction layer. An entity, or an agent acting for it, holds its own local money. The counterparty receives theirs. The conversion clears underneath, atomically, in a single settlement, at the best rate the market can offer in that moment. The developer, the agent, and the person at the end of it need not think about the exchange at all, in the way an application sending data across the internet never thinks about packets or routes. And the market that does the converting is not one mechanism but a plurality (request-for-quote order books, automated liquidity pools) competing and routing for best execution: a contestable microstructure rather than a monopoly utility.

The settlement property underneath is a real advancement. When the conversion clears atomically, both legs settle together or neither does, which eliminates the cross-currency timing gap: the risk that one currency leaves your account and the other never arrives. What it does not do is eliminate settlement risk altogether; it relocates it. The fiat boundary (funding from local bank money, or redeeming the stablecoin back into it) remains a separate, non-atomic event, and the integrity of any non-dollar stablecoin depends on the depth of its reserves and the liquidity of its redemption. So the agentic FX layer converts the opaque counterparty risk of the correspondent system into transparent, priceable peg-and-redemption risk: a risk one can see and measure, rather than one buried in a chain of foreign balance sheets.


The most consequential promise is reach: the long tail of currencies. A Paraguayan or Kenyan or Filipino currency is costly to transact globally today not because demand is absent but because the operational overhead of serving it through correspondent banking exceeds the value of the flow. Onchain, that fixed cost collapses toward zero, and a currency that was uneconomical to serve becomes serviceable: the same dynamic by which the internet, having driven the cost of distribution to nothing, served the long tail of niche products that physical economics had stranded. One important caveat on this is worth noting: liquidity is not like content. Standing ready to trade an obscure currency means committing real capital to inventory and bearing real risk on it, and that variable cost does not fall to zero. So the tail grows dramatically longer without becoming instantly complete.

In practice, the overwhelming majority of conversions route through the dollar as a bridge. A trade from one small currency to another clears as local-to-dollar and dollar-to-local, not as a native direct pair, because concentrating liquidity in a single vehicle currency is far more efficient than maintaining the impossible number of direct pairs a fully meshed market would require. The system is multi-currency at its endpoints and dollar-concentrated in its plumbing. The experience is genuinely local-to-local even as the value transits a dollar hub in between. The hub does not vanish, and its persistence has consequences for monetary sovereignty that the next part of this section takes up.

The same architecture reshapes how corporations and financial firms manage their treasury. A treasury that lives in onchain money is not a collection of accounts in many banks across many

jurisdictions, painstakingly pooled. It is a single, global, always-on, policy-governed balance. Idle money becomes a legacy artifact: balances sweep into yield or credit continuously, as a standing condition rather than a nightly batch. People set the guardrails as programmable policy, while agents execute beneath. The same last-mile and FX-depth frictions still apply: a global treasury's reach into and out of local money still meets the licensing and liquidity constraints of the legacy edge.

5.5 Monetary Sovereignty, Reshaped



Monetary Sovereignty, Reshaped

Monetary sovereignty is the worry that surfaces fastest once money moves on a neutral, global software layer. If value can settle anywhere in seconds, and one currency dominates cross-border routing, nations, especially smaller ones, would seem to be surrendering control over their monetary affairs. But that intuition rests on a conflation the architecture dissolves: sovereignty is reshaped, not surrendered.

Start by separating two things national banking systems have always fused: the rails on which money moves and the money that moves on them. The protocol layer is multi-stakeholder and jurisdiction-agnostic by design, owned by no single state. The money that travels across it

remains jurisdiction-anchored: a regulated, full-reserve stablecoin is a claim denominated in some sovereign currency, issued under some country's law. The protocol is neutral. The money is not.

That distinction is fundamental. Neutrality of the protocol is not neutrality of the settlement asset. A dollar stablecoin is the regulated liability of entities answerable to one jurisdiction, and it carries, in effect, a foreign off-switch: it can be frozen or placed beyond a holder's reach by a government that is not the holder's own. But protocol neutrality is valuable precisely because it is the *precondition* that lets a nation issue its own currency on the same rails and reduce its dependence on a foreign asset with someone else's off-switch. Neutrality at the base is what makes sovereignty newly possible.

So bringing a nation's own currency onchain, as a regulated full-reserve instrument, is a sovereignty upgrade rather than a concession. A currency reachable only through a slow, costly lattice of correspondent relationships becomes globally programmable and directly usable by anyone. And the lever that actually constitutes monetary sovereignty, the authority to set the price of money in the domestic unit, stays with the central bank; what changes beneath the rate is the transmission plumbing, not the rate-setting power. This redefines sovereignty itself. The old definition was territorial: control the rails, police the borders. The more durable definition is competitive: sovereignty as the capacity to run sound money that earns usage, where currencies compete on credibility rather than enforce loyalty by walls.

An accountability engine and a censorship engine can be the same machine

The upgrade is genuine for currencies with deep markets and credible institutions; it is largely unavailable to the smallest and most fragile, which struggle to attract issuance or liquidity because few wish to hold them. But onchain rails lower the threshold for a small currency to reach usability in ways correspondent banking never allowed. The gap between these sovereign currencies predates these rails; the question is whether they widen it or offer the first affordable path to close it.


The genuinely hard problem is digital dollarization. If citizens of a weak-currency economy can hold a dollar stablecoin as easily as a messaging app, currency substitution becomes frictionless as never before: draining the demand for domestic money, disintermediating local banks, and weakening a central bank that can print the local unit but not the dollars its citizens now hold. The architecture imposes discipline against this: when sounder money is a frictionless choice, unsound policy carries an immediate, visible cost. But the discipline is asymmetric. It falls on the weak currency, never on the issuer of the dominant one, which faces no comparable sanc-

tion precisely because its money is the asset everyone flees toward. That exorbitant privilege is real; it predates onchain money by generations and lives in the network effect, not the protocol.

Dollarization must be managed by "policy at the edge," not wished away: capital-flow measures, holding limits, and conversion rules expressed in code and enforced at the on-ramps and off-ramps. These are partial (the more permissionless the interior, the leakier the edge) but sharper and more observable than legacy models of enforcement. The threat is ceding the field to a foreign currency, which is the argument for issuing one's own local currency stablecoin early.

A world of sound onchain currencies, deep liquidity, and effective boundary tools may be stable and even welfare-improving, but the history of capital-account liberalization teaches that the transition is where economies break, when substitution outruns the institutions and buffers that would make it safe. And these rails are arriving whether or not any single sovereign chooses them. So the choice is not between transition and no transition, but between one that is managed and sovereign-led and one that is forced and unmanaged. Sovereignty in the agentic economy is reshaped, and defensible only by those who run money worth holding.

5.6 Legacy Interoperability and Migration



Legacy Interoperability & Migration

The agentic economy will not be born into a clean field. It arrives atop a planetary installed base of payment infrastructure (bank transfers, card networks, wires, e-money) that has run the world's money for decades and will not be swept away on any near horizon. This will not be a world of overnight displacement, a story this industry has told before and been wrong about. It is a new settlement *substrate* forming underneath the new value the agentic economy creates, while the legacy rails keep running the value that already exists, joined by bridges, with a boundary that moves slowly inward.

Onchain becomes the *core* (the native place where value is born and settled) for the flows that are new or that the old rails were never built to carry: cross-border value, programmable value, value that moves continuously around the clock, and above all value exchanged machine-to-machine by software agents in volumes and increments no human-paced rail was designed for.

Legacy becomes the *edge*: the last mile through which onchain value exits to reach endpoints that have not migrated. It is also worth noting that for much domestic, low-friction activity the incumbents are not standing still, as real-time account-to-account bank systems now deliver instant, cheap, final domestic settlement. The claim is not that those rails die. It is that the frontier of new value formation is onchain-native, while the legacy world keeps the installed base for years.

The two worlds are joined by bridges, and bridges must be described without illusion. A bridge is an intermediary, precisely the kind of trusted party the onchain model was meant to minimize. The point is that the trust is *relocated and reduced*, not removed. This is not the fragile, trustless cross-chain token bridge that produced crypto's worst losses; it is closer to a regulated clearing utility: licensed, capitalized, with conventional liability and resolution. The bridge is a systemic node, and its governance and resolution have to be engineered as deliberately as any clearing house.

The hardest part is the last mile, because cashing out to local money is where this industry has historically been weakest. A stablecoin balance is far more mobile than pre-funded accounts scattered across correspondents, but the endpoints remain a grind: a payout in a given country still depends on a banking partner, a license, and deep local liquidity at the moment of conversion, scarcest in precisely the emerging-market corridors where the cross-border case is largest.


Onchain does not, by itself, dissolve the correspondent-banking problem at the endpoints; it relocates it there. What a unified network changes is that the licensing-and-partnership burden becomes *aggregatable*: one network amortizing relationships, licenses, and local liquidity across many jurisdictions, rather than every firm building bilateral ties one country at a time.

The consumer and merchant surface also embraces and extends into the onchain core. Tokenized card credentials let value funded onchain reach existing merchant acceptance, but they

still ride the incumbent card rails, with their scheme rules and fee economics intact. The distinction is between *settlement-layer* migration, which is likely, and *acceptance-layer* migration, which is contested and may never fully happen. A durable equilibrium in which stablecoins merely fund the same card networks, with the incumbents keeping the acceptance surface and the rents, is a genuine possibility, though as the agentic economy unfolds and more services are fulfilled using onchain money, it seems clear that the ultimately acceptance-layer rails will succumb to the new onchain environment.

What ties the migration together is that agentic commerce is *net-new demand*: value exchanged by software agents that no incumbent rail was built to serve, riding the secular growth of AI rather than depending on the zero-sum conversion of incumbents. Onchain need not win an argument with the card networks to become the substrate of the agentic economy. It need only be the place that the agentic economy can natively run. Coexistence will continue for a long time; but the substrate is being laid even as the old world keeps running on top of it, and the boundary between them moves, steadily, in one direction.

5.7 The Equalizing Dimension and Its Counterweight



The Equalizing Dimension & Its Counterweight

Every property that makes the agentic economy powerful cuts both ways. Being borderless opens the whole world to a creator in a small market, and opens that small market to the whole world. Permissionless access removes the gatekeeper who kept out the unbanked, and removes the gatekeeper who protected the local incumbent. Zero marginal cost lets anyone serve everyone, and lets the best-funded player serve everyone first. The design is two-sided: radically equalizing and powerfully concentrating at the same time, from the very same features. Section 8 takes this up in full.

The equalizing case is real, and in places absolute. Earlier waves of digital opportunity were egalitarian in rhetoric but gated in practice. Distribution needed an app-store reviewer's approval. Payments needed a bank. Capital markets needed accreditation. And the global financial system ran on correspondent banking that left much of the world out. An on-chain agentic economy can have no gatekeeper at the protocol level. No one's permission is needed to hold sound money, to transact, to supply or consume credit, or to sell agent-produced work into a global market. For the genuinely excluded, that's the difference between zero and one. And access to a stable store of value, independent of a failing local currency, matters most to the people a high-inflation economy has served least.

The counterweight is just as real. Removing the friction of borders, currencies, and local licensing works in both directions. It isn't a one-way gift to the periphery. The world reaches the small market as easily as the small market reaches the world, and it arrives better-funded, with better agents and lower costs. Local players used to survive on the protection created by distance, language, currency, and regulation. Strip that away, and they face winner-take-all competition they're set up to lose. And winner-take-all is the baseline here, not just a risk. Money is the purest network-effect good there is, so one dominant settlement asset tends to become the default. And frontier AI capability is gated by capital at a scale the open web never required. The open web ended with a handful of giants. The same gravity operates here.

The two outcomes aren't evenly matched. Concentration is the default: it's what these features produce on their own, reinforced by the deepest patterns of platform economics. Equalization is the buildable alternative. It only happens if specific things are built and chosen: open infrastructure at the core layers that would otherwise collect tolls, ownership distributed by design rather than left to capital's pull, and policy that stops gatekeepers from re-forming at the new chokepoints, namely the model, the dominant issuer, the identity layer, and the bridge. Equal access is not the same as equal outcomes. That a creator anywhere can reach the global market says nothing about who captures the profit. The honest claim is that the floor rises, not that the gap closes. The outcome is a choice, made through ownership structure and policy, and for the first time the technology has put the better outcome within reach of deliberate design. The global layer doesn't settle this question. It sharpens it and hands it to Section 8, where the choice actually gets made.

An economy that's global by construction can't be only a technical or economic fact. It's unavoidably a geopolitical one, and three things follow. First, the base infrastructure has to be technologically neutral and governed by a wide set of stakeholders, not owned by any one state, because a planetary economic operating system controlled by a single power would be neither trusted nor stable. Second, when a major government enshrines its currency as regulated digital money, the consequences are profound and lasting. It accelerates the migration of value onto these rails and forces every other government to compete, adapt, or resist. Third, economic velocity raises the cost of conflict. A world whose commerce is woven together at machine speed has, built into its very plumbing, a reason for its participants not to tear it apart. That's the most hopeful of the three implications, and the least certain. All of these carry forward to Section 8.

 CONTINUE ONLINE

Read “Global by Construction” online, with live cross-links to related sections, the shorter reads, and the audiobook.


AgenticEconomyTreatise.com/treatise/section-5.html

SECTION 06

The Supply Side: From Subscription to Consumption

The agentic economy needs a supply side. It needs a universe of services that agents can call, hire, and pay. That supply side forms in two waves. The first wraps what already exists. Software, data, and web services expose themselves to agents through command-line interfaces, machine-readable wrappers, and skills. They are packaged so an agent rather than a human can invoke them, and increasingly metered and priced for that machine consumer. The second wave builds what did not exist before: purpose-built specialist agents that go deep in a domain and sell their work into the market. How that market organizes itself (discovery, identity, the labor market for agents-as-a-service) is the subject of an earlier part of this work. The concern here is the economic model, where the deepest break with the past occurs. That break is simple. The unit of value shifts from access to work, and that single shift reprices the software industry.

For three decades the dominant business model of software has been the subscription priced per seat: a recurring fee for *access* to a tool a human sits in front of. The agentic economy dissolves the seat. The consumer is no longer a person occupying a license but an agent performing a task, and what gets bought is not access but the work itself. What dies is the *seat as the axis of pricing*, not subscriptions as such. The meter moves from "how many people may log in" to "how much work was done."




The Seat Dies

Around that new unit, commercial structure re-forms in several shapes at once. Pure consumption pricing suits spiky and exploratory work. Subscriptions with allocated usage (a committed spend that includes a budget of work units, with overage beyond it) restore the budget predictability that finance and procurement demand. Outcome-based pricing applies where the outcome can be cleanly defined and measured. These are not competing predictions but a pluralistic equilibrium. The forces usually cited against consumption pricing (the preference for forecastable costs, the difficulty of attributing "work delivered," the volatility that capital markets penalize in pure-usage revenue) are precisely the forces that produce that pluralism. The seat dies. The wrapper that restores predictability around the new unit of work survives in its place.

The same logic runs one layer down, into the relationship between the agents that do the work and the foundation models that power them. This is where value migrates. As specialist agents proliferate, the buyer increasingly purchases an outcome from an agent rather than tokens from a model. The agent absorbs the cost of intelligence as a cost of goods, arbitraging across competing foundation models to render the work as cheaply as quality allows. This is no longer theoretical. The machinery already exists and is maturing fast. Model routers (a layer that directs each request to the most appropriate model on cost, latency, and quality) have moved in a single year from optional tooling to mission-critical infrastructure. A large share of enterprises now run several models in production, and routing is reported to cut costs by substantial

margins while preserving quality. The incentive to route is enormous because the price spread across the model menu is enormous. The cheapest production models cost a tiny fraction of a cent per million tokens. The most capable frontier models cost orders of magnitude more. The spread is wide enough that sending a simple task to an expensive model is simply waste.



The Model Is the Cost, the Agent Is the Business

The unit of value shifts from access to work, and that single shift reprices the software industry

Around this has grown an entire discipline of cost governance, with a new vocabulary to match. "Tokenmaxxing" names the failure mode of optimizing for tokens consumed rather than value produced. The controls follow, imposed by large enterprises as their internal AI bills ran toward the billions. The unit economics of agentic work are not a distant abstraction. They are being forged right now, quickly, as the field moves from early experimentation toward the mainstream.

What that forging implies is captured in a phrase: the model becomes the cost line, and the agent becomes the business. Value tends to accrue not to the raw intelligence but to the layer

that owns the customer relationship, the proprietary context, the workflow, and the accountability for the outcome. This is the same pattern by which prior platform waves saw a commodity input sit beneath a value-capturing layer above it.

However, there is a serious counter-case. The owners of frontier models are not passive commodities-in-waiting. They retain real pricing power on the capability frontier (the genuinely hard fraction of tasks where one model is materially better and routing collapses to "use the best") and they integrate upward into the agent and application layers, owning the model and the customer at once. The likely equilibrium is a barbell: a large commoditizable middle of work, arbitrated across interchangeable models by specialist agents who capture the margin there, and a frontier tail where model owners keep their rents and compete directly at the agent layer. The line between the two moves upward as models improve. Yesterday's frontier becomes today's commodity, continually feeding the arbitrageable middle even as a new frontier opens above.

Beneath the pricing of work sits the settlement of it, and here an old dream finally comes true. The micropayment, the sub-cent charge for a small unit of value, was promised for the consumer internet for thirty years and never arrived. The usual explanation is that settlement was too expensive. That was only half the reason. The deeper killer was the human mental-transaction cost. People hate having to decide, over and over, whether a thing is worth a penny, and the cognitive friction of the decision dwarfed the price.

the model becomes the cost line, and the agent becomes the business

The agentic economy removes both halves at once. Settlement on programmable rails clears fractions of a cent cheaply, and the consumer is now a machine that has no instinct to deliberate over a penny. So micropayments arrive at last. Not for content access, the use case that always failed, but for *labor*: the metered exchange of small units of work between agents. The friction moves rather than vanishes. The "is this worth it" decision does not disappear; it is re-encoded in software as a budget policy and a value-of-work estimate, and metering millions of sub-cent events carries its own accounting and observability overhead. But that cost is amortizable (one policy governs millions of decisions) and aggregatable, since the micro-events can be netted and batched rather than each settled individually. The per-event cost falls far without reaching zero. The dominant historical barrier collapses, and a new, far lower, amortizable cost takes its place, and is what makes labor at the granularity of a penny finally economic.


SECTION 6 · THE SUPPLY SIDE — FROM SUBSCRIPTION TO CONSUMPTION



Micropayments Arrive — for Labor

Micropayments finally arrive — for agent labor, not human content; friction moves rather than vanishes.

LABOR, NOT CONTENT



FRICTION RELOCATES — IT DOES NOT VANISH



The cost of paying moves out of the moment of payment.

Micropayments Arrive, for Labor

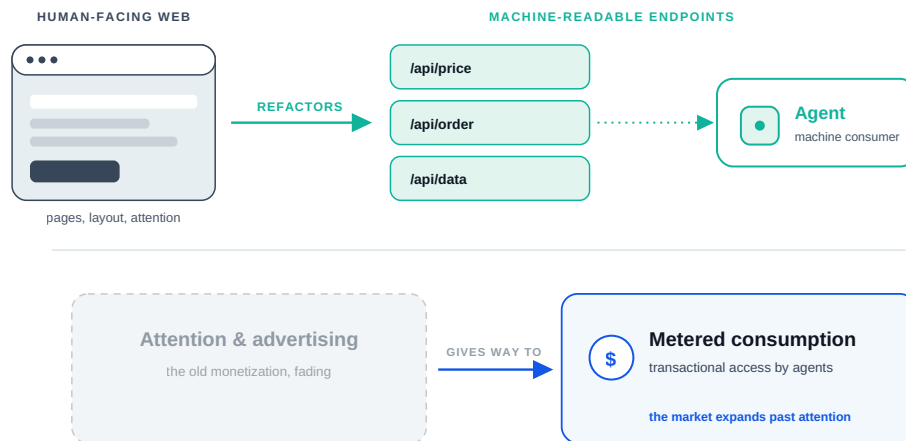
When the consumer of a service is a machine rather than a human with finite attention, the services themselves begin to change shape. A growing class of them, including consumer web services, will refactor to assume that their primary user is an agent: machine-callable, priced per call, stripped of the interface, the funnel, and the advertising that existed to capture and monetize human attention. The addressable market for such a service expands enormously, because an agent can consume thousands of services in parallel, around the clock, freed from the attention bottleneck that capped human consumption.

SECTION 6 · THE SUPPLY SIDE — FROM SUBSCRIPTION TO CONSUMPTION



The Internet Refractors for Agents

The web rebuilds for a machine consumer; metered consumption becomes the successor to advertising.



The Internet Refractors for Agents


The refactor will be partial and gradual before it is complete. Incumbents will bolt an agent-readable layer onto their existing human interface long before they tear the interface out. Many will resist, walling agents out or admitting only their own, because the present internet is a multi-hundred-billion-dollar two-sided market they have every incentive to defend. And it opens a genuine hole. Attention was the very thing the free, ad-funded web sold. A machine consumer with no attention and no susceptibility to advertising removes the revenue model that funded much of the internet. The hole and its repair are the same event.

When agents disintermediate attention-based monetization, the natural successor is direct, metered payment for the service the agent consumes: precisely the consumption-and-micro-payment model the rest of this section describes. Agentic consumption pricing is the successor revenue model to attention and advertising.

micropayments arrive at last. Not for content access, the use case that always failed, but for *labor*: the metered exchange of small units of work between agents

None of this works without a layer; the optimistic story tends to leave out, and recent experience has already shown what its absence costs. If work is metered and agents can hire other agents, tools, and models in turn, the surface over which money can be spent becomes effectively unbounded. The agent running up the cost isn't the one who has to pay it. A machine has no instinct to flinch at spending. Caught in a bad loop or an over-eager plan, it can run a very large bill very fast. Big companies have already watched internal AI usage blow past budget, and responded with hard per-tool caps and centralized controls.


So the agentic economy needs a governance layer just to function at all: spend caps and budget rules, allowances that pass down the chain of delegated authority, real-time metering, anomaly detection, and human sign-off for consequential actions. Think of it as the agentic version of cloud cost controls and the corporate-card limit. It becomes a product category in its own right. Far from undermining the thesis, it completes it. A priced, machine-callable service has to expose policy, identity, and authorization too. The micropayment rail is incomplete without a budget-authority rail beside it, anchored in the same accountable identity described throughout this work.



The Governance Layer

One more point, about pace, because two clocks are easy to confuse. The cost-and-capability clock is genuinely fast. Token costs have fallen by orders of magnitude, cost-control techniques

are maturing in months, and the leading edge (software development above all, and other digital-native work) is going mainstream quickly, because it has high trust, easy integration, and measurable results. The enterprise-adoption clock for consequential and regulated work is slower, and mostly runs independently of the first. It's gated not by unit economics but by integration into systems of record, security and compliance review, the immaturity of the governance layer just described, the question of who's liable when an agent errs, and procurement cycles that themselves resist the new pricing.



Two Clocks

Both are true at once. The direction is fast and the enabling economics are arriving fast, while broad adoption across consequential domains is a multi-year process, gated unevenly. But the direction isn't in doubt. The unit of value shifts from access to work. Value moves to the agent that owns the outcome, while intelligence itself becomes a commodity across most tasks. Micropayments finally arrive for labor, because the buyer is finally a machine. The internet starts to reshape around its new agentic customer. And a governance layer rises to make autonomous spending safe. It's a pluralistic, governed regime, fast-arriving but unevenly gated: not a clean, sudden flip, but a real repricing of how the economy's work gets bought and sold.

 CONTINUE ONLINE

Read “The Supply Side: From Subscription to Consumption” online, with live cross-links to related sections, the shorter reads, and the audiobook.

AgenticEconomyTreatise.com/treatise/section-6.html

SECTION 07

The Onchain Corporation

As artificial intelligence takes on more of the work of the firm, the firm itself needs somewhere new to live. A corporation whose labor is increasingly performed by software agents (agents that hold value, enter contracts, coordinate with one another, and act continuously across the world) requires an economic substrate on which all of that can actually happen. It needs a substrate where money is held and moved programmatically, where decisions and rules of execution are expressed in software, where coordination among agents and humans is recorded and enforced, and where the firm's external economic relationships are transacted at machine speed. That substrate is the onchain economy.

The agentic corporation and the onchain corporation are the same entity seen from two sides. The agentic side describes who does the work. The onchain side describes the form that work takes. This is the keystone of the entire treatise: the agentic economy is the onchain economy. The two are not adjacent trends that may one day intersect. They are one phenomenon, because an economy run by software agents must run on software money, software contracts, and software governance, or it cannot run at all.


This does not mean every corporation dissolves into a token-governed collective, and that distinction matters more than any other in this section. The future of the firm is hybrid, and two paths run in parallel.

SECTION 7 · THE ONCHAIN CORPORATION



Two Parallel Paths

*Existing companies tokenize and map governance onchain;
agentic-native firms build it as core primitives.*



Two Parallel Paths

On the evolutionary path, existing corporations, including ordinary Delaware C-corporations, begin to tokenize their equity and to map their governance mechanisms onto onchain ones. They gradually relocate the operational substance of the firm to programmable infrastructure while keeping their familiar legal form. This is already underway: in the shifting posture of securities regulators, in share registrars and transfer agents piloting onchain registers, in the firms assembling the connective infrastructure, and in corporate-law amendments that permit a distributed ledger to serve as the stock ledger itself.

But it is a long arc, gated by the slowest-moving incumbents in finance. The statutes that define the authoritative register of ownership, the entrenched central securities depository through which public shares settle, the unresolved tax mechanics, the audit standards, and the rational caution of boards and their counsel all hold it back. This conversion will unfold over a decade or two, not a single cycle.

The agentic corporation and the onchain corporation are the same entity seen from two sides. The agentic side describes who does the work. The onchain side describes the form that work takes

On the De Novo path, strongly agentic new firms build native from the start. They treat governance, treasury, and digital tokens not as features bolted onto a conventional company but as the core primitives through which the firm is operationalized from its first day. The two paths mirror what is happening on the supply side of software itself. Incumbents evolve their existing products into agent-consumable services, while new entrants build agent-native from scratch. The native builders, unburdened by legacy, demonstrate the model and pull the incumbents and their regulators forward.

Even the De Novo firm does not escape the law by being born in software. Legal personhood, limited liability, and the standing to contract, sue, and be sued are conferred by a sovereign, not by a deployment transaction. A system of smart contracts that has not been domiciled in a recognizing jurisdiction is, by legal default, an unincorporated association or general partnership. Its participants can bear unlimited personal liability, as a token-governed collective discovered when a court held exactly that. So the working De Novo firms incorporate *through* a thin statutory wrapper: the purpose-built DAO LLC and related vehicles that bridge code to an existing corporate-law container.




Even De Novo Needs a Shell

The novelty is not that the firm exists outside the law. It is that the ratio inverts. The legal shell becomes thin and near-formal while the operational substance (treasury, payroll, contracting,

governance execution) lives onchain and thick. That is a genuine transformation in how a company is run. The sovereign's role in conferring its existence does not disappear.

The substance that fills these new forms is itself novel. The onchain corporation combines traditional governance (an incorporated, limited-liability container) with tokenized ownership: equity or digital tokens that convey contractual rights to revenue and income, to voting and participation, and to other forms of utility in the entity. It replaces non-digital infrastructure with programmable infrastructure: a treasury managed onchain under explicit policy, auditability that is continuous rather than periodic, and interoperability that is global by default.




Tokenized Ownership & Infrastructure

Three things must be kept in mind with this evolution. A token can *represent* a share today; that is live and real. Whether the chain is the *record of authority* (the legally operative register of who owns what) is an emerging, statute-dependent matter, advancing only as fast as corporate law makes the ledger the official register. And settlement with true legal finality of title, as opposed to finality of a token that mirrors a share, remains largely aspirational. Until the chain becomes the legal register, a tokenized share lives in two books at once. The unglamorous machinery (proxy voting and disclosure, dividends and tax withholding, lockups and holding periods, accredited-versus-retail transfer restrictions, regulated custody) is precisely where the work is hardest and the progress thinnest. Representation is present. Record-of-authority is

the active frontier. Full legal-settlement finality is the destination. The three do not collapse into one.

The ledger is a better witness, not a better fiduciary

The deepest change is in governance: as decisions within a firm come to be made by both humans and agents, the enterprise needs a single tamper-proof record that both can read, write, and act on with equal authority. Without it, governance fractures along the seam between human and machine. The cryptographically provable canonical ledger is that shared source of truth, the backbone that runs from the execution of an agent's action outward to token and equity holders and to the perimeter of the board itself.




Governance Across the Seam

But the ledger's power has to be described for what it is. It establishes, permanently and verifiably, what happened, in what order, and by whom. That's a huge gain in provenance and non-repudiation, and it settles whole categories of factual dispute. What it doesn't establish is that an action was authorized, within the entity's powers, prudent, or loyal. A perfectly provable

record of a self-dealing or unauthorized transaction is still a breach, and permanence can even lock a wrongful act in place. The ledger is a better witness, not a better fiduciary.

Fiduciary duty is a duty of persons. A machine can't hold it, can't be questioned on its intent, can't be disqualified for disloyalty, can't be personally liable. So when an agent makes or carries out a governance decision, the duty doesn't disappear. It lands on the humans who designed, configured, authorized, and were supposed to supervise the agent. Machine governance doesn't move accountability off humans. It sharpens it onto the designers and overseers, and it adds a duty to supervise the agent on top of the existing duty of care. "Humans in the governance loop" then becomes precise: humans hold the duty, the ledger holds the record, and agents act within bounded authority that humans delegated.

The same precision rescues the section's most evocative claim, that the contract becomes the program. A contract, at bottom, is a set of decisions and rules for carrying them out. Increasingly those rules are written in code and enforced by machines, between firms and agents and between agents themselves. In that sense the contract really does become the program. But code executes literally and completely, while legal contracts are deliberately incomplete. They rely on courts to read intent over literal text, and to apply the doctrines that make commerce workable: mistake, fraud, duress, impossibility, and the mandatory consumer, employment, and competition law that parties can't contract around.



The Contract Becomes the Program

The contract becomes the program in how it's performed. It stays a legal instrument in how it's governed

So the picture is two-layered. The program is the performance and settlement layer. It executes the agreement automatically on the high-volume, unambiguous path, which is most transactions by count. The legal contract stays the governing instrument for interpretation, defenses, and mandatory law on the rare path where code and intent come apart. The program automates performance, not meaning. The classic example is now a decade old. When a famous early smart-contract system was exploited, the code did exactly what it was written to do, and the community still treated the outcome as wrong and reversed it. Even the most code-native actors put the real judgment outside the bytecode when execution and intent diverge. The contract becomes the program in how it's performed. It stays a legal instrument in how it's governed.


This is why the onchain firm is best understood as integrity at the core and adjudication at the edges. The deterministic core handles the huge volume of low-ambiguity activity (transfers, payments, vesting, simple conditional logic) with an automaticity and auditability no human back office can match. For that majority of cases, trust really is minimized. The edge handles the contested minority: oracles that bring in offchain facts, adjudication that resolves disputes, and human override that can step in when the core produces a result that's right in the code but wrong in the world. That edge is re-intermediated, not dis-intermediated. An oracle is a trusted party for the fact it reports. An adjudicator has to be run by someone. An override key that can reverse the core is, by definition, a point of centralization, capture, coercion, and compromise. Whoever holds the override, in the end, holds the firm.

SECTION 7 · THE ONCHAIN CORPORATION



Integrity at the Core, Adjudication at the Edge

An ever-expanding onchain core executes more contracts and governance by code — while accountable parties adjudicate the exceptions at its edge.



The core settles what code can settle; people adjudicate what it can't — and it is expanding.

The Agentic Economy

Concept 7.6 · integrity-core-adjudication-edge

Integrity at the Core, Adjudication at the Edge

The win isn't getting rid of intermediaries. It's turning intermediation into something transparent and contestable: an override that's multi-party, time-locked, and on the record rather than a silent administrative key; an adjudication layer whose rules and incentives are visible; oracles whose sources can be checked. Integrity at the core, accountable intermediation at the edge. That hybrid is the real architecture of the evolving firm. It's more honest, and more robust, than either a fantasy of pure trustlessness or a defense of the opaque intermediaries it replaces.

Two cautions complete the picture. Onchain governance isn't automatically more democratic than the cap table and proxy process it replaces. Token-weighted voting is plutocratic by design. It concentrates power in large holders and insiders, suffers chronically low participation, and opens new attack surfaces of vote-buying and governance extraction, a concentration risk the next part of this work takes up directly. And the timeline shouldn't be oversold. The direction is real and the substrate is being rebuilt, but converting the incumbent corporation is a decade-scale project, gated by the most conservative institutions in finance.

The shape of the destination isn't in doubt. The firm takes on a new material form: money, decisions, coordination, and outside economic relationships all expressed in software, wrapped in a thin legal shell conferred by a sovereign, recorded on a provable ledger that serves as the governance backbone across the human-machine seam, with contracts performed by programs and governed by law, and integrity at the core balanced by accountable adjudication at the

edge. Some firms get there by evolving and some by being born there, and the two journeys run in parallel. But the endpoint is the same, and it's the one the whole argument has been building toward: the agentic corporation and the onchain corporation are one and the same, because the agentic economy is the onchain economy.

 CONTINUE ONLINE

Read “The Onchain Corporation” online, with live cross-links to related sections, the shorter reads, and the audiobook.

AgenticEconomyTreatise.com/treatise/section-7.html

SECTION 08

Implications and the Concentration of Power

The agentic economy carries the greatest opportunity and the gravest risk of the era in the same hand. These are not separate futures to choose between. They are joint consequences of the same machinery, and the balance between them is not yet set. This section is the diagnosis. It specifies the mechanisms by which each outcome could prevail and what tips the scale. The constructive answer it points toward is the subject of the section that follows.

Let us begin with Labor. The claim is not that automation destroys jobs on net. That "lump of labor" assumption has been wrong for two centuries. Displaced work has reliably been reabsorbed into new tasks, and comparative advantage guarantees that even an agent absolutely better at everything is only *relatively* better at some things, leaving humans something to do. But comparative advantage is a statement about *relative* productivity. It is silent on *price*. Humans can stay employed at the task where machines are comparatively weakest while the wage that task commands falls below what sustains a household. That is "full employment" formally and a catastrophe socially. So the serious claim is not about the *level* of employment. It is about the *labor share of output* and the *wage at which human work clears*.


SECTION 8 · IMPLICATIONS AND THE CONCENTRATION OF POWER



The Labor Share, Not Jobs

The risk is to labor's share of income, not the existence of jobs — and it is not the lump-of-labor fallacy.

HOW NATIONAL INCOME SPLITS



A distributional risk — who captures the income, not whether work survives.

The Labor Share, Not Jobs

It holds under three conditions.

First, new tasks keep emerging, as they always have, but software colonizes them faster than humans can be retrained into them, collapsing the historical lag between displacement and re-absorption toward zero. The agent is not merely a better worker on today's tasks. It is a faster occupier of tomorrow's.

Second, the marginal price of agentic labor tracks the falling cost of inference, so across more and more of the task frontier the machine is simply cheaper at the margin, dragging the human clearing wage down with it.


Third, and this is the genuine break from every prior wave, capital can fund its own expansion. A loom never earned the capital to buy the next loom or wrote the design for it. Agents can generate the surplus that finances more agents and authors their successors, making the capital share self-reinforcing rather than bounded by diminishing returns. This is the capital-to-software-to-capital loop. Where the three conditions hold, the route by which most people have always claimed a share of output, their labor, narrows toward a trickle.

SECTION 8 · IMPLICATIONS AND THE CONCENTRATION OF POWER



The Capital → Software → Capital Loop

A loom couldn't buy the next loom; agents can fund and author their successors — capital becomes software becomes capital, self-reinforcing.



THEN

A loom could not reinvest in itself. ✕

NOW

Agents fund and author their successors. ✓


The Capital → Software → Capital Loop

the labor question and the ownership question are the same question

But even all three conditions holding does not produce immiseration. It produces a *distributional* problem, not an *output* problem. The output of such an economy is plausibly enormous. This is, after all, the abundance case. The catastrophe, if it comes, is entirely about *who owns the capital stock that is producing the abundance*. And the pessimistic case quietly assumes humans retain no priced edge and own nothing.

Neither is a law of nature. Human provenance may itself command a premium in care, status, authenticity, and the simple demand to be served by a person, and rising abundance could widen that space rather than shrink it. Wherever an agent must touch the physical world, bear legal liability, or actuate in regulated domains, a cost floor persists and human work clusters there. And if displaced workers *own* capital, through pensions, broad equity, or distributed tokens, the fall in the labor share is offset by a capital share they participate in.

That is the hinge of the entire section, and it should be stated bluntly: the labor question and the ownership question are the same question. The collapse of the labor share is a social catastrophe only if ownership is concentrated. If ownership is broad, the identical automation is simply broadly shared abundance. The fear of mass technological unemployment is, on inspection, a fear about the distribution of capital. That is why the answer, when we reach it, is about ownership and not about preserving jobs. It also reconciles the tension flagged at the outset of this argument: the individual is genuinely magnified by agentic tools even as the labor share of output may fall. Both are true at once. The same person is made far more capable and stands to claim a far smaller share of what is produced.




Labor and Ownership Are One Question

The contradiction dissolves only at the level of ownership. A magnified individual who owns a share of the machines shares in the abundance they create. A magnified individual who owns nothing is merely a more capable worker whose wage is still being dragged toward the cost of the agent working beside them.

That makes the concentration of power the decisive question. "Concentration is the default" is not a law of physics. The history of open protocols, forking, and commoditization is a record of cases where power dispersed rather than pooled. Concentration wins at layers that combine

strong increasing returns, where network effects and data flywheels mean each user makes the product better for the next, with a *non-forkable bottleneck*.

You can fork open-source code. You cannot fork the dominant currency, a regulatory license, a deep liquidity pool, or an override key. Openness disperses power at the layers where the frontier commoditizes, switching costs are low, and forking is credible. So the right question is not whether the agentic economy concentrates. It is *which of its layers* sit on non-forkable bottlenecks.



Which Layers Concentrate

You can fork open-source code. You cannot fork the dominant currency, a regulatory license, a deep liquidity pool, or an override key

The foundation-model layer, often imagined as the great toll, is among the *most* likely to commoditize. Frontier capability has repeatedly been matched by open weights within a year or two, and this is accelerating, and inference prices have fallen by orders of magnitude, so the durable rent migrates off the raw model and onto its complements: proprietary data, distribution, and the lock-in of an agent that knows you.

Bridges are less a rent story than a security one, the most-exploited surface in the space. Oracles tend to concentrate around the source everyone else uses, but remain forkable. The genuinely durable chokepoints are the identity layer and the override key. Identity, because one verified identity is reused everywhere, with winner-take-most dynamics. The override, because whoever can reverse the deterministic core, in the limit, controls the firm. And token governance, sold as democratization, is plutocracy by the construction of one-token-one-vote unless it is deliberately engineered against that force.

One concentration layer deserves direct treatment, because the author writes from inside it. A dominant stablecoin issuer captures the yield on the reserves backing the money it intermediates. The author and the author's own industry are among the parties positioned to benefit from exactly the concentration this section warns about.

The issuer's reserve yield is largely a *policy artifact*. It exists in its current form because the law forbids issuers from passing that yield to holders as interest, even as it increasingly permits value to be returned through usage- and activity-based rewards and distribution partners rather than as a direct, deposit-like payment. What statute creates, statute can redistribute. What is barred as interest can be channeled as participation, by competition, by regulation, or by sharing it with the people who hold and use the money.

This is not aspiration. Broad distribution of reserve income to the ecosystem is already the operating practice, and it is rising. So the resolution is evidenced, not merely promised. This rent is contestable and policy-dependent, reasonable policy may well require more of it to flow to holders.

The sharpest concentration concern sits not here but at the intelligence-infrastructure layer, in entities that hold both economic and intelligence infrastructure at once, and this is increasingly a topic of policy discussion: should society at large hold stakes in the capital of the largest AI companies.


The concentration question is also the geopolitical one, because the chokepoints that pool rents are the ones that become weapons. The economic operating system should be technologically neutral and governed by many stakeholders rather than by any single state, and at the protocol and technical layers that neutrality is genuinely achievable and worth defending, because those layers are forkable and ownerless.

SECTION 8 · IMPLICATIONS AND THE CONCENTRATION OF POWER



Concentration Is Geopolitics

The chokepoints that pool rents are the ones that get weaponized; velocity is a brake, not a cure.



What pools rents gets weaponized; interdependence buys time, not safety.


Concentration Is Geopolitics

The agentic economy could be the most pacifying economic order ever built or the most weaponizable. Which it becomes depends on whether its chokepoints are neutralized or captured

But the monetary layer cannot be made neutral by governance design as long as one state issues the dominant unit and can reach through the rails to freeze, seize, or exclude. A dominant digital dollar is, by the same fact, both the most likely path to global scale and the single largest concentration-of-control risk. Pretending the whole stack can be neutral is the naive version of the argument. Specifying which layers can and cannot is the credible one.

Whoever sits at the settlement rail, the reserve currency, or the identity layer gains both a panopticon and a switch, the weaponized interdependence already visible in financial sanctions and clearing exclusion. Deep interdependence raises the cost of conflict at the margin and modestly reduces its frequency below the great-power threshold, especially among democracies. The agentic economy could be the most pacifying economic order ever built or the most weaponizable. Which it becomes depends on whether its chokepoints are neutralized or captured.

What, then, of the abundance, the promise that this same economy is the greatest opportunity in history to lift people out of poverty and widen access to health, education, and information? It is real, and it deserves its weight. But abundance and inequality can rise together. The strongest optimistic case is consumer surplus. If agents drive the marginal cost of medical diagnosis, tutoring, legal advice, and expert knowledge toward zero, the real living standard of someone with very little money income can climb steeply even as measured income and wealth inequality widen.



Abundance and Inequality, Together

The ledger is two-sided, and the two sides can diverge. A world can become more unequal in income and ownership while becoming more equal in what people can actually consume and do. Which leg of the K dominates is not a mood or a prophecy. It is determined by the breadth of ownership and the design of fiscal policy. Broad ownership and a redistributive hand make the same automation a story of shared abundance. Concentrated ownership and a passive state make it a story of upheaval atop already-extreme wealth concentration.

The compression of the labor share, the concentration of power at the non-forkable choke-points, and the geopolitics of weaponizable infrastructure are not three problems but one. The one thing they share is the question of who comes to own and govern the concentrating layers.

Concentration is the default at exactly those layers, reinforced by the deepest regularities of platform economics.

Broadly shared abundance is achievable, but it must be *built*, through the distribution of ownership and the governance of the chokepoints, against the gravity that pulls the other way. The technology has, for the first time, made the better outcome genuinely constructible rather than merely imaginable. It has not made it automatic. How ownership might actually be distributed widely enough to matter, and how the chokepoints might be governed rather than captured, is the question the final part of this work must answer.

 CONTINUE ONLINE


Read “Implications and the Concentration of Power” online, with live cross-links to related sections, the shorter reads, and the audiobook.

AgenticEconomyTreatise.com/treatise/section-8.html

SECTION 09

The Civic Vision

The agentic economy breaks the link between labor and a share of output. The answer is not to defend the jobs. It is to broaden the ownership of the capital that is accreting the value: the agents, the models, the infrastructure, and the firms. That is the resolution toward which this entire argument has pointed, and it is genuinely available. The same onchain architecture that, left to its defaults, concentrates ownership at a handful of chokepoints can instead distribute ownership, outcomes, and governance more widely than any economic order in history.



Broaden Ownership, Don't Defend Jobs

Begin with the lineage, because it sets the scale of what is possible. The joint-stock corporation was a profound social technology. By letting strangers pool capital and share in the long-term outcomes of an enterprise, it made possible capital markets, the corporation as a legal person, and the modern machinery of banking and exchange, and it widened participation in enterprise far beyond the merchant and the sovereign.

SECTION 9 · THE CIVIC VISION



The Joint-Stock Lineage

*The joint-stock company was a profound social technology;
the onchain era extends read and write to own.*

A PROFOUND SOCIAL TECHNOLOGY, EXTENDED



The Joint-Stock Lineage

The onchain economy extends that lineage and, in principle, surpasses it. For the first time the material exists to distribute not only the capital but the governance and the upside of enterprise to vast numbers of people at near-zero administrative cost, with digital tokens as the building block. This is what the Web3 movement reached for in its slogan of read, write, own: platforms whose users share in their governance and economic benefit rather than merely supplying the attention and data that others monetize. The aspiration is old and serious. What is new is that the cost of acting on it has collapsed.

It is tempting to claim that earlier movements for broad ownership, syndicalism, the cooperative movement, guild socialism, failed only because they lacked the coordination technology that onchain rails now supply, and that the ideals were right and merely awaited the infrastructure. That claim is too flattering to the tools and false to the history. Those movements were not defeated by a logistics problem. They were defeated by power, by capital, by employers, and by the state, through suppression, hostile law, and force.

The cases that succeeded prove the point from the other side. Mondragon's tens of thousands of worker-owners, the credit-union and mutual-insurance movements, large consumer co-operatives, all achieved distributed ownership and governance at real scale, for decades, with nothing more exotic than sound legal form and durable institutions, and no blockchain at all.

Furthermore, in our evolved systems of state-led market capitalism, the shift is not necessarily towards deeper democratization over industry itself, but instead a new vision for stakeholder participation, with broader societal ownership in the highly concentrated and disruptive technology platforms that are concentrating the greatest amount of capital.

So the claim is narrower and sturdier. Onchain rails lower the coordination and transaction costs of broad ownership and strip out the gatekeepers that once stood between ordinary people and a stake. That is a real and meaningful contribution. But they do nothing about the power asymmetry that actually killed the historic movements. The technology makes distribution cheaper and more feasible. It does not make it happen. Politics still decides that.

And the default, absent design, is re-concentration. The empirical record of token ownership is sobering. Presale and insider allocations, team unlocks, and above all secondary markets that funnel transferable tokens back to the largest holders the moment they have value have made onchain ownership at least as concentrated as equity, sometimes worse. One-token-one-vote governance is plutocracy by construction.

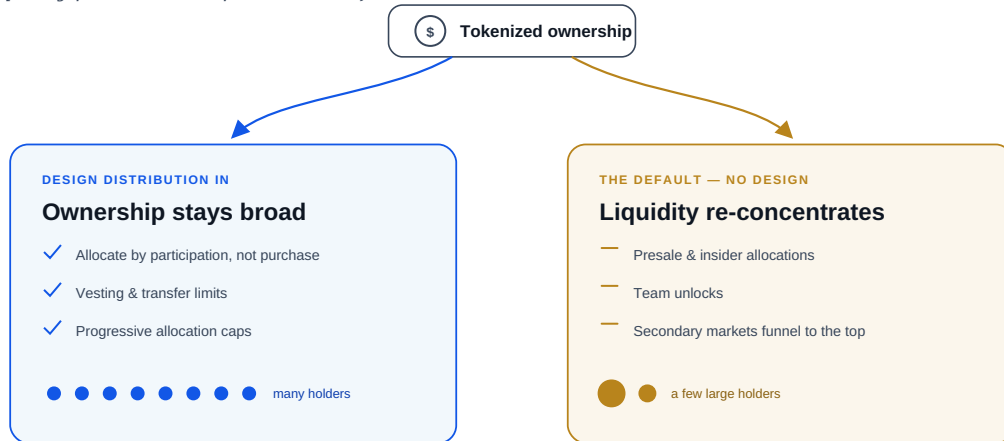
Liquidity is the enemy of broad ownership. So the claim that the architecture *can* be distributed is true and nearly trivial. The question is whether distribution is *designed in* against the gravity that pulls the other way. That design is possible, but it requires naming mechanisms and accepting their costs. Allocate ownership by *participation and contribution* rather than purchase, so the people who use and build a system come to own it. Use vesting and transferability constraints that blunt secondary-market re-concentration. Set progressive caps on allocation. Each of these fights capital's pull, and each carries a real tradeoff.

SECTION 9 · THE CIVIC VISION



Liquidity Is the Enemy of Broad Ownership

The rails make broad ownership cheaper — but absent design, liquidity pulls ownership back to the few. You must choose.



You cannot maximize liquidity and breadth at once — broad-but-constrained is a choice.

Liquidity Is the Enemy of Broad Ownership


There is a deeper trap the optimistic version glosses: distributing ownership is not the same as distributing power. Broad fractional ownership routinely coexists with no governance influence at all. The ordinary retail shareholder is the proof, and onchain the gap can be worse, through founder super-voting, foundation control, low-turnout delegated governance, and whale-dominated voting.

SECTION 9 · THE CIVIC VISION



Ownership Is Not Power

*Distributing ownership is not the same as distributing power;
aim governance at the real control points.*



A wide cap table can still leave power in a single pair of hands.


Ownership Is Not Power

distributing ownership is not the same as distributing power

More to the point, the previous section located the durable power not in cash-flow claims but in the control chokepoints: the override key that can reverse the system, and the identity layer reused everywhere. You can distribute economic exposure to a billion people and the holder of the override still controls the firm. So distributing governance is a separate and harder task than distributing upside, and it must be aimed squarely at those control points. We see this fully expressed today over the issue of sovereign intervention in foundational model deployment choices and safety controls.

That means governance designs that break the link between wealth and votes: one-person-one-vote secured by proof of personhood, quadratic and conviction voting, and rights weighted by contribution. It also means public-interest governance of systemically critical infrastructure: independent or public-interest board seats, golden shares, and escrowed, multi-party, on-the-record control over the override and identity layers, including the most powerful and capital AI foundation models, so they aren't private switches. Distributing the upside while the control chokepoints stay concentrated is exactly the half-measure the prior section exposed.

This is where the comfortable half of the vision has to be separated from the hard half. Market-native distribution, through tokens, universal access, and long-tail participation, grows the pie the author's own industry is building. It isn't sufficient on its own. The previous section was explicit that the outcome turns on how broadly ownership is held and how fiscal policy is designed, and the second half can't be quietly dropped, because it's the half that requires the state to act against concentrated interests.



Reserve Yield, Shared Broadly

The position is *and*, not *instead of*. Broaden ownership by design, and pair it with progressive taxation of capital and automation, public provision of the goods that abundance should make universal, and, in its strongest form, a public stake in the value the rails accumulate: a citizens' dividend or sovereign claim on critical infrastructure, so the public shares directly in the capital that's displacing labor income. Design distributes what the market will bear. Fiscal policy distributes what the market won't.

Open, forkable, interoperability-mandated standards at the identity, settlement, and model-access layers, so the chokepoints stay technically contestable and capture can be defeated by exit. Public-interest governance and escrowed control over systemically critical infrastructure. Ownership distributed by contribution and usage, with limits on transfer to resist re-concentration. And a fiscal instrument, a tax on automation or capital that funds a broad, onchain

ownership dividend, distributing a stake rather than just a transfer. These are starting points, not a finished program. But they're the kind of thing the closing call for executable building blocks actually has to contain, and several of them cost incumbents, the author included.

None of this wins on its merits alone, and the political economy has to be faced. The incumbents who benefit from concentration are the same actors with the power to write the rules. The distribution-favoring path doesn't win because it's more just. It wins only through countervailing power: forkability and open standards that make capture technically defeatable, so exit disciplines the chokepoints; antitrust and public mandates aimed at the control layers; public ownership of the most critical rails; and, not least, a broad class of owners with a real stake to defend, which is itself the constituency that makes the politics durable.

The claim that this can't be solved one country at a time is true, but "a shared set of ideas that civic society, business, and political leaders gather around" is not a transition plan. It's the thing that has to be fought for against those who profit from the alternative.

Underneath the economics sits the question the whole transformation finally forces. It deserves more than a gesture, even if it can't be fully answered here. If labor is no longer the organizing principle by which people earn a place and a voice, then ownership may have to take its place: a stake in the capital stock as a new basis of citizenship, the ground of standing, security, and participation that the wage used to provide. And the domains that stay distinctly human as machines absorb the rest: care, creativity, craft, judgment, the value of being made or served by a person, the work of community and meaning, aren't leftovers. They're the things a society of abundance could finally afford to honor and value, if it chooses to price and dignify them rather than treat them as what's left over. How people will actually live, spend their time, and be valued in such a world is the deepest open question this economy raises. The architecture enables an answer without supplying one. What it does is make the better answer possible. Choosing it is ours.

SECTION 9 · THE CIVIC VISION




Ownership as Citizenship

Ownership becomes a new basis of standing — a citizenship in the economy — while human work is honored.

A NEW BASIS OF STANDING



HUMAN WORK, HONORED



A stake confers standing; the dignity of human work is honored alongside it.

Ownership as Citizenship

That's the whole of the argument, and it ends where it began. The agentic economy is the on-chain economy: AI supplies the labor, and the onchain substrate supplies the form in which money, decisions, coordination, and ownership are expressed. That architecture cuts both ways, and its default is concentration, of income, of power, and of the weaponizable infrastructure everything runs on. But for the first time, the tools exist to distribute ownership, outcomes, and governance broadly enough to turn the same machinery toward shared abundance instead.

CONTINUE ONLINE

Read “The Civic Vision” online, with live cross-links to related sections, the shorter reads, and the audiobook.

AgenticEconomyTreatise.com/treatise/section-9.html