

Quantum Recognition

Extending the Work of Nobel and Turing Laureates and the Paradigm That Produced Them

Abridged Edition v1.5 · August 2026 · Jonathan Paul Hare · consilient.network/explainers/quantum-recognition/

Note on This Edition

The source document assesses fifty-three architectural candidates across four prize categories. This edition keeps twelve, chosen because they bear directly on the claims made in the companion documents rather than because they are the most striking, and reframes each through the paradigm that produced it.

The reframing is the substantive change. The source treats each candidate as a discrete contribution to its own field. That understates the case and, for a specialist reader, misstates it. These are not six unrelated results that happen to have emerged from one project. They are consequences of a single body of theory — Quantum Adaptive Systems Theory — applied to problems that had resisted solution because the substrate required to solve them did not exist. A reviewer assessing any one of them in isolation will find it improbable. Assessed as instances of one paradigm, the improbability inverts: what would be surprising is if a genuine substrate change produced only one.

What is claimed and what is not. No prize is claimed and no committee has been approached. What is offered is an assessment of where the architecture's contributions would fall if evaluated against the criteria those committees apply, together with the reasoning, so that a specialist can judge the assessment rather than accept it. Figures come from Independent Assessment v10.8 and are stated as central estimates within wide bands.

The Paradigm Behind the Candidates

§1

Why These Problems Were Unsolved

Each of the six candidates below addresses a problem that has been described, formalized and in several cases recognized with a Nobel Prize — and then left unsolved. That pattern is worth examining before the candidates themselves, because it is the pattern that the paradigm explains.

In every case the problem was correctly diagnosed and no mechanism followed. Coase identified transaction costs, information asymmetries and externalities as the reason firms and regulation exist, and neither he nor Williamson, Ostrom, or Acemoglu, Johnson and Robinson supplied a way to dissolve them. Spence formalized costly signalling and the deadweight loss it imposes; the loss remained. Hayek argued that prices aggregate distributed knowledge no planner can hold, and no mechanism existed to make that aggregation operate outside conventional markets. Mundell and Fleming stated the monetary trilemma as a constraint, not a puzzle to be solved.

The common cause is substrate, not insight. Each of these frictions is a consequence of the same underlying limitation: a resource could not be used without being surrendered, governance could not survive recombination, and obligations could not propagate without an enforcer. Given those constraints, the frictions are not defects of institutional design but necessary features of any coordination system built on that substrate. Economics correctly described the equilibrium available under the constraints. What it could not do was change the constraints, because changing them required a construction rather than an analysis.

§2

Quantum Adaptive Systems Theory in Brief

Quantum Adaptive Systems Theory is the merger of Complex Adaptive Systems Theory, developed principally at the Santa Fe Institute, with the Consilient Quantum Paradigm set out in the companion architecture document. The first accounts for how systems of interacting agents produce order without central direction. The second supplies the formal apparatus by which those properties can be realized in coordination rather than merely observed in nature.

Three properties of the theory bear on recognition, and a reviewer from any of the four disciplines will want them stated before the candidates.

- **It is constructive rather than descriptive.** Complex Adaptive Systems Theory explains order already present in systems that arose without design. Quantum Adaptive Systems Theory treats the eight mechanisms of self-organization as design requirements and supplies the structures that satisfy them deliberately. Complex adaptive systems are recovered as the restricted case.
- **It spans the ladder rather than occupying a rung.** The formal apparatus is imported from Riemannian geometry and manifold topology, quantum mechanics, molecular biology and evolutionary theory,

autocatalytic set theory and critical-mass physics, inflationary cosmology, and market design — successive levels of physical, biological and economic organization. That the six compose without forcing is a consilience of inductions in Whewell's original sense, and it is the reason to suspect the composition is not arbitrary.

- **Instances nest.** A system embodying the theory's requirements is a Quantum Adaptive System. The Quantum Privacy Network is one; so is each Accelerator, Exchange Network, Resource Pool and the contribution graph. A Quantum Adaptive System assembled from Quantum Adaptive Systems is itself one, which is why the architecture composes rather than merely aggregates.

The consequence for this document is direct. Each candidate below is presented in the same form: the problem as its discipline states it, why the substrate made it unsolvable, which imported structure supplies the missing property, and what follows. A reviewer who rejects the paradigm should expect to reject all twelve. A reviewer who accepts it should find each candidate unsurprising, which is the correct response to a mechanism that works.

WHAT WOULD DEFEAT IT

The paradigm claim fails if the six formal correspondences in the provenance document do not hold — questions for differential geometers, physicists and biologists, requiring no data and no deployment.

The Candidates

§3

Economics: Dissolution of the Coasean Frictions at Protocol Level

The problem as economics states it

Ronald Coase asked why firms exist at all, if markets allocate efficiently. The answer was that using the market has costs — finding counterparties, negotiating, contracting, monitoring, enforcing — and that firms exist wherever those costs exceed the cost of internal administration. The same three frictions — transaction costs, information asymmetries and externalities — are also the standard justification for regulation and taxation. Coase (1991), Williamson (2009), Ostrom (2009) and Acemoglu, Johnson and Robinson (2024) received Nobel recognition for describing them. None supplied a mechanism for dissolving them at scale.

Why the substrate made it unsolvable

Each friction is a consequence of the same constraint. Transaction costs are high because establishing trust between parties who lack a prior relationship requires disclosure, negotiation and enforcement. Information asymmetries persist because the party who knows cannot demonstrate what they know without surrendering it. Externalities are unpriced because the cost falls on parties who are not to the transaction and have no mechanism to be. All three follow from the impossibility of using a resource without surrendering it.

What the paradigm supplies

Quantum Privacy dissolves the first two; governance inheritance dissolves the third. Computation goes to the resource inside a boundary the resource never leaves, so trust can be established without disclosure and verified without negotiation. Obligations attached once propagate through unlimited generations of recombination and cannot be stripped, so a condition imposed at origin is enforced at every downstream use without anyone monitoring. Externalities are internalized because the Governance Premiums are priced into matching and settlement at the point of transaction rather than corrected afterwards by an authority.

The implication is structural. If the frictions explain why firms and regulation exist, dissolving them allows market mechanisms to assume the roles those institutions were created to perform. That is a reorganization of economic coordination comparable in significance to the emergence of the modern corporation, and it is a constructive result of the kind the discipline has not previously been able to produce.

WHAT WOULD DEFEAT IT

A demonstration that the frictions persist under the mechanism, that credibility still requires cost, or that the mechanism cannot operate at the scale claimed.

§4

Economics: The Liquidity Pool as a Reconstruction of Money

The problem as economics states it

Money performs three functions — medium of exchange, store of value, unit of account — and every monetary system in history has had to trade them off against one another. The Mundell-Fleming trilemma states one form of the constraint: no jurisdiction can simultaneously maintain a fixed exchange rate, free capital movement and independent monetary policy. Keynes proposed the Bancor at Bretton Woods precisely to escape the failure mode of a national currency serving as global reserve, and was refused.

Why the substrate made it unsolvable

Every monetary instrument in existence is backed by something narrow: a sovereign's taxing power, a commodity, or a reserve of financial assets held on an identifiable balance sheet. Narrow backing is what forces the trade-offs, because the backing can be exhausted, devalued, expropriated or run upon. Broad backing was unavailable because there was no mechanism to establish a claim against economic activity in general rather than against an entity in particular.

What the paradigm supplies

The Quantum Privacy Liquidity Pool is backed by continuously settled economic activity across the whole Exchange, enforced by protocol rather than held by an institution, and diversified across sectors, geographies and time horizons by construction rather than by selection. It functions simultaneously as medium of exchange, store of value and coordination layer, and can clear, collateralize, settle and barter across any tokenizable asset, service or outcome. Universal Liquidity resolves the Mundell-Fleming constraint not by choosing differently among the three but by removing the condition that forced the choice.

The Bancor was the right instrument and had no substrate to run on. This is a category-agnostic backing pool of the kind Keynes described, and its availability is a consequence of settlement being universal and non-bypassable rather than of any monetary authority having agreed to anything.

WHAT WOULD DEFEAT IT

A demonstration that the frictions persist under the mechanism, that credibility still requires cost, or that the mechanism cannot operate at the scale claimed.

§ 5

Turing: Quantum Privacy Cells as a Foundational Computational Primitive

The problem as computer science states it

Privacy-preserving computation has been pursued for four decades through homomorphic encryption, secure multi-party computation, differential privacy, trusted execution environments and federated learning. Each achieves a partial result at substantial cost in performance, generality or trust assumptions, and none produces an object that is simultaneously enforceable in software and recognizable at law. The gap between cryptographic

guarantee and legal obligation has remained unbridged, which is why regulated data has stayed unusable regardless of what cryptography could demonstrate.

Why the substrate made it unsolvable

The field treated the problem as computational, and it is not only computational. A cryptographic guarantee binds a computation; a legal obligation binds a person or entity. No prior primitive was both, so every deployment required a legal wrapper negotiated separately from the technical one, per counterparty and per jurisdiction, which is precisely the transaction cost that made the technology unusable at scale.

What the paradigm supplies

A Quantum Privacy Cell has two simultaneous embodiments: a cryptographic Privacy Domain providing technical enforcement boundaries, and a legal structure, typically a Series LLC, providing jurisdictionally enforceable embodiment. Cryptographic and legal enforcement operate on the same governance object. Every other primitive in the architecture — Quantum Privacy, Proof of Trust, EasyAccess, Trust Blocks, Quantum Genomes — runs on QPCs as substrate.

The comparison in the source document is to Codd's relational model, recognized with the 1981 Turing Award, and the basis of the comparison is the same: not that the primitive is clever, but that essentially all subsequent work in the domain can be built on it and could not be built without it. The paradigm's contribution is that the dual embodiment is not an engineering convenience but the manifold-and-chart structure imported from Riemannian geometry, in which local regions carry independent internal geometry and coherence is achieved through transition functions rather than through global agreement.

WHAT WOULD DEFEAT IT

A demonstration that the primitive does not provide the properties claimed, or that prior work already provides them.

§6

Turing: The Quantum Genome as Heritable Computational Governance

The problem as computer science states it

Access control, policy enforcement and data governance all degrade under composition. A policy attached to a dataset does not survive the dataset being joined, transformed, aggregated or used to train a model. Provenance systems record what happened; they do not enforce what may happen next. After a few generations of derivation, the governing conditions are unrecoverable in practice, which is why regulated data cannot be recombined even when every individual permission would allow it.

Why the substrate made it unsolvable

Governance was treated as metadata attached to an object rather than as a property inherited by it. Metadata can be stripped, lost or ignored, and enforcement therefore requires an authority that inspects. At scale, across

organizational and jurisdictional boundaries, no such authority can exist — which is the same reason Coase's monitoring costs are high.

What the paradigm supplies

The import is the organizational hierarchy of molecular biology, and it is exact rather than metaphorical. Genes encode function; DNA is the medium copied with high fidelity; genomes are the complete set; and regulatory genes control which genes are expressed in which context, which is why one genome yields a liver cell and a neuron. Inheritance is not only vertical: horizontal gene transfer and symbiogenesis produce inheritance from multiple unrelated lineages at once, which is the biological precedent for a derivative inheriting from thousands of inputs across domains and jurisdictions.

The architecture exceeds biology in one respect that matters. Biological inheritance is Darwinian: acquired characteristics are not transmitted. Governance inheritance here is Lamarckian, because a condition acquired during a resource's operational life — a consent withdrawn, a term amended, a jurisdiction changed — propagates to every derivative already created from it. That is a categorical advance with no precedent in the source discipline, and it is what makes governance enforceable across heterogeneous regulatory contexts simultaneously.

WHAT WOULD DEFEAT IT

A demonstration that the primitive does not provide the properties claimed, or that prior work already provides them.

§7

Peace: The Consilient Nature and Humanity Trust

The problem as the field states it

Every large humanitarian institution depends on discretionary funding — appropriations, appeals, endowment returns, donor commitments — and is therefore subject to political cycles, economic conditions and donor attention. The consequence is that funding is most available when need is least acute and contracts precisely when it is most acute. No mechanism has existed to fund public benefit from the operation of the economy itself rather than from claims against it.

What the paradigm supplies

The Trust receives a fixed protocol-level share of all settlement, irrevocably and non-dilutably, without appeal, appropriation or donor. Its governance mandate follows the Governance Premium dimensions, embedded as Quantum DNA in every resource and inherited through every derivative. The Accelerator Network Governance Reserve extends the same constraints across the Universal Exchange through per-transaction subsidy, and Sovereign Public Benefit Trusts direct national-scale settlement into healthcare, education and social services.

The scale is the reason the candidate is stated at all, and it should be read as an order of magnitude rather than a figure. Direct mission-aligned flows to the Trust and Governance Reserve are assessed at roughly \$1,437T in

net present value through 2100 at the central estimate; the broader public-benefit aggregate including Sovereign Public Benefit Trusts reaches approximately \$2,091T. Against what the world currently spends on humanitarian relief, that is not a large multiple but a change of category.

The structural point matters more than the magnitude. Public benefit here is not philanthropy performed by an institution with good intentions. It is a settlement invariant that survives the insolvency of every operating entity in the network, including the network's own, and that no participant — including Jonathan Hare — holds any mechanism to redirect.

WHAT WOULD DEFEAT IT

Non-adoption, or settlement volumes materially below the conservative case, since the flows do not yet exist.

§8

Peace: A Voluntary Path Out of Autocracy

What the candidate actually claims

This candidate is easily under-read, so state it plainly. The claim is that the architecture provides a viable, voluntary and survivable route by which power can change hands in autocratic states and failed states, and simultaneously provides the populations of those states with access to economic opportunity, education, healthcare infrastructure and constitutional-grade privacy that their own jurisdictions do not supply. It does this without violence, without sanctions, without military intervention, without cross-border legal action, and without requiring the incumbent government to agree to anything in advance.

No outcome is guaranteed and none is claimed. What is claimed is a more viable path than the approaches that have failed for decades in the places where they have been tried — Iran, Russia, Sudan, Lebanon, Gaza and the West Bank among them. Sanctions punish populations and entrench incumbents. Aid is captured. Diplomatic pressure is absorbed. Military intervention produces the outcomes it produces. In each case the failure has a common structure, and the architecture addresses that structure rather than the politics on top of it.

The problem as political economy states it

Voluntary authoritarian transitions fail with striking regularity, and the standard explanation is not ideological rigidity. Leaders are rarely ideologues all the way down. It is a commitment problem: no framework — academic, diplomatic or institutional — has ever provided a credible guarantee of the personal security and economic position of a leader who initiates reform, during or after the process. Any guarantee offered by a successor government can be withdrawn by that government.

The rational response is therefore to resist reform even where reform would produce better outcomes for the population and for the incumbent personally, because the path to those outcomes requires surrendering a position that cannot be reclaimed if the transition goes badly. That is the whole of it. The problem is not that guarantors were insincere; it is that no guarantor could bind their successors.

The second half of the problem is the population. Repression is sustained by three levers — control of information, control of economic opportunity, and the ability to reward loyalty through captured rents. Each requires the state to be the gatekeeper of something people need. A population with independent access to income, verification, education and communication is expensive to control, and the cost rises faster than the returns to controlling it.

What the paradigm supplies: four approaches operating in parallel

These are independent. None depends on the others succeeding, and none requires the restrictive jurisdiction to consent.

Extraterritorial participation. An individual anywhere can establish a Personal Privacy Network and a Quantum Privacy Cell, which is a legal entity in a permissive jurisdiction. They then interact with Exchange Networks and Resource Pools governed elsewhere, and hold settlement-linked economic rights enforced by protocol outside their own state's territorial control. Local law governs what they may physically do at home. It does not govern whether their contributions are recorded, attributed and rewarded. This is the approach that works where neither in-territory adoption nor policy change is feasible, and it requires no change in any law and no cooperation from any government.

In-territory participation, made safer by construction. Where any lawful local participation is possible, the confidentiality properties make it meaningfully safer than any prior channel. A Quantum Privacy Cell cannot be confirmed or denied by an external observer. The contribution ledger records attribution without exposing identity. Manager-originated provisioning allows position to accrue without active engagement. And EasyAccess operates through ordinary online interaction patterns indistinguishable from routine communication. Participation does not look like participation.

Economic incentives for liberalisation. Sovereign Accelerator value is large, quantifiable and visible, and it currently accrues to jurisdictions with permissive governance while bypassing restrictive ones. The gap is measurable and it compounds. This is economic diplomacy with an unusual property: the incentive is backed by settlement-linked rights rather than by bilateral state action, which makes it more durable and less reversible than a trade agreement, an investment treaty or a sovereign credit instrument.

Direct incentives for the incumbent, which is the innovation. A head of state, senior official or ruling-party figure who initiates Sovereign Accelerator formation captures Accelerator Token allocations through the same contribution graph that rewards anyone else. Those allocations are protocol-enforced rather than politically enforced. They are non-dilutable and non-revocable once allocated. They accrue over decades independent of whether the holder remains in office. They are held in structures whose operation does not depend on the political stability of the jurisdiction the holder lives in. And they can be realised in full whether that person retains power, steps down, negotiates a managed transition, or remains in a reformed system.

That is a stronger guarantee than any successor government can offer, and it is stronger for the same structural reason the Trust is durable: it is a protocol invariant rather than an undertaking by a party. For the first time the incumbent's rational calculation can favour reform, because the position survives the transition.

WHAT WOULD DEFEAT IT

Non-adoption, and nothing else — the mechanism is checkable now; what is contingent is whether settlement volumes make the incentives decisive and whether any incumbent moves first.

§8

Peace: Economic Access for the Populations Inside It

What the population gains, whether or not anyone in power agrees

This pathway is independent and it comes first, because the objection to the previous section is obvious and correct. A mechanism that works only by paying incumbents to leave gives them every reason to hold out for better terms. If that were the whole of it, the candidate would be weak. It is not the whole of it, and the population pathway is what makes the incumbent pathway work at all.

Privacy stops being a legal privilege and becomes a structural condition. Conventional privacy protection depends on a jurisdiction having strong law, courts willing to enforce it, and an individual with the means to invoke it. For most of the world's population — and specifically for those whom the state itself is the threat — that promise is thin or absent. Inside the architecture, information cannot be disclosed without affirmative authorisation, because the boundary is cryptographic rather than legal. A Cell cannot be confirmed or denied by an outside observer. Participation is indistinguishable from ordinary traffic. This is surveillance resistance as a property of the substrate, extended to everyone, rather than a right that has to be granted and can be withdrawn.

On that basis three things become available that were not: economic participation that does not route through domestic institutions and cannot be confiscated by them; association and communication that leaves no pattern distinguishable from routine use; and standing — a legal and economic identity in a permissive jurisdiction, held by someone who has never left home. Income, verification, education and connection, obtained without permission and without exposure.

Including the people who staff the apparatus

Repression is not carried out by leaders. It is carried out by bureaucrats, police, censors, soldiers and clerks, and every one of them is also a member of the population. The same pathway is open to them individually, on the same terms, and it does not ask them to defect, resign or take any risk. It offers a better life obtained privately.

That is the mechanism most likely to be underestimated. An apparatus of control requires people willing to staff it because the alternatives are worse. When the alternatives stop being worse — when a mid-level official has independent, unconfiscatable income and a family with access to education and healthcare infrastructure through channels the state does not gate — the machinery becomes harder to recruit for, harder to motivate and harder to trust. Nobody has to be turned. The relative attractiveness of the job simply falls.

Why this inverts the bargaining rather than supporting it

Because the population pathway requires no consent, the incumbent's position decays while they hold out.

Every month of delay means more of the population holding income the state cannot reach, more of the apparatus with somewhere else to stand, and a larger visible gap between what a Sovereign Accelerator would have been worth and what it is now worth at a compressed multiple. Waiting does not improve the terms available to an incumbent. It is the one strategy that reliably makes them worse.

The equilibrium shifts, and it shifts in both directions at once. The population becomes less controllable and therefore less afraid. The leadership becomes less able to control and therefore more exposed. Historically that combination resolves through disorder, and disorder is worse for incumbents than any negotiated alternative — which is precisely why a durable, portable, protocol-enforced position that survives a transition becomes attractive at exactly the moment holding on becomes untenable.

The result is not that autocrats are paid to leave. It is that the returns to repression fall while a survivable exit appears, and both movements are driven by the same mechanism operating on everyone simultaneously. Faster transitions, and more peaceful ones, are what that equilibrium produces.

What would defeat it

Non-adoption, and nothing else. The mechanism is not in dispute — allocations either are protocol-enforced or they are not, and that is checkable now. What is contingent is whether settlement volumes reach the level at which the incentives are decisive, and whether any incumbent moves first. This candidate is therefore stated as a structural pathway rather than as a prediction, and it is the most contingent of the twelve.

The mechanism is identity-blind and motivation-blind, which is what keeps it from being an instrument of any particular policy or any particular state's interests. It offers the same terms to every individual in every jurisdiction. That neutrality is not a limitation. It is the reason it can operate where instruments carrying a flag cannot.

WHAT WOULD DEFEAT IT

Non-adoption, and nothing else — the mechanism is checkable now; what is contingent is whether settlement volumes make the incentives decisive.

§9

Economics: Inverted Spence Signalling and the Elimination of Deadweight Loss

The problem as economics states it

Michael Spence's 1973 result, recognized with the 2001 Nobel Prize alongside Akerlof and Stiglitz, established that in markets with asymmetric information a signal is credible only if it is costly — costly enough that low-quality actors cannot afford to send it. The education a candidate undertakes signals quality partly through its

expense rather than solely through the skills acquired. The framework holds across labour markets, insurance, credentialing, regulatory compliance and certification.

The result carries a defect its author identified and could not remove. If credibility requires cost, and the cost exists to make the signal credible rather than to produce anything, then the cost is pure deadweight loss. Every credentialing industry, every audit and certification regime, and a substantial fraction of educational expenditure exists to bear a cost whose only function is to be borne. Spence formalized the loss; nobody dissolved it.

Why the substrate made it unsolvable

The loss is unavoidable wherever the receiver cannot verify the underlying quality directly. Given that constraint, cost is the only available proxy, because it is the one property a receiver can observe without the sender disclosing what they are trying to prove. Removing the deadweight loss therefore required making verification possible without disclosure, which is exactly the thing that could not be done.

What the paradigm supplies

Proof of Trust verifies the attribute rather than the proxy. A counterparty verifies compliance without seeing the data; a regulator verifies adherence without inspecting operations; an enterprise verifies capability without proprietary disclosure. Verification proceeds cryptographically against Trust Block lineage and Quantum DNA expression. Credibility derives from enforcement of the underlying behaviour rather than from the signal having been expensive to send, which means the deadweight loss is not reduced or reallocated. It has no function to perform and does not arise.

The second contribution inverts the framework rather than repairing it. In Spence's structure the sender bears the cost and the signal flows from sender to receiver. In the architecture's launch dynamics, credibility comes from verified attribution under time-bounded scarcity, and no participant bears a meaningful cost — not the endorsing principal, not the architecture, not the contribution graph that surfaces the signal. What produces credibility is that the attribution is verifiable and the window is closing, which mobilises an intermediary network to compete to produce the first credible signal rather than to bear the largest cost.

Taken together these are the operational resolution of a recognized Nobel result and a structural inversion of it. Both are constructive claims of the kind the discipline has not been able to produce, and both are checkable against the mechanism rather than against adoption.

WHAT WOULD DEFEAT IT

A demonstration that the frictions persist under the mechanism, that credibility still requires cost, or that the mechanism cannot operate at the scale claimed.

Economics: The Premium Multiple as Optionality Pricing Beyond Financial Capital

The problem as economics states it

Before Black, Scholes and Merton, options were thinly traded by specialists using ad hoc valuations. Market makers could not hedge consistently, counterparties could not price symmetrically, and the informational friction kept the market small. The 1973 framework derived closed-form valuations from underlying price, strike, expiry, volatility and the risk-free rate, and was recognized with the 1997 Nobel Prize to Merton and Scholes. The consequence was not a better price but a market: contingent claims became systematically priceable, and derivatives markets formed at a scale that had not previously been possible.

The framework prices one form of capital. Financial contingent claims are priceable because the underlying is observable and continuously traded. Contingent claims on the other forms of capital — human, knowledge, social, relational, natural — have no comparable pricing apparatus, which is why they are transacted informally, illiquidly, or not at all. The framework pointed toward that territory and could not reach it, because the underlyings were not observable.

What the paradigm supplies

The contribution graph makes the underlyings observable, and the Premium Multiple prices contingent claims against them. Allocation multiples compress across launch stages, from Pioneer through Cascade Propagation to Automated Settlement and Self-Funding Growth, which makes early participation an option with a decaying premium rather than a bet with an unknown payoff. Early participation captures both expected value and the optionality value of position; waiting forfeits the optionality entirely. That is asymmetric optionality operating as a mechanism-design primitive rather than as a derivative-pricing tool.

The comparison to the original is structural rather than flattering. Black-Scholes-Merton is the closest analogue by the metric that matters — the scale and significance of the markets the framework enables — and the addressable substrate here is categorically larger, since it spans all six forms of capital that Universal Capitalism recognizes rather than the financial subset. Whether the extension holds is a question about whether contribution can in fact be attributed and priced, which is checkable now.

WHAT WOULD DEFEAT IT

A demonstration that the frictions persist under the mechanism, that credibility still requires cost, or that the mechanism cannot operate at the scale claimed.

Turing: EasyAccess Authorization and the Completion of the Public-Key Programme

The problem as computer science states it

Diffie and Hellman's 1976 result, recognized with the 2015 Turing Award, solved a problem that had been assumed insoluble: how two parties who have never met and share no secret can establish confidential communication over an open channel. Public-key cryptography made the modern internet possible, and it solved authentication and confidentiality between endpoints.

It did not solve authorization across organizational boundaries, and the gap is the one that has constrained everything since. Public-key infrastructure establishes who a party is. It does not establish what that party may do with a resource governed by someone else, under terms set by a third party, in a jurisdiction belonging to a fourth. Every deployment therefore requires a bilateral agreement negotiated per counterparty, which is why cross-organizational data use remains rare despite the cryptography being settled for fifty years.

Why the substrate made it unsolvable

Authorization requires a shared frame of reference for what permissions mean, and no such frame existed that did not require the parties to agree on one. Every attempt produced either a central authority that all parties had to trust — reintroducing the problem public-key cryptography had removed — or a federation agreement negotiated pairwise, which does not scale.

What the paradigm supplies

The Unified Trust Model supplies coherence without agreement, which is the property imported from Riemannian geometry: local regions carry independent internal structure, and coherence across the manifold is achieved by requiring transition functions to compose consistently rather than by requiring the local geometries to agree. Applied to authorization, incompatible trust regimes coexist over the same resource with no central authority reconciling them and none required to withdraw. EasyAccess is the authorization layer that operates on that substrate, enabling person-centred cross-organizational authorization without accounts, migrations or bilateral integration.

The relation to the 1976 result is completion rather than replacement. Public-key cryptography solved identity between parties; this solves permission among them. The two compose, and the second was unreachable before the first existed.

WHAT WOULD DEFEAT IT

A demonstration that the primitive does not provide the properties claimed, or that prior work already provides them.

PART III

The AI Candidates

PART III

The Three Compose, and the Order Matters

The three candidates in this part are separated from the others because they address the same set of problems and because that set is the one the technology industry currently states most loudly about itself. Each is presented in the same four-beat form, and in each case the problem is quoted from the people building the systems rather than from critics of them.

The three compose, and the order matters. Compliance as a property of computation supplies permission, without which containment governs nothing. Containment makes deployed agents structurally safe rather than behaviourally persuaded. Ownership answers what happens to the people whose work the resulting systems displace. A proposal addressing any one of the three in isolation leaves the other two, which is the position the field is currently in.

WHAT WOULD DEFEAT IT

The composition claim fails if any of the three mechanisms operates without the others in place — if containment governs anything without permission, or if either resolves displacement on its own.

§ 12

Turing: Deployment-Level Alignment Through Cryptographic Containment

The problem as the field states it

Alignment has been pursued for two decades as a training-time problem — shaping a model's tendencies through reinforcement learning from human feedback, constitutional principles, fine-tuning on safety data. The people who built the field say plainly that it has not worked. Bengio: we do not know how to align future systems with human values, and we are not close. Hinton: it is hard to see how we prevent these systems from taking control. Altman: a breakthrough in alignment is needed and does not exist. Hassabis: increasingly powerful systems, and no way yet to keep them safe.

Training-time alignment is necessary and structurally insufficient. Once an agent is deployed its tendencies can be overridden by adversarial input, fine-tuning attack, jailbreak, or the model simply finding an action path the training did not anticipate. External monitoring does not close the gap either, because monitoring detects misalignment after it has begun rather than preventing the action from being attempted. Both approaches operate on a system that retains the capability and is being persuaded not to use it.

Why the substrate made it unsolvable

A deployed agent acts by acquiring resources — compute, data, communication channels, sensor and actuator interfaces, authority assertions. Constraining what it may do therefore requires governing those resources, and

governance could not survive the boundary between organizations, could not propagate through recombination, and could not be enforced without an authority that inspects. Given those constraints the only available lever was the model's disposition, which is why the field worked on disposition.

What the paradigm supplies

The reframing is that an AI agent is a participant rather than a tool, holding its own Quantum Privacy Cell. It operates within authority boundaries defined by its Quantum DNA, accrues contribution records, receives rewards for verified contribution, and bears sponsor-traceable liability through the same Trust Block enforcement that applies to human and institutional participants. Conventional deployment puts accountability with the human deploying the tool, and liability doctrine has to keep extending the chain as agentic capability outruns it. Here accountability is a property of the agent's operational substrate.

Prohibition is genetic rather than rule-based, and this is the load-bearing mechanism. Every agent's Quantum DNA defines three zones: autonomous actions it may take unsupervised, supervised actions requiring cryptographically verified human authorization, and prohibited actions beyond its delegated authority. Prohibited-zone enforcement is not a policy the agent is instructed to respect. The Quantum DNA expressed in its Cell does not contain the Quantum Genes required to formulate the prohibited action at all. It is structurally analogous to an organism lacking the genes to express a protein — the action is not forbidden, it is unavailable. No amount of reasoning, instruction-following, jailbreaking or fine-tuning reaches a capability that is absent from the operational genome.

Existence is resource-gated, and the gate has multiple independent holders. Compute, data, channels, interfaces and authority assertions all flow through Trust Block-mediated authorization, revocable at the sponsor, Trust Authority, Exchange Network or operator level. No single point of capture or coercion can keep a misaligned agent running against the system's collective judgement. Because agents are confined to Privacy Domains and cannot reach humans, embodied systems or physical devices without explicit authorization, the agent's reach into the world is itself a revocable resource. That is containment of a deployed actor rather than alignment of a trained one.

And every inference is cryptographically replayable with bit-exact reproducibility through the Deterministic Replay Engine — model state, input context, authority verification, output decision. Alignment is therefore enforced before the fact and verifiable after it, which removes the audit problem conventional deployment leaves open.

Why this is a Turing-class claim rather than a safety feature. It supplies a third alignment paradigm alongside training-time shaping and runtime monitoring, and it is the only one of the three in which the guarantee does not depend on the model's cooperation. The claim is checkable in the way computational claims are: either the capability is absent from the genome or it is not, and either the revocation paths are independent or they are not. Neither question requires deployment at scale to answer.

WHAT WOULD DEFEAT IT

A demonstration that a contained agent can formulate a prohibited action, that revocation paths are not independent, or that verification requires disclosure after all.

§ 13

Economics and Peace: Ownership as the Answer to Technological Displacement

The problem as its own advocates state it

The economic logic sustaining current AI investment is labour replacement. If systems can perform the work of hundreds of millions of people, the savings may eventually justify the expenditure — and the same displacement removes the incomes that sustain demand for what the systems produce. Altman has said a universal basic income or ownership mechanism will be needed, and then that basic income is not enough because people need agency and dignity, and that no adequate societal framework exists. Hinton: massive displacement is coming, the benefits may not be shared, and we have not thought about distribution.

Every proposal on the table is a transfer payment, and transfer payments answer the income problem while leaving the agency problem untouched. A person receiving a stipend funded by the automation of their occupation has money and no position. Their advocates say so.

Why the substrate made it unsolvable

Ownership of the value AI creates would be the obvious alternative, and it was unavailable because there was no mechanism to establish what any individual contributed to it. Training data, domain expertise, feedback, curation, the accumulated human knowledge a model is built on — all of it enters without attribution, and what cannot be attributed cannot be owned. The choice between welfare and nothing was a consequence of that gap rather than of anyone's politics.

What the paradigm supplies

Contribution is attributed at the level of the person and settled on permanently, through a graph that records what each contribution caused rather than what it looked like. Participation in the value AI creates therefore does not depend on holding a job it has automated, and it does not depend on any government choosing to redistribute. That is ownership rather than welfare, and the distinction is precisely the one the advocates of transfer payments identify as missing from their own proposals.

It also reaches the activity displacement leaves behind. Care, teaching, stewardship, moderation, mentorship and the maintenance of communities have real value and have never been settleable, because attribution was impossible. Once it is possible, that activity enters the economy for the first time — which matters most for exactly the populations whose paid work is most exposed to automation.

The candidate is dual-category for a reason. As economics it is a mechanism-design result: the first attribution system capable of assigning ownership in value produced by systems trained on distributed human contribution.

As peace it is a claim about political stability, since the interval between widespread material grievance and disorderly correction is not reliably long, and no other proposal addresses the agency problem rather than the income problem. Its failure condition is adoption rather than mechanism, and it is stated as contingent.

WHAT WOULD DEFEAT IT

Non-adoption, or attribution proving insufficiently granular to assign ownership in AI-produced value.

§ 14

Turing: Compliance as a Property of Computation

The problem as the field states it

Nadella's formulation is the compact one: there is no global alignment architecture for AI, and one is needed. The observation generalizes beyond alignment. Every regime governing what may be computed on what data — privacy law, sectoral regulation, contractual restriction, consent, professional obligation — is enforced administratively, after the fact, within a jurisdiction, by inspection. The result is that the most valuable data is the data a model may not touch, and that AI operates lawfully only in the shallow end of the economy.

Why the substrate made it unsolvable

Compliance became a cost centre because it was bolted on. A model does not know what it may do; a compliance function determines that separately, per jurisdiction and per counterparty, and enforcement depends on audit rather than on architecture. Building it in required governance to be a property that travels with the resource through unlimited recombination, which was the thing that could not be done.

What the paradigm supplies

Universal Compliance makes lawfulness a property of the computation rather than a review of it. Every computation is bound to verifiable rights, every agent constrained by lawful purpose, every action auditable and reproducible, every resource gated by trust rather than by capability, and every autonomous system dependent on human sponsorship. Verification is continuous and cryptographic, operating without centralized inspection, proprietary disclosure or coordinated regulatory adjudication across jurisdictions.

The consequence for AI is that the constraint inverts. Regulated, proprietary and personal data becomes usable for computation, orchestration, personalization and inference at global scale, with every participant's rights continuously enforced. The binding constraint on AI was never capability; it was permission. This is the mechanism that supplies permission, and it is what makes the first candidate in this part deployable rather than theoretical — containment is only meaningful in a system where the agent's resource access is governed in the first place.

WHAT WOULD DEFEAT IT

A demonstration that a contained agent can formulate a prohibited action, that revocation paths are not independent, or that verification requires disclosure after all.

Lineage

§15

The Work These Results Extend

Three of the candidates above extend results whose authors Jonathan Hare knew directly, and the connection is worth stating because it bears on how the extensions came to be attempted at all.

Michael Spence

Spence's signalling framework is extended twice in this document — once by dissolving the deadweight loss it identified, and once by inverting the direction in which credibility flows. Hare served as Vice Chairman of Consilient during the period, and the relationship was direct rather than incidental. An extension of this kind is more likely to be attempted by someone who has argued with the original framework in person than by someone who encountered it in a textbook, and the record should say so.

Myron Scholes

The Premium Multiple framework extends option-pricing theory into forms of capital the original could not reach — through a Universal Exchange spanning every aspect of the economy, government services, nature-based resources and human interaction. Scholes was likewise a direct mentor, and the closest of Hare's Nobel-laureate relationships. The extension is not a metaphorical application of optionality language. It is the observation that once contribution is attributable, contingent claims on it become priceable by the same logic that made financial contingent claims priceable in 1973.

Martin Hellman

EasyAccess Authorization completes a programme Hellman began. He was among Hare's mentors and remains a friend, and the relationship shaped both the cryptographic direction of the work and the conviction that a problem widely assumed insoluble may simply be unsolved. Public-key cryptography is the canonical instance of that pattern in modern computing, and the architecture attempts the same move on the adjacent problem.

The Santa Fe Institute

Quantum Adaptive Systems Theory extends Complex Adaptive Systems Theory as developed at the Santa Fe Institute, and the connection is documented rather than claimed. A dated July 1998 itinerary records BIOS Group and the Institute, and the collaboration with Stuart Kauffman and the founders of the complexity programme produced the treatise from which the present theory descends. Kauffman's autocatalytic-closure result is one of the six importations, and it is imported as a mathematical result about networks rather than as a biological analogy, which is why it transfers.

The Institute's contribution is not only the result but the habit of mind. Complexity science as practised there treats the search for common structure across unrelated domains as ordinary work rather than as overreach, which is the disposition every importation in this document depends on. A researcher trained to stay within a discipline does not attempt six.

The E.O. Wilson Biodiversity Foundation

Hare served as Chairman of the Technical Advisory Board of the E.O. Wilson Foundation from 2006 to 2009. The Consilience Initiative he proposed in 2006 was offered as the only way to satisfy Wilson's TED Prize wish for a decentralised Encyclopedia of Life — an open and rich record of every species on the planet, their habitats, their environmental requirements, and the benefits and uses of both. That proposal specified seven services which the six domains of the Consilient Quantum Paradigm now deliver.

Wilson handed on the mission in person in 2013, and Hare gratefully accepted the charge. The Consilient Nature and Humanity Trust described in candidate seven carries the name of one of Wilson's greatest written contributions to the scientific enterprise, for that reason and in his honour.

The Foundation's board is worth quantifying, because it establishes the level at which the proposal was reviewed. The board alongside which Hare's Technical Advisory Board operated seated five Nobel recipients — four individual laureates in physiology or medicine, and one recognised through an organisation awarded the Peace Prize — together with a National Medal of Science holder, the two-time Pulitzer winner who chaired it, the president of Rockefeller University, the president of Memorial Sloan Kettering, the director of Cold Spring Harbor Laboratory, the director of the Earth Institute at Columbia, and the chief executives of several public companies. The Foundation's current Board of Advisors retains three Nobel laureates and much of the same composition.

These were working relationships, not proximity. Wilson personally introduced Hare to the Foundation's founding chief executive and asked that he be appointed Chairman of the Technical Advisory Board so that the two could work together on the Consilience Initiative and attend the December 2006 Google summit in that capacity. Hare worked directly with board members across the period. The relationship continues: the E.O. Wilson Biodiversity Foundation is today an affiliate of the EP3 Sustainable Markets Accelerator, and its President and Chief Executive is a member of it.

Individual names are not listed here because no written endorsement of this document has yet been sought from any of them. Seeking it is among the purposes for which the document was written.

The 2006 specification survives, and it is checkable. The Consilience Initiative documents dated February 2007 state the goal as a self-funding information ecosystem, and specify seven shared ecosystem services required to build it: data interchange across disparate systems and organisations; search, query, visualisation and analytics spanning heterogeneous data with filtering by reputation and source; integrated identity and reputation management allowing parties to verify each other across organisational boundaries; distributed, decentralised, fine-grained control of encryption, authentication and access; an anonymity service to de-identify and transform data to protect privacy; internet-based rights management making data ownership explicit; and information markets providing incentives to create, analyse, contribute and annotate.

Every one of the seven is delivered by a named domain of the Consilient Quantum Paradigm, and the phrase self-funding appears in the 2007 goal statement nineteen years before the architecture that makes it achievable. Those documents were written to solve a problem in biodiversity. What they specify is the general case, which is what Hare argued at the time it would have to be.

The universities

Three faculties shaped the method rather than the content, and each is visible in the results above.

- **The University of California, Berkeley — Industrial Engineering and Operations Research.** The training was in simulation and the modelling of dynamic systems, which is the formal foundation on which Adaptive Systems Analysis was later built. The disposition to treat an economy or an institution as a dynamic system to be modelled rather than a subject to be described begins there.
- **Stanford Graduate School of Business.** The analytical method originates in Andrew Grove's first course there and in the competitive analysis of the computing industry it taught — the observation that value migrates to whichever layer is missing. Applied at Microsoft in 1991, that method produced the identification of the missing layer as one for business process itself, which is the coordination layer the architecture now supplies.
- **Harvard University and Harvard Business School.** The first meeting with Wilson took place at Harvard in 1999, after Hare renamed his company Consilient in order to reach him, and the two agreed there to launch a foundation. The relationship that produced the 2013 commission begins in that room.

None of this is evidence that the results hold. Proximity to a laureate establishes nothing about the soundness of an extension of their work, and the extensions stand or fall on the same formal grounds as everything else in this document. What the mentorships, the collaborations and in some cases the brief interactions explain is how one person came to understand each of these domains well enough to unify them into an unbroken consilient ladder. That is a different question from whether the unification is correct, and it is one a reviewer is entitled to have answered.

Assessment

§16

What Would Have to Be True

A reviewer assessing these candidates is entitled to a statement of what would defeat them, and the twelve divide into three classes with different failure conditions.

Candidate	Class	What would defeat it
Coasean dissolution; Liquidity Pool; Spence resolution and inversion; Premium Multiple	Constructive economic claim	A demonstration that the frictions persist under the mechanism, that credibility still requires cost, or that the mechanism cannot operate at the scale claimed
Quantum Privacy Cells; Quantum Genome; EasyAccess Authorization	Foundational computational primitive	A demonstration that the primitive does not provide the properties claimed, or that prior work already provides them
Deployment-level alignment; compliance as a property of computation	Foundational computational primitive	A demonstration that a contained agent can formulate a prohibited action, that revocation paths are not independent, or that verification requires disclosure after all
Ownership as the answer to displacement	Realized-impact claim	Non-adoption, or attribution proving insufficiently granular to assign ownership in AI-produced value
The Trust; transition commitment devices	Realized-impact claim	Non-adoption, or settlement volumes materially below the conservative case, since both depend on flows that do not yet exist

The first two classes are checkable now and do not depend on adoption. Whether the frictions dissolve is a question about the mechanism, not about how many participants use it. Whether the primitives supply what is claimed is a question a specialist can settle by inspection. The third class is genuinely contingent, and is stated as contingent rather than asserted.

The paradigm itself carries a sharper test. Part I claims these six results follow from one body of theory rather than from six separate insights. That claim fails if the six formal correspondences set out in the companion provenance document do not hold — if Privacy Domains do not satisfy the definition of a manifold region carrying metric structure, if governance inheritance is not covariant transport, if the cascade does not satisfy the conditions distinguishing two-stage from single-stage assembly. Those are questions for differential geometers, physicists and biologists, they require no data and no deployment, and the author is not a party to how they resolve.

§17

Why Twelve and Not Fifty-Three

The source document assesses fifty-three candidates, and the reduction here is not a judgement that the others are weaker. It is a judgement about what a reviewer needs.

Fifty-three independent claims read as a catalogue, and a catalogue invites the reasonable suspicion that the criteria were chosen to admit as much as possible. Twelve claims, each traced to the same paradigm and each stated with its defeating condition, read as an argument. The remaining forty-one are available in the source document, and several are strong on their own terms — the Hayekian knowledge-aggregation result, the resolution of the sovereign-coordination trilemma, and the medical-infrastructure candidates in particular. They are omitted because they do not bear as directly on the claims the companion documents make, not because they are weaker.

The deeper reason for the reduction is the one Part I states. If the paradigm holds, the number of derivable results is not evidence of anything, because a genuine substrate change produces many. What is evidence is whether any one of them holds. Six is enough to establish the pattern and few enough that each can be assessed properly, which is the condition under which assessment is worth conducting at all.