Cheap Translation, Scarce Commitment

Standards for Software, Tokenization, and Finance in the Age of Machine-Mediated Interoperability

Zakaryae Boudi Jiuling Teng

Intelligence Economy Institute

July 2026

Abstract

Advanced artificial intelligence is approaching the capacity to inspect software interfaces, infer intent, translate between heterogeneous systems, detect ambiguity, and generate integration logic on demand. This raises a provocative question: do standards retain a function when machines can mediate interoperability dynamically? The provocation deserves arguments rather than dismissal, and the report therefore steelmans it before answering it. The answer is that the question dissolves once standards are decomposed into their constituent functions. Standards do two separable things. They compress the cost of compatibility, and they fix the locus of commitment. Advanced AI collapses the first function and cannot, by construction, supply the second, because commitment is a relation between accountable persons, not a property of representations. The report develops four scenarios in which standards weaken, strengthen, change role, or bifurcate by domain, and steelmans each. Using tokenized finance as the limiting case, it argues that the binding constraint on institutional tokenization has never been syntactic interoperability but the absence of shared institutional meaning around rights, controls, obligations, lifecycle events, finality, and accountability. It then analyses sixteen operational domains, from data mapping to model governance, distinguishing in each what AI can plausibly automate from what still requires standards, governance, legal authority, or market convention. Finally, it projects the future forms of standards: shared ontologies with formal semantics, executable specifications, conformance test suites, control-property standards, machine-readable law, agent mandate and negotiation protocols, evidence-grade audit logs, settlement and finality standards, and certification regimes for models and adapters. The central conclusion is that AI changes the economics of making and enforcing standards far more than it changes the need for having them. The finding is a redistribution of the standards function, not its retirement: standard-setting institutions emerge from the analysis with work that is scarcer and more consequential, not less. Expect fewer, deeper, more executable standards. Translation becomes cheap; commitment remains scarce.

Keywords: standards; interoperability; tokenization; financial market infrastructure; artificial intelligence; settlement finality; semantic interoperability; smart contracts; regulatory technology; formal verification; model governance.

Contents

1 Executive Thesis

The question this report addresses is usually framed as whether machines that can translate anything still need agreements about everything. Framed that way, the question answers itself, and answers itself wrongly. It treats a standard as a workaround for the absence of intelligence: a fixed schema tolerated only because no one could afford to interpret every counterparty's idiosyncratic format. On that view, once interpretation becomes nearly free, the workaround should retire. This report argues that the view rests on a category error, and that correcting the error yields a more interesting and more actionable forecast than either the abolitionist or the preservationist position. One point of framing should be fixed at the outset, because a report that steelmans the abolitionist case invites misreading: nothing in what follows concludes that standards, or the institutions that make them, are becoming ornamental. The conclusion is a redistribution of the standards function across its elements and layers, not its retirement, and the bodies that set standards end this analysis holding the scarcer half of the work.

A standard performs two functions that are usually fused in practice and must be separated in analysis. The first is *compression*: a standard reduces the cost of achieving compatibility between systems by replacing n^2 bilateral negotiations with n conformances to a common reference. The second is *commitment*: a standard fixes, in advance and for all relevant parties, what a message, an interface, a term, or an act shall be taken to mean, so that parties can rely on that meaning, build on it, price it, litigate over it, and be held to it. Compression is an engineering economy. Commitment is an institutional fact. Artificial intelligence attacks the first function with real force, because translation, inference of intent, reconciliation of mismatched schemas, and generation of adapters are precisely the tasks at which large models excel. It does not and cannot attack the second function, because commitment is not a property of representations that a sufficiently clever interpreter can recover. It is a relation among accountable persons, sustained by governance and, in finance, by law.

From this decomposition the report derives its central forecast. Standards will not disappear; they will be refactored along three migrations. First, standards migrate *up the semantic stack*: syntax is released to machines, while semantics is consolidated into ontologies precisely so that machine translation can be checked rather than trusted, and pragmatics, meaning the obligations, finality, and recourse that attach to messages, is hardened into law-adjacent artifacts. Second, standards migrate *from prose to executable form*: the normative core of a standard shifts from a document that humans read to a reference implementation, a conformance test suite, a set of formal properties, and machine-readable rules that both humans and machines can run. The standard of the future is a test you can fail. Third, standards migrate *from the edges to the core*: peripheral, low-stakes integration becomes an AI-mediated free-for-all, which is broadly desirable, while the commitment core of markets, namely settlement, custody, collateral, compliance determinations, and lifecycle events with distributive consequences, becomes more rigidly and more formally specified than today, because the surrounding automation raises the speed and coupling of the system.

Tokenized finance is the limiting case that makes these dynamics visible. The report's answer to its own focal question is direct: in tokenized finance, the binding constraint is not, and has not been, syntactic interoperability. Bridges, APIs, adapters, and message translators are abundant and cheap. The constraint is the absence of shared institutional meaning about what a token is, what rights it carries, what controls govern its movement, when a transfer is final, who keeps the authoritative register, what happens on default, and who is liable when any of these answers is contested. Evidence for this claim is presented in Section 7, including the revealed behaviour of institutions, which invested first in legal definitions, taxonomies, and control frameworks rather than in translators, and the failure taxonomy of past interoperability programmes, in which delivered technical platforms waited a decade for the harmonisation of market practice.

Two further conclusions follow and are defended at length. The first concerns adversarial interpretation. Interpretation between counterparties with adverse economic interests is not a translation problem that scale solves; it is a conflict that authority settles. Two highly capable models, each advocating for its principal, will disagree more precisely, not less. Ambiguity detection is a capability; ambiguity resolution is an office. Markets therefore need, and will continue to need, reference points whose interpretation is fixed by a body with the authority to bind, whether that body is a standards organisation, a determinations committee, a market infrastructure rulebook, a regulator, or a court. The second concerns determinism. Settlement finality is a legal construct that requires a definite moment at which a transfer becomes unconditional and irrevocable. A process whose state transitions depend on probabilistic inference cannot supply that moment. Probabilistic systems can advise on state transitions; they cannot constitute them. The architectural consequence, developed in Sections 5 and 9, is a pattern this report calls *negotiate, freeze, execute*: AI may negotiate, interpret, and propose upstream, but its output must compile into a deterministic, certified, replayable artifact before it touches the commitment path.

Finally, the report inverts the usual framing of AI's effect on standardisation. The historical bottleneck of standards bodies has been drafting labour and slow consensus. AI compresses drafting, formalisation, test generation, and conformance monitoring toward negligible cost, leaving consensus and authority as the residual. That residual was always the point. When drafting becomes free, agreement becomes the entire cost of a standard, and the comparative advantage of standards institutions shifts from producing documents to conferring bindingness. The practical implications for founders, regulators, financial institutions, and standards bodies are set out in Section 10, and the synthesis, a layer-by-layer allocation of tasks among AI, standards, and governance, together with ten defensible conclusions, closes the report.

2 Why the Question Is Real

2.1 The capability trajectory

It would be a mistake to treat the premise of this report as science fiction, and an equal mistake to treat it as accomplished fact. As a matter of evidence, contemporary AI systems already perform, in production settings, a large share of the tasks that historically justified syntactic standards. They read undocumented interfaces and produce working client code. They map fields between mismatched schemas with high accuracy when given examples. They translate legacy message formats into modern ones, including the extraction of structured meaning from free-text fields that were never designed to be parsed. They draft integration logic, reconcile data discrepancies across systems, summarise contractual documents, and flag inconsistencies between a document and a data record. In financial operations specifically, machine assistance is now routine in reconciliation break investigation, in mapping proprietary formats to ISO 20022 messages, in extracting terms from credit agreements, and in parsing unstructured corporate action announcements, a domain whose data quality problems have resisted three decades of standardisation effort.

By inference from the current trajectory, it is reasonable to expect systems that can do these things with fewer examples, across longer contexts, with tool use, and in agentic loops that test their own output against live endpoints. The scenario sketched in the introduction to this report, in which an agent inspects two systems, infers the intent behind their interfaces, detects the ambiguities, negotiates a working mapping with a counterpart agent, and generates a running adapter, is an extrapolation, not a fantasy. Where this report relies on that extrapolation rather than on observed capability, it says so explicitly.

2.2 The economic logic of the challenge

The challenge to standards is at bottom an economic argument, and it deserves to be stated in its strongest form before it is examined. Standards emerged as a response to costly coordination. The classical economics of compatibility (Katz and Shapiro, 1985; Farrell and Saloner, 1985; David and Greenstein, 1990) treats a standard as a device for internalising network externalities: each party's adoption raises the value of adoption for others, and a common specification prevents the market from fragmenting into islands. But the same literature is explicit that standardisation has costs, including the risk of converging on an inferior specification (David, 1985), the suppression of variety that would otherwise generate information about which designs are best (Farrell and Saloner, 1986), and the creation of installed bases that resist beneficial change. The historical bargain was that these costs were worth paying because bilateral, bespoke integration was prohibitively expensive at market scale.

If AI reduces the cost of bespoke integration by one or two orders of magnitude, the terms of that bargain move. The option value of postponing convergence rises. The penalty for heterogeneity falls. The argument that a thousand incompatible systems impose a deadweight n^2 translation burden weakens if translation is generated on demand at negligible marginal cost. It is intellectually lazy to wave this away with the observation that standards are about more than translation. The correct response, which this report attempts, is to identify precisely which functions of standards are translation-like, and therefore exposed to AI substitution, and which are not.

2.3 Why finance is the right test bed

Much of commercial software integration tolerates interpretive failure, because errors there are typically reversible and their costs idiosyncratic: a dashboard shows stale data, a webhook is retried, a user re-enters a form. The tolerance is not universal, and avionics, medical devices, and industrial control stand well outside it; the point is narrower, that where integration errors are cheap to reverse, machine-mediated improvisation is at its most defensible. Finance sits at the opposite pole of commercial computing: it is the domain where the consequences of interpretive failure are maximised and where the institutional machinery for preventing, allocating, and adjudicating those consequences is most developed. Payments, securities settlement, custody, collateral, derivatives, and fund administration involve enforceable obligations, regulated processes, systemic risk, requirements of auditability and finality, and dense webs of liability. The international regulatory community has codified what infrastructure in this domain must guarantee, including a well-founded legal basis, clear and certain final settlement, and rigorous operational risk management (CPMI-IOSCO, 2012). If the abolitionist case for AI-mediated interoperability holds anywhere, it should hold least here; if the preservationist case holds anywhere, it should hold most here. Tokenization sharpens the test further, because it is the live construction site where new instruments, new infrastructures, new law, and new standards are being built simultaneously, and where the question of what must be standardised is being answered in real time by legislatures, central banks, standards bodies, and protocol developers.

2.4 A note on epistemic status

This report is an exercise in institutional analysis, not prophecy, and it tries to keep three registers separate. *Evidence* refers to observed facts: what standards exist, what they cost, how past interoperability programmes fared, what AI systems demonstrably do today, what law currently provides. *Inference* refers to conclusions drawn from evidence via mechanisms the report makes explicit, such as the incentive analysis of adversarial interpretation or the base-rate arithmetic of error at scale. *Speculation* refers to projections about capabilities or institutional forms that do not yet exist, including several of the future standard forms in Section 9. Where the register shifts, the text says so. The reader should discount the three registers differently, and the report's conclusions are constructed to survive even if the most speculative capability claims are realised in full, which is the scenario least favourable to standards.

3 What Standards Actually Do

Any forecast about the fate of standards must begin from an accurate account of what they do. The popular account, that standards exist so that systems can interoperate, is true but radically incomplete, and its incompleteness is exactly what makes the AI challenge seem stronger than it is.

3.1 What a standard is: definition and elements

Because the argument that follows turns on which functions of standards survive, the term itself must be fixed first, and fixed broadly enough to prevent a misunderstanding that distorts much of the current debate. As used throughout this report, a standard is a reference fixed in advance of the transactions that use it, maintained by an identifiable steward, and intended for repeated use by parties who did not author it, such that conformity with it can be assessed. Seven elements recur across the population, and they supply the analytical vocabulary for everything that follows: (i) a *subject*, the thing being fixed, which may be a format, a meaning, a process, a property, or a required conduct; (ii) *normative content*, the requirements and conformance criteria that distinguish compliant from non-compliant; (iii) an *artifact form*, historically prose, increasingly schema, code, and test suite; (iv) a *steward*, the body with authority over the reference, whether a formal standards organisation, an industry consortium, a market infrastructure, a regulator, or a legislature; (v) an *adoption mechanism*, ranging from voluntary uptake through contractual incorporation to regulatory mandate and network compulsion; (vi) a *binding force*, running from social convention through contract to statute; and (vii) a *change regime*, the governed procedure by which the reference itself may be altered.

Two consequences of this definition matter for the whole report. First, interface protocols are one species of standard, not the genus. Rulebooks, definitions booklets, statutory technical rules, accounting policies, identifier schemes, and documented market practice all satisfy the definition, and an analysis that silently equates standards with interoperability protocols will misjudge the AI era, because, as Section 6 develops, AI bears on the species very differently. Second, the elements can be affected independently. A technology that transforms the artifact form and the economics of drafting, as AI does, need not touch the steward, the binding force, or the change regime at all, and much of the confusion in the current debate comes from inferring the fate of the whole institution from the disruption of one element.

3.2 Six functions

Observation of standards in software and finance suggests at least six distinct functions, which different standards combine in different proportions.

Coordination. A standard is a focal point in the sense of Schelling (1960): it lets many parties converge on one of many possible equilibria without pairwise negotiation. The value here is not that the chosen point is best but that it is common. This is the function most exposed to AI substitution, because machine negotiation makes pairwise coordination cheap.

Compression. Related but distinct: a standard replaces n^2 mappings with n conformances, and replaces continuous renegotiation with a one-time investment. This too is exposed to AI, with an important reservation developed in Section 5: compression also compresses the *verification* problem, and verification does not become free just because generation does.

Semantic fixation. A standard fixes what terms mean. ISO 4217 fixes what `EUR` denotes; the ISDA Definitions fix what a Failure to Pay is; Article 12 of the amended Uniform Commercial Code fixes what it means to have control of a controllable electronic record (Uniform Law Commission, 2022). Semantic fixation is not translation. Translation maps one representation to another; fixation establishes the reference against which any mapping is right or wrong. Without fixation, translation quality is undefined.

Commitment and reliance. By conforming to a standard, a party binds itself in advance to interpret and to act in a specified way, which allows counterparties to rely, to invest, and to price. The economic content of a standard is often precisely this self-binding: a settlement system's rulebook is valuable because participants cannot deviate from it unilaterally, not because deviation would be technically difficult.

Liability allocation and verification substrate. Standards allocate responsibility ex ante. Conformance functions as a safe harbour and as evidence of due care; certification against a standard is how auditors, insurers, and supervisors convert an open-ended question, was this system adequate, into a decidable one, did it conform. Audit, assurance, and certification are only possible against a stable reference. A moving specification cannot be audited, only observed.

Institutional memory. Standards encode lessons paid for in losses. Payment-versus-payment settlement in foreign exchange encodes the 1974 Herstatt failure (Galati, 2002). The Legal Entity Identifier encodes the discovery, in 2008, that no one could aggregate exposures to Lehman Brothers across its hundreds of legal entities (FSB, 2012). Risk data aggregation standards encode the same discovery inside banks (BCBS, 2013). To ask whether AI makes such standards unnecessary is to ask whether intelligence substitutes for memory, and the honest answer is only if the intelligence is instructed by the memory, which is to say, only if the standard persists in some form.

The pattern to hold onto is this: the first two functions concern the *cost* of compatibility, and AI attacks them directly. The last four concern the *authority* behind meaning and conduct, and AI does not touch them, because they are not information-processing problems.

3.3 Syntax, semantics, pragmatics

A second decomposition cuts across the first. Borrowing from the philosophy of language, any communicative standard operates at three levels. *Syntax* governs form: field layouts, encodings, data types, message grammars. *Semantics* governs meaning: what a field denotes, what an event type is, how a business concept maps to data. *Pragmatics* governs force: what sending or receiving a message commits the parties to, what acting on it does to legal state, what recourse exists when expectation and reality diverge.

The distinction matters because AI's leverage is sharply unequal across the three levels. Syntax is where models are strongest: parsing, converting, and generating well-formed structures is close to a solved problem. Semantics is where models are strong but unanchored: a model can propose that field `X` in system A corresponds to concept `Y` in system B, and the proposal will usually be plausible, but plausibility is not correctness, and correctness is only defined relative to a semantic reference that someone fixed. Pragmatics is where models have no purchase at all, because pragmatic facts are not inferable from representations. Whether a transfer is final is not a property of the transfer message; it is a property of a rulebook, a statute, and a designation (European Union, 1998). A model can report what the rulebook says. It cannot make the rulebook true.

3.4 The families of standards

It is also worth cataloguing what kinds of artifact the word standard covers, because the AI challenge lands differently on each. *API and session standards* (FIX session layer, REST conventions) govern how systems converse. *Schema and message standards* (ISO 20022, FpML) govern the structure of what they say. *Ontologies and taxonomies* (the Financial Industry Business Ontology, the ISDA and FINOS Common Domain Model's product and event model, the ICMA Bond Data Taxonomy) govern the concepts behind the structures (EDM Council, n.d.; FINOS and ISDA, n.d.; ICMA, 2023). *Identifier standards* (LEI, ISIN, the ISO 24165 Digital Token Identifier) give the market shared names for shared things (ISO, 2021). *Legal definition standards* (ISDA Definitions booklets, statutory definitions of control, transfer order, or electronic money token) fix meaning with legal force. *Process standards* (corporate action market practice, collateral management rulebooks such as the Eurosystem's SCoRE standards) govern sequences of obligations (ECB AMI-SeCo, 2020). *Execution and settlement standards* (delivery-versus-payment models, finality rules, default management procedures) govern the commitment core itself. As a rough law, AI's substitution potential declines monotonically down this list, and Section 6 will argue that the value of standardisation correspondingly rises down it.

3.5 Static versus adaptive; bilateral versus market-wide

Two further axes complete the analytic frame. Standards differ in their *change regime*. Some are effectively frozen; some evolve on governed release cycles, as ISO 20022 does through its annual maintenance process; some are living standards in the sense pioneered by the WHATWG for HTML, continuously updated under version control with conformance defined by tests rather than by dated documents. The choice of change regime is itself a governance decision, and one of this report's conclusions is that AI shifts the optimal regime toward governed velocity: faster than committee cycles, slower and more accountable than model updates.

Standards also differ in *scope*. Bilateral interoperability requires only that two parties agree; market-wide interoperability requires that all parties share not merely pairwise compatibility but a common state. The difference is fundamental and is repeatedly elided in the AI-abolitionist argument. Translation is bilateral; truth in a market is multilateral. A securities market does not need every pair of participants to understand each other; it needs a single answer to the question of who owns what, an answer that survives the failure of any participant. Pairwise consistency among n^2 AI-negotiated mappings does not compose into a shared state; a register does. This is why the infrastructures at the core of finance, central securities depositories, payment systems, trade repositories, are not translation layers but authoritative records, and why their standards are not conveniences but constitutions.

3.6 Precedents and their lessons

The report will draw throughout on a set of precedents, introduced here with the lesson each contributes.

TCP/IP versus OSI. The internet's protocol suite, developed through rough consensus and running code, defeated the more complete, more official OSI stack (Abbate, 1999; Russell, 2014). Lesson: de facto standards anchored in working implementations beat de jure standards anchored in documents. This is a point *for* the executable-standards thesis of Section 6, not against standards as such: TCP/IP was still a standard, and its interoperability was proven in plugfests, the ancestors of conformance test suites.

The robustness principle and its revision. Postel's injunction to be conservative in what you send and liberal in what you accept enabled the early internet to grow across imperfect implementations

(Postel, 1981). Four decades later, the IETF's considered position is that liberal acceptance breeds ossification, divergence, and security vulnerability, and that protocols are healthier when deviations fail fast and visibly (Thomson and Schinazi, 2023). The formal-language security literature reached the same conclusion from the exploit side: parsers that accept ill-defined inputs become weird machines that attackers program (Bratus et al., 2011). Lesson: tolerant interpretation at scale is not a free lunch, and AI is Postel's principle with a research budget. The internet already ran this experiment.

ISO 15022 to ISO 20022. The securities industry has spent more than two decades migrating between its own message generations, and ISO 20022's deeper contribution is not its XML syntax but its methodology: a business model and data dictionary from which syntaxes are derived (ISO, 2013). Yet identical syntax did not produce identical usage; cross-border payments required an additional layer of market practice guidelines to constrain how the standard is used. Lesson: syntax does not carry semantics, and even a semantic repository does not carry pragmatics; usage must be governed.

FIX. The FIX protocol standardised the session and message layer of electronic trading and then discovered that every bilateral connection still required negotiated rules of engagement specifying which fields, values, and workflows each counterparty actually supports. Lesson: standards and negotiated profiles coexist; the sustainable pattern is negotiation *within* a certified envelope, a pattern Section 9 generalises to AI agents.

FpML to the Common Domain Model. FpML standardised the representation of derivatives; the CDM was created because representation proved insufficient, and what the market needed was a standard for the *events and functions* of the trade lifecycle, expressed in executable form (FINOS and ISDA, n.d.). Lesson: the frontier of financial standardisation was already moving from data to process and from prose to code before modern AI arrived.

ACTUS. The Algorithmic Contract Types Unified Standards project demonstrated that the cash-flow obligations of most financial contracts can be reduced to a small set of contract types with deterministic, executable payoff functions (Brammertz et al., 2009; ACTUS, n.d.). Lesson: semantics can be made computable, and when it is, valuation, risk, and reporting cease to be interpretive exercises. ACTUS is the existence proof for the executable-standard category this report projects forward.

ERC token standards. The ERC-20 interface, a few function signatures, created composability at a scale no committee ever achieved, and simultaneously demonstrated the cost of standardising interface without semantics: non-conforming return values, fee-on-transfer and rebasing tokens that silently violated integrator assumptions, and approval race conditions became endemic (Vogelsteller and Buterin, 2015). ERC-4626 standardised vault accounting and still required post-hoc property guidance on rounding direction and inflation attacks (Santoro et al., 2022). Lesson: interface standards without property standards compose bugs as efficiently as they compose features.

Package ecosystems. Semantic versioning is a machine-actionable compatibility promise, a social contract enforced by tooling. The 2016 removal of a trivial package, left-pad, broke builds across the industry and revealed the fragility of unpinned dependency graphs; lockfiles restored determinism by freezing resolution at a moment in time; software bills of materials emerged as provenance standards after supply-chain compromises. Lesson: when automated resolution became too clever to reproduce, the ecosystem responded by *pinning meaning at the point of commitment*. The reader is invited to hold that phrase in mind; it recurs as this report's architectural refrain.

Post-crisis identifiers. After 2008, regulators mandated derivative trade reporting, and the reports proved unaggregatable: without common transaction and product identifiers, supervisors could not match the two sides of the same trade. The remedy was not smarter parsing but the global harmonisation of critical data elements, the UTI, and the UPI (CPMI-IOSCO, 2018). Lesson: supervision at market scale requires shared reference points; it cannot be reconstructed by inference over heterogeneous data, because the question of whether two records denote the same trade is precisely what inference cannot settle and identifiers do.

4 The Strongest Case That AI Weakens the Need for Standards

This section states the abolitionist case as its proponents should state it, without the strawman framing that standards advocates prefer to attack. The case is not that meaning is unnecessary. It is that *ex ante, committee-produced, document-shaped* agreements are an increasingly inefficient technology for producing the coordination that markets need, and that machine mediation supplies a superior technology for a growing share of that need. Seven arguments carry the case.

4.1 The collapse of the cost asymmetry

The economic foundation of standardisation is a cost asymmetry: bespoke integration is expensive, so paying the fixed costs of convergence, committee time, migration projects, retesting, is worth it. Estimates of what large financial institutions spend on integration, reconciliation, and data remediation vary, but no practitioner disputes that glue work consumes a substantial fraction of technology budgets and that much of it is semantically trivial: moving values between representations that differ for historical rather than substantive reasons. If AI reduces the marginal cost of a working, tested adapter from months of specialist effort to hours of supervised generation, the asymmetry that justified convergence weakens across a large class of integrations. The standard economics of compatibility does not say standardise always; it says standardise when coordination is costlier than variety is valuable (Farrell and Saloner, 1985; David and Greenstein, 1990). AI moves the threshold, and honesty requires admitting that it moves it a long way.

4.2 Standards freeze bad abstractions, and semantics freezes hardest

Every standard is a bet that today's abstraction will fit tomorrow's world, and the historical record of such bets is mixed. QWERTY is the folkloric example (David, 1985); the financial examples are more instructive. ISO 15022 froze a message model that the industry then spent twenty years co-existing with its successor, running dual infrastructures at enormous cost. FIX accumulated thousands of tags whose meanings drifted across venues. The deeper point cuts against the standards-maximalist reading of this very report: if the future of standards is semantic, then the lock-in risk migrates to semantics, and semantic lock-in is worse than syntactic lock-in, because an ontology that misdescribes the world propagates its misdescription into every system that conforms to it, and unwinding a concept is harder than unwinding a field layout. A world of AI translators, by contrast, holds abstractions loosely: each mapping is regenerable, each integration is a hypothesis that can be revised when the world changes. Dynamic translation is, on this argument, a form of institutional humility that premature convergence forecloses.

4.3 Standards lag, and the lag is getting more expensive

Formal standardisation runs on committee time: multi-year cycles of proposal, ballot, comment, and publication, followed by multi-year adoption. Product cycles in AI-adjacent software run on weeks. When the two clocks diverge this far, the standard published at the end of the cycle describes a world that no longer exists. The history of web services is the cautionary tale: the WS-* stack produced dozens of formally impeccable specifications that the market abandoned wholesale for informal REST and JSON conventions, which were never standardised in the ISO sense at all and interoperated better in practice. The lesson generalises: where change is fast, the de facto conventions of running systems, discoverable and learnable by machines, may coordinate better than de jure documents, and AI dramatically improves the discoverability of de facto conventions. On this view, the model does not replace the standard; it replaces the *committee*, by learning the emergent convention directly from the corpus of live behaviour.

4.4 Capture, politics, and the lowest common denominator

Standards bodies are political economies. National delegations pursue industrial policy; incumbents pursue moats; vendors pursue standard-essential positions; and consensus processes systematically produce lowest-common-denominator outcomes, because the veto players are those with installed bases to protect. The result is familiar to every practitioner: standards with optional everything, extension points that swallow the specification, and profiles of profiles. A machine-mediated regime removes the choke point. No delegation can capture a translation layer that any party can regenerate; no incumbent can slow-roll an interface that an agent reverse-engineers in an afternoon. There is a genuine liberal argument here: standards concentrate power in whoever governs the standard, and dissolving that concentration into a competitive market of interpreters is, other things equal, a gain.

4.5 Diversity is resilience; standards are monoculture

A market in which everyone parses the same format with the same reference implementation fails together. Heartbleed and log4shell were catastrophic precisely because OpenSSL and log4j were de facto standards: single implementations embedded everywhere, so a single defect became a systemic event. Heterogeneous, independently generated adapters do not share failure modes; an input that confuses one is unlikely to confuse all. On this argument, the standardisers have the systemic-risk analysis backwards: it is convergence, not diversity, that concentrates fragility, and a population of diverse machine interpreters is a portfolio where a monoculture of conformant parsers is a single position. (Section 5 will contest the empirical premise that AI-mediated interpretation is in fact diverse, but the argument as stated is coherent and historically grounded.)

4.6 Dynamic translation as discovery

Standardisation is often premature because no one yet knows the right abstraction, and the way markets learn the right abstraction is by running variety. Screen scraping preceded and effectively designed open banking APIs: the aggregators' bilateral hacks revealed what data users actually wanted, and the eventual standards codified the discovered practice. On this reading, AI-generated adapters are not the enemy of standards but their research programme: a vast, cheap, parallel search over the space of workable mappings, whose observed convergences become the draft of the next standard. Suppressing that search with early convergence destroys the information it would have produced. The strong form of the argument says: let a thousand adapters bloom, instrument them, and standardise the survivors.

4.7 The sharpest version: the model is the standard

The most sophisticated abolitionist argument concedes that shared meaning is indispensable and claims that AI supplies it by a different route. Large models trained on overlapping corpora converge on substantially shared latent semantics; two competent models asked what a coupon payment date is will agree, not because a committee fixed the answer but because the training distribution did. On this view, the foundation model *is* the standard: a continuously updated, empirically grounded, universally accessible semantic reference, produced without ballots and adopted without migrations. Interoperability through shared weights is standardisation without standardisation, and it scales with capability rather than with committee throughput.

The claim can be given operational teeth, and its strongest advocates give it exactly that. A pinned model version together with a system prompt and a handful of exemplars is, on this view, an executable specification: the exemplars are the schema, the embedding space is the interlingua, and the pair of model and prompt can be versioned, hashed, and distributed exactly as reference implementations are. Concentration completes the argument. If most institutions integrate through the same two or three frontier models, coordination happens through the weights whether anyone intends it or not: the model becomes the focal point of Section 3.2, selected not by ballot but by capability rankings and procurement gravity, and adopted at the speed of an API key rather than of a migration programme. On settled vocabulary the convergence claim is plausible and, as far as it can be observed, largely holds; the contested and the novel are another matter, and the reply below takes them up. As inference from the trajectory, expect the informal version of this function to grow regardless of what institutions decide: models will operate as the market's default glossary, its first-pass translator, and its discovery engine for latent conventions, in exactly the territory that Section 4.8 marks as conceded.

This is the argument the rest of the report takes most seriously, because it is the only version that addresses semantic fixation rather than mere translation. Note, however, what the argument has done: it has not eliminated the standard; it has *relocated* it, into an artifact that lacks every property that historically made references institutionally trustworthy. Running the elements of Section 3.1 across the model-as-standard makes the deficit precise. The steward is a commercial laboratory with no duty to the reference's users. The change regime is silent: weights are retrained and behaviour shifts without a change log, a comment period, or a deprecation notice, so the one property a reference must have, that it stays put while parties rely on it, is structurally absent; a standard that learns is a moving target. The normative content is unverifiable except behaviourally, which means anyone needing assurance must write test suites against the model, thereby reinventing conformance standards one level up. The binding force is disclaimed in the terms of service, where a standards body's text carries an author of record and a liability trail. The adoption mechanism is exit-hostile: a proprietary semantic reference cannot be forked, archived, or independently maintained, so dependence arrives without the membership rights that standards governance exists to confer. And the fixation the model does supply is fixation of the training distribution's past: it standardises yesterday's usage as of a cutoff, which is the lag problem of Section 4.3 rebuilt without the amendment procedure that makes lag correctable. A de facto semantic standard without a change log, a conformance test, a governance process, or a liability regime is not the end of standardisation. It is standardisation with every institutional safeguard removed.

The residue of the argument is nonetheless real and should be banked rather than argued away. Models will function as informal semantic conventions and as instruments of convention discovery, and in low-stakes territory they will displace written standards outright, as the next subsection concedes. But wherever the reference must be relied on against an adverse party or an auditor, the market's observable response to the model-as-standard is the opposite of abolition: it is the demand that the missing safeguards be supplied, through version pinning, evaluation suites, change notification, provenance, and certification, which is to say, through standards *for* models. Section 5 develops the consequences; Section 9 develops the remedy.

4.8 Where the abolitionist case simply wins

Intellectual honesty requires marking the territory this case captures outright, because it is large. Internal integration within a single legal entity, where one party bears all interpretive risk and can correct errors administratively, needs no external standard; AI-mediated glue is strictly superior to the undocumented bespoke glue it replaces. The long tail of SaaS and data integrations, prototypes, analytics pipelines, legacy migration, one-off data onboarding, and the reading of dying formats (an AI that fluently parses thirty years of mainframe extracts is a gift to every bank with a merger history) are all domains where the stakes are recoverable errors and the alternative to machine mediation was never a standard but a backlog. There is also a genuine emergency case: in a resolution weekend or an operational crisis, the counterparty's systems must be understood *now*, and an interpreter that works tonight beats a harmonisation programme that completes in five years. Even here, though, the institutional response to the last crisis is instructive: after 2008, authorities did not conclude that ad hoc interpretation under fire was acceptable; they standardised the data that must be producible in resolution. Mature systems standardise even their emergencies.

The abolitionist case, taken whole, is strongest exactly where this report's six-function analysis predicts: where the coordination and compression functions dominate and the commitment, liability, and memory functions are absent. The question is what happens where they are present. That is the subject of the next section.

5 The Strongest Case That AI Strengthens the Need for Standards

The preservationist case is routinely stated badly, as an appeal to caution or to the cultural authority of existing institutions. Stated properly, it is a set of structural arguments about what markets are, and each argument gets *stronger*, not weaker, as AI capability rises. That inversion is the section's through-line: most of what follows shows that the more capable the machines, the more valuable the fixed points.

5.1 Probabilistic inference cannot constitute deterministic settlement

Settlement finality is not an engineering property; it is a legal construct. Under the European settlement finality regime, and under equivalent provisions elsewhere, there is a designated moment at which a transfer order becomes irrevocable and unconditional, protected even against the insolvency of the sender (European Union, 1998; CPMI-IOSCO, 2012). Everything in wholesale finance is built on the existence of that moment: netting, collateral, intraday credit, default management. The moment must be definite, because rights vest at it, and it must be common knowledge, because parties on both sides must be able to act on it simultaneously.

A process whose state transitions depend on probabilistic interpretation cannot supply such a moment. This is not a complaint about current model quality; it is a statement about kinds. A settlement system that is extremely confident it settled has not settled; it has produced a well-calibrated opinion. Determinism at the commitment boundary is not a limitation of settlement systems that better technology will relax. It is their product. The architectural consequence is not that AI is excluded from settlement, but that its role is bounded: probabilistic cognition may surround the deterministic core, proposing, checking, and explaining, while the core itself, the state machine that vests rights, must be specified, certified, and replayable. Specified, certified, and replayable is a description of a standard.

5.2 Adverse interests: interpretation is a conflict, not a computation

The abolitionist case models interpretation as a cooperative task: two systems want to understand each other, and a sufficiently capable interpreter closes the gap. Finance is not like that. Counterparties have adverse economic interests, and interpretive questions in finance are distributive: whether a credit event occurred, whether a transfer restriction applied, whether a valuation input was permissible, decides who pays whom. In a distributive dispute, each party's AI will, rationally and without any malfunction, construct the most defensible reading favourable to its principal. Greater capability does not converge the readings; it improves both briefs. Two very capable models with adverse mandates disagree more precisely, not less.

Markets solved this problem long before AI, and the solution is instructive because it is not a translation layer. It is an *office*: a body whose determination binds. The ISDA Determinations Committees exist because whether a Failure to Pay occurred is frequently not inferable from public facts to the satisfaction of both sides, and someone must decide with finality so that a hundred thousand contracts can settle uniformly. Exchanges maintain clearly-erroneous-execution rules because even mistakes must have standardised consequences. Ambiguity detection, at which AI will excel, is a capability. Ambiguity resolution is an office. AI will surface more ambiguity, faster, in more contracts, than human review ever did, and every surfaced ambiguity is a dispute in waiting. A world of machine-scale ambiguity detection therefore needs *more* authoritative resolution capacity, not less: more precise definitions, more executable rules whose outputs are decidable, and more determination mechanisms whose findings bind. This is also the answer to the negotiated-compatibility framing: negotiated compatibility governs how two systems talk; enforceable finality governs what the talking did to legal state; and no volume of the former substitutes for the latter.

5.3 Error correction versus legal accountability

In ordinary software, error correction is a virtue: detect the mismatch, patch the mapping, replay the job. In finance, correcting an executed error is not a patch; it is an *unwind*, and unwinds are governed events with distributive consequences. Reversing a settled transfer disturbs finality; adjusting a struck NAV triggers compensation frameworks; amending a submitted regulatory report is itself a reportable event. The relevant question when AI-mediated interpretation fails is therefore not can the error be fixed but *who bears it, under what procedure, with what effect on third parties who relied.* Those questions are answered by rulebooks, by contracts, and by law, none of which can be generated at the moment of failure, because their entire function is to have been agreed before it. The better AI gets at acting, the more acting it does, and the more the binding constraint becomes the pre-agreed allocation of the consequences of action. Self-correction is a property of systems; accountability is a property of institutions.

5.4 AI reasoning versus audit trails

Regulated finance runs on reconstructability. A supervisor, an auditor, or a court must be able to determine, years later, why a decision was taken and whether the process followed was the process required. Model-mediated interpretation strains this requirement in three specific ways. First, non-determinism: the same inputs can yield different outputs across runs, so the decision cannot be replayed. Second, version drift: the model that made the decision no longer exists in the fleet six months later, so the decision cannot be re-derived. Third, context dependence: the output depended on a prompt assembled from live data whose state was not preserved. None of these is fatal, and all three have remedies: pin model versions, log prompts and outputs, fix decoding parameters, archive weights, hash the whole. But observe what the remedies are. They are record-keeping and reproducibility *standards* for models, of exactly the kind that model risk management guidance has required of quantitative models for over a decade (Federal Reserve and OCC, 2011) and that the European AI regulation now requires of high-risk systems in the form of logging and traceability obligations (European Union, 2024a). The audit-trail argument does not show that AI cannot be used; it shows that using it manufactures fresh demand for standards, aimed at the machines themselves. The evidence-grade audit log, hash-chained, signed, model-versioned, is a standard form this report expects to become as fundamental to AI-era finance as the trade confirmation was to the paper era.

5.5 Liability, certification, and the procurement reality

When an AI-generated adapter misinterprets a payment instruction and value is lost, existing law does not recognise the translator erred as a category that absolves anyone. Loss will be allocated across the deployer, the counterparty, and possibly the vendor, along lines fixed by contract, by negligence doctrine, and by regulatory expectations about outsourcing and operational resilience. Institutions know this, which is why the adoption question is not decided by capability demonstrations. An institution can procure against a certified interface: the specification defines the promise, the certification evidences conformance, the contract allocates residual risk, the insurer prices it, the auditor attests it. An institution cannot procure the model will figure it out, because there is no specification to warrant, no conformance to certify, no stable object for the contract to describe. This is the concrete meaning of the contrast between AI-generated adapters and certified interfaces: the artifact may be identical code, but only one of them is a *legal object*. The likely and desirable synthesis is not either-or: adapters will be machine-generated and then certified, with the certification, conformance evidence against a published test suite, doing the institutional work. That synthesis presupposes the test suite, which presupposes the standard. Note also the public-trust dimension: private parties are free to run uncertified interpretation between themselves and eat the losses; the case for mandatory reference points arises exactly where third parties, clients, markets, taxpayers, are exposed to interpretations they never agreed to. Private-sector flexibility and public trust are not opposites to be balanced rhetorically; they are jurisdictions to be demarcated, and the demarcation line is externalised risk.

5.6 Model drift, hallucination, and adversarial manipulation

Three failure modes distinguish machine interpretation from the deterministic parsing it would replace, and each has systemic character. *Hallucination and base rates:* an interpreter that errs once in a hundred thousand messages is extraordinary by human standards and produces tens of thousands of errors a day at payment-system scale; when the erring interpreter is embedded in an autonomous loop, errors execute rather than queue for review. *Drift:* model updates silently change interpretive behaviour across every integration simultaneously, which is a change-management event of market-wide scope occurring on a vendor's schedule; the financial industry's hardest-won operational lesson, that uncontrolled deployment is how a firm loses hundreds of millions in forty-five minutes, as Knight Capital did in 2012, applies with full force to weights. *Adversarial manipulation:* an interpreter that reads free text is a machine an adversary can program with inputs; instruction-injection through the unstructured fields of financial messages is the modern instance of the formal-language security thesis that every permissive parser is an attack surface (Bratus et al., 2011). The mitigation for all three is the same and is instructive: constrain the input language, validate against strict schemas, canonicalise before commitment, and treat interpretive components as models under governance. Strict schemas, canonical forms, and governed models are standards. The security case for standards was strong before AI; adversarial AI makes it decisive, because the attacker now has the same interpreter the defender does.

There is also a monoculture rejoinder to Section 4.5. The premise that machine interpretation is diverse is empirically doubtful: the market is converging on a small number of foundation models, so interpretive errors will be *correlated* across institutions in a way that a population of independently written parsers never was. Correlated interpretation error across a market is a financial-stability exposure of a new kind, kin to the correlated model behaviour behind the 1987 portfolio-insurance cascade and the 2010 flash crash (CFTC and SEC, 2010), and supervisors currently have no instrument for measuring it. The realistic near future is the worst of both framings: monoculture at the model layer, heterogeneity at the behaviour layer. Standards at the commitment boundary are the available hedge.

5.7 Supervision and the aggregation problem

A supervisor overseeing a market of n institutions cannot maintain n bespoke semantic relationships, and, more fundamentally, cannot aggregate risk across reports whose terms do not share reference. The post-crisis reporting experience settled this empirically: mandated derivative reporting into repositories produced data that could not be matched or summed until identifiers and critical data elements were harmonised globally (CPMI-IOSCO, 2018). It is sometimes suggested that AI dissolves this problem, since a capable model can read n heterogeneous submissions. It cannot dissolve it, for the reason given in Section 3: whether two records denote the same trade, the same entity, the same exposure, is an identity question, and identity in an adversarial reporting environment is fixed by identifiers and attestations, not inferred from resemblance. Inference over heterogeneous regulatory data does not produce supervision; it produces a well-read guess about the system's state, held by the one actor whose statements about the system's state must not be guesses.

5.8 The commitment argument: adaptive agents raise the value of invariants

The deepest version of the preservationist case is almost paradoxical. The classical function of a standard is to constrain behaviour so that others can rely on it. Human institutions could partly substitute reputation, habit, and slow change for explicit constraint, because humans are not very adaptive. Autonomous agents are maximally adaptive: they can renegotiate, re-optimise, and re-interpret continuously. A market of such agents in which everything is renegotiable is a market in which nothing can be relied upon, and reliance is what markets sell. The more flexible the agents, the more valuable the invariants: the properties guaranteed not to be renegotiated, the interpretations guaranteed not to drift, the boundaries within which optimisation is licensed. Machine autonomy does not relax the need for explicit, testable constraints; it converts that need from an efficiency preference into a precondition for the market's existence. It is, on reflection, unsurprising that the very first act of the agentic-AI ecosystem was to write standards for itself: the Model Context Protocol and the Agent2Agent protocol arose within months of agent deployment, produced by AI firms, because even machine-to-machine interaction was discovered to need fixed points (Anthropic, 2024; Google, 2025). The ecosystem that was supposed to make standards obsolete began by demanding them. That is not an irony; it is the thesis.

Finally, the tension between innovation speed and operational resilience, often presented as the decisive argument against standards, is better understood as an argument about *where* standards sit. Speed at the edges is compatible with rigidity at the core; indeed each enables the other, because experimentation is only safe when the commitment boundary it presses against is well-defined. The regulated future is not slow AI. It is fast AI inside certified envelopes, a construction Section 9 makes concrete.

6 The Middle View: Standards Migrate Up the Stack

Sections 4 and 5 are both, in this report's judgement, substantially correct. Both views have merits, but on different levels, and the appearance of contradiction dissolves once the object of the dispute is stated precisely. The question is not whether standards *about* syntax lose value while standards *about* semantics gain it, as though these were different documents with separate fates. Nearly every consequential standard bundles several aspects at once: a syntactic aspect, how content is serialised and carried; a semantic aspect, what the content means; and a pragmatic aspect, what rights, obligations, and processes attach to it. ISO 20022 is the canonical illustration, since its own methodology separates a business model of meanings from the message syntaxes that carry them (ISO, 2013), and a market rulebook likewise couples defined terms to procedures to formats within one instrument. What AI redistributes is value across the aspects *within* each standard, and across the standard population weighted by aspect: the abolitionist case of Section 4 bites hardest on the syntactic aspect away from commitment boundaries, and the preservationist case of Section 5 bites hardest on the semantic and pragmatic aspects and wherever the commitment functions of Section 3 are engaged. Three migrations follow from this redistribution. The first two are, in the report's judgement, well supported by evidence and inference; the third is more contested, and the discussion marks it so.

6.1 First migration: from the syntactic aspect toward the semantic and pragmatic aspects

The value of the syntactic aspect declines. It does not reach zero, for the security and determinism reasons of Section 5.6: at the commitment boundary, strict grammars and canonical serialisations remain the cheapest and most adversarially robust validation technology available, and canonical form is what gets hashed, signed, and archived. But away from that boundary, the historical rationale for syntactic uniformity, that heterogeneous forms were expensive to bridge, lapses. What appreciates is the semantic aspect, and it appreciates for a reason that deserves emphasis because it inverts the abolitionist inference. Machine translation between systems is only *checkable* against a semantic reference. An ontology is not what the translator replaces; it is what makes the translator auditable, the difference between a mapping that is plausible and a mapping that is correct. AI without shared ontology produces fluent, coherent, institutionally unanchored mappings, and fluent, coherent, wrong is the most dangerous failure mode in finance, because it passes review. The more translation machines do, the more valuable the fixed points of meaning against which their output is verified. Semantics is consolidated so that syntax can be liberated.

A clarification matters enough here to state bluntly, because the debate is routinely conducted as if it were false. Rulebooks, definitions booklets, statutes and the technical rules made under them, accounting policies, identifier schemes, and documented market practice are standards in the full sense of Section 3.1: references fixed in advance, under stewardship, for repeated use by parties who did not write them. They are not interoperability protocols, and an analysis that quietly equates standards with interface protocols will systematically misread the AI era, because the first migration shifts weight precisely onto this non-interface population. Intellectual honesty requires the concession that travels with the clarification: most of Section 4's arguments apply to these non-interface standards too, and several apply with greater force. A definitions booklet freezes an abstraction harder than a schema does, since renaming a field is refactoring while redefining default is litigation (Section 4.2). Legal and quasi-legal standards lag worst of all, on amendment cycles measured in years and treaty processes measured in decades (Section 4.3). Capture is most valuable exactly at the semantic layer, because whoever fixes the definition moves the money (Section 4.4). The middle view does not claim that the upper aspects are immune to the critique; it claims that the remedy differs by aspect. For the syntactic aspect there is now an exit: machine translation makes heterogeneity cheap, so uniformity can be relaxed. For the semantic and pragmatic aspects of a multilateral market there is no exit, because a private meaning of default is an oxymoron where obligations are multilateral; the live choice is between governed shared meaning and ungoverned shared meaning, never between shared meaning and none. The remedy for the real defects of semantic and legal standards is therefore not their abandonment but their reformed production, which is the second migration, and a reformed change regime, which is the contested third.

Above semantics, the pragmatic aspect, what messages and acts do to rights and obligations, was never really carried by technical standards at all; it was carried by exactly the rulebooks, definitions, statutes, and market practice just discussed, which the technical standards presupposed. The first migration makes this visible. When the machinery below becomes fluid, the fixed points above stand out, and institutions discover that what they actually depend on is not the message format but the finality rule, not the schema but the definition of default, not the API but the allocation of loss. Expect pragmatic standards, previously implicit in legal boilerplate and operational custom, to be made explicit, precise, and, where possible, executable, precisely because autonomous systems cannot be trusted to absorb them from custom.

6.2 Second migration: from prose to executable form, and from drafting cost to ratification politics

The second migration changes what a standard physically is. The historical standard is a document: human-readable, normative in prose, tested only socially. The emerging standard is a bundle of artifacts: a machine-readable data model, a reference implementation whose behaviour is normative, a conformance test suite that operationally defines compliance, formal properties that certified implementations must be shown to satisfy, and, increasingly, executable regulatory logic. This migration predates modern AI, and the finance-native evidence is decisive: the Common Domain Model expresses lifecycle events as executable functions rather than described procedures (FINOS and ISDA, n.d.); ACTUS reduces contract semantics to deterministic payoff algorithms (Brammertz et al., 2009); ISDA's Digital Regulatory Reporting programme compiles reporting rules into open-source code that firms run, converting a compliance interpretation exercise into a software distribution, and is live in production for major reporting regimes (ISDA, 2024); supervisors themselves piloted machine-executable rulebooks in the Bank of England and FCA digital regulatory reporting work (Bank of England and FCA, 2019); and the rules-as-code movement has generalised the idea to legislation (Mohun and Roberts, 2020). The direction of travel was set by the discovery, repeated across all these programmes, that prose does not survive contact with automation: every firm's lawyers read the same paragraph differently, and the divergence surfaces as unmatchable data years later.

AI's effect on this migration is straightforwardly complementary, and this is the report's most consequential inversion of the framing question. One reason executable and formal standards have been rare is cost: writing a formal specification, a reference implementation, and an exhaustive test suite is expensive, and formal verification was economically viable only for the most critical components, a fact documented even by its industrial champions (Newcombe et al., 2015). Large models collapse exactly these costs. They draft specifications from corpora of practice, translate prose rules into candidate logic, generate conformance tests adversarially, formalise properties, and check implementations against them, at a small fraction of historical effort. AI does not compete with standards at this layer; it is the enabling technology for a generation of standards that were previously unaffordable: richer, more formal, more testable, and maintained at software velocity.

But cost was never the whole story, and an honest statement of this migration says so. The other binder is politics. Prose can hold a coalition together precisely because it lets incompatible readings coexist; constructive ambiguity is not always a drafting failure but sometimes the price of agreement, in statutes as in negotiated settlements. Executable form refuses to provide that service. Encoding forces a choice among readings, the reference implementation is therefore a distributional act, and the diff shows exactly whose interpretation lost. AI does not remove this politics; it concentrates and exposes it, by collapsing the cost of every stage except the one that was political all along. The bottleneck moves from drafting, which machines now do, to ratification, which they cannot, and ratification carries the consensus-building, the side payments, and the legitimacy work that committee processes always quietly performed between the drafting sessions. The practical consequence, developed for standards bodies in Section 10.4, is that consensus institutions become more central to standardisation as their drafting function shrinks. When drafting is free, agreement is the entire cost of a standard, and agreement is what standards institutions are for.

6.3 Third migration: from static documents toward governed velocity, and its limits

The third migration concerns the change regime, and it is the point at which the report's synthesis is most exposed and is therefore stated with the most caution. Static standards impose the costs catalogued in Section 4: frozen abstractions, lag, dual-running migrations. Fully adaptive interpretation imposes the costs catalogued in Section 5: drift, irreproducibility, unaccountable change. The synthesis on offer is *governed velocity*: standards maintained as versioned, machine-readable artifacts under explicit change control, evolving on short cycles, with conformance defined continuously by test suites rather than periodically by documents, and with the authority to change concentrated in an accountable process rather than diffused into model updates. Precedents exist: the living-standard model of the modern web replaced dated documents with continuously maintained specifications disciplined by test suites; semantic versioning made compatibility promises machine-actionable; lockfiles reconciled adaptive resolution with deterministic execution by resolving freely and then pinning. But the precedents should be read for what they are: software communities governing software artifacts among parties with broadly aligned interests.

The objection to generalising them writes itself, and it is partly right: transplanted naively, governed velocity is a technologist's theory of institutions. Standards processes are slow for reasons that better tooling does not touch. Due process, consultation, and consensus are not overhead on legitimacy; for voluntary standards they largely constitute it, being the reason non-authors treat the reference as their own. National mirror committees, translation, and international coordination impose cadence floors on de jure standards. Installed bases convert every breaking change into a negotiation with sunk capital. And where a standard is incorporated into contracts or law, its stability is not a defect but the reliance interest it exists to create; a rulebook that released weekly would destabilise the very thing it is for. The evidence therefore supports a narrower claim than the slogan suggests, and the report states the narrower claim as its position. Velocity is migrating to the *artifact* layer of standards, where it demonstrably works: message schemas on regular maintenance cycles, protocol extension packs, executable reporting rules released and patched as software (ISDA, 2024), conformance suites in continuous integration. Ratification, by contrast, retains institutional cadence, and in the legally binding population it should. The unsolved design problem, and it is unsolved, is the coupling: keeping fast artifacts subordinate to slow authority without either paralysing the artifacts or letting the code become the de facto rule while the ratified text ages into commentary. The lockfile pattern, resolve and negotiate freely, then pin a signed, versioned profile before anything touches a commitment path, is offered here as a suggestion for that coupling, not as a demonstrated institutional equilibrium.

Within that suggestion sits a candidate resolution of the bilateral-versus-market-wide tension of Section 3.5, likewise offered as a design proposal rather than a forecast. FIX's rules-of-engagement experience showed that market-wide standards and bilateral profiles are complements: the standard defines the envelope, the profile selects within it. The AI-era generalisation would be negotiation within certified envelopes: agents negotiate parameters, formats, and workflows freely inside an option space the standard defines and certification has bounded, and the negotiated profile is compiled, hashed, and pinned by both parties before execution, so that what runs is deterministic even though what was agreed was negotiated. Section 9.4 develops the pattern as a candidate standard form. Table 1 summarises the three migrations, with the change-regime row read subject to the caveats just stated.

Aspect or layer	**Direction of change**	**Resulting standard form**
Syntax and transport	Commoditised by machine translation; retained at commitment boundaries for validation, canonicalisation, and signing	Strict grammars and canonical serialisations at settlement, reporting, and audit interfaces; free form elsewhere
Semantics	Appreciates; becomes the reference against which machine translation is verified	Shared ontologies with formal semantics; identifier regimes; machine-readable definitions
Pragmatics and process	Made explicit and executable; previously carried implicitly by prose and custom	Executable lifecycle models; machine-readable rulebooks and legal definitions; determination mechanisms
Change regime	Velocity migrates to the artifact layer; ratification retains institutional cadence	Versioned, signed, test-defined artifacts under accountable change control; negotiation within certified envelopes and pinned profiles, offered as a design pattern
Assurance	From document conformance to continuous, evidence-based certification	Conformance suites in continuous integration; property certification; evidence-grade logs

Table 1: The migration of standards under machine-mediated interoperability.

6.4 What the middle view suggests

The middle view yields expectations that can be checked against events over the coming years. They are offered in the register of Section 2.4 as suggestions rather than inferences: directions the analysis makes plausible, not outcomes it can guarantee, because each depends on institutional choices that have not yet been made. If the view is right, the number of actively maintained syntactic message standards should fall while the number of maintained ontologies, conformance suites, and executable rule sets rises. Standards bodies should publish less prose and more code, and should begin measuring adoption by conformance telemetry rather than by citation. The sharpest standardisation activity of the coming decade should target the layer AI itself creates: model governance, agent identity and mandate, evidence-grade logging, and certification of machine-generated components. Bilateral, AI-mediated integration should expand enormously at the periphery of finance while the specified core, settlement, custody, compliance determinations, and contracts, becomes more formally specified than it is today. And the institutions that thrive should be those that understood early that the scarce resource had shifted: from compatibility, which machines now manufacture, to commitment, which only institutions can confer.

7 Tokenized Finance as the Hard Case

Tokenization is the right hard case for this report's question because it is the domain where the two rival intuitions collide at full strength. On one side, tokens are the most machine-legible financial objects ever created: their state is public or shared, their interfaces are self-describing, their behaviour is code. If AI-mediated interoperability works anywhere, it should work here. On the other side, tokenized instruments are claims, and claims are creatures of law, contract, and institution. The collision exposes, with unusual clarity, where the real difficulty lies.

7.1 The anatomy of a tokenized instrument

A tokenized financial instrument is not a token. It is a stack of five bindings, each of which can fail independently, and each of which requires a different kind of standard to hold.

The first binding is *record to asset*: the relationship between the on-chain entry and the thing of value, which may be native to the ledger or may exist off it. The second is *record to rights*: the legal architecture, terms of issue, prospectus, trust deed, fund documents, statute, that determines what the holder of the record may claim, from whom, and in what priority. The third is *rights to controls*: the mechanisms, in code and in agreement, that constrain who may hold and transfer, including whitelists, transfer restrictions, lock-ups, and permissioning. The fourth is *controls to lifecycle*: the events through which the instrument lives, issuance, coupons and distributions, corporate actions, collateral calls, defaults, redemptions, maturity, which typically span on-chain and off-chain processes. The fifth is *lifecycle to cash*: the money leg, which may be a stablecoin, a tokenized deposit, central bank money on or off the ledger, or a conventional payment, and whose coordination with the asset leg determines whether settlement is atomic or merely hopeful.

Syntactic interoperability addresses none of these bindings. A perfectly translated message about a token whose relationship to its rights is uncertain is a perfectly translated uncertainty. This is the structural reason the section's focal question has the answer it has.

7.2 The focal question: syntax or shared institutional meaning?

Is the bottleneck in tokenized finance the lack of syntactic interoperability, or the lack of shared institutional meaning around rights, controls, obligations, lifecycle events, risk, and accountability? The evidence points one way with unusual consistency: the constraint is meaning, and it is not close.

Three lines of evidence support this. The first is the revealed behaviour of institutions. If syntax were the constraint, the formative investments of the institutional tokenization effort would have been translators and bridges. They were not. They were legal statements and legislation establishing what a digital asset *is* as property (Law Commission, 2023; Uniform Law Commission, 2022); private-law principles for digital assets harmonised internationally (UNIDROIT, 2023); model laws making electronic records capable of functioning as documents of title through the concept of exclusive control (UNCITRAL, 2017; United Kingdom, 2023); contractual definition booklets extending netting and close-out to digital asset derivatives (ISDA, 2023); regulatory taxonomies distinguishing e-money tokens from asset-referenced tokens and fixing the redemption rights of each (European Union, 2023); bond data taxonomies (ICMA, 2023); and an ISO identifier so that market participants and regulators can agree which token they are discussing at all (ISO, 2021). Every one of these artifacts fixes meaning; none of them moves bytes. The market, spending its own money, diagnosed its own bottleneck.

The second line is the failure taxonomy of past interoperability programmes in conventional securities markets, which tokenization inherits. TARGET2-Securities delivered a single technical settlement platform for Europe in 2015; a decade later, the binding work remains the harmonisation of corporate action processing, collateral management, and billing practice through the Eurosystem's SCoRE standards and the AMI-SeCo harmonisation agenda (ECB AMI-SeCo, 2020). The platform was the fast part. Meaning was the slow part. There is no reason to expect the causality to reverse merely because the platform is a distributed ledger.

The third line is native to the token world itself and was introduced in Section 3.6: the ERC experience. ERC-20 gave the ecosystem perfect syntactic interoperability, one interface, every wallet, every venue, and the result was semantic chaos wherever the unspecified mattered: tokens that return no value where integrators expect a boolean, tokens that deduct fees in transit, tokens that rebase balances, each one conformant in form and treacherous in meaning (Vogelsteller and Buterin, 2015). The ecosystem's response, security-token standards binding transfers to identity and eligibility such as ERC-3643 (Lebrun et al., 2021), vault standards subsequently annotated with rounding-direction properties (Santoro et al., 2022), modular architectures whose upgrade discipline is enforced by kernels rather than by convention (Mudge, 2020; Boudi, 2026), is a movement from interface standards toward property and control standards. The purest laboratory of machine-readable finance concluded, from its own losses, that readable syntax without fixed semantics is a liability, which is the section's thesis in one sentence.

None of this implies that syntactic work is worthless; it implies an ordering. The binding constraints on institutional tokenization rank as follows: first, legal semantics, what the token is, what control means, when transfer is final, how the holder's claim survives an intermediary's insolvency; second, control and lifecycle semantics, how restrictions, entitlements, and events are represented and executed consistently; third, process standards, how the surviving off-chain machinery, transfer agents, paying agents, registrars, fund administrators, interlocks with the ledger; and only fourth, syntax, where AI assistance is already abundant. AI's arrival does not relax the constraint; it makes the constraint visible, by clearing away the syntactic underbrush that allowed institutions to believe the problem was plumbing.

7.3 The asset walk

The general claim gains texture instrument by instrument.

Tokenized deposits. A deposit token is a claim on a specific bank, and interoperability between deposit tokens of different banks is not a translation problem but a *settlement* problem: transferring the token across banks requires an interbank obligation to arise and be discharged, which is what correspondent accounts, common rulebooks, and central bank money are for. This is why every serious tokenized-deposit initiative, from regulated liability network experiments to the BIS-convened Project Agorá and the unified-ledger programme of work (BIS, 2023; BIS Innovation Hub, 2024), is at bottom the drafting of a rulebook and the engineering of a common settlement asset, not the construction of adapters. An AI cannot mediate between two deposit tokens whose issuing banks have not agreed how to settle, because there is nothing to mediate; the missing object is an obligation, and obligations are made by agreement.

Stablecoins. The economically salient property of a stablecoin is not its ERC interface but its redemption right: who may redeem, at what value, in what timeframe, backed by what assets under whose custody. Jurisdictions have begun standardising exactly this, by legal category rather than by message format (European Union, 2023), and the international supervisory community has applied the settlement-finality expectations of financial market infrastructure to systemically used arrangements (CPMI-IOSCO, 2022). Two stablecoins with identical syntax and different redemption law are different instruments, and no interpreter can make them the same.

Tokenized fund shares. Money-market and other fund tokenizations pose the register question in its cleanest form: is the on-chain record the register of members, or a mirror of a register kept by a transfer agent? The distinction determines whose entry is authoritative when they diverge, what a court will enforce, and what a fork or an exploit does to ownership. Live institutional products have answered it by keeping a regulated transfer agent in the loop as keeper of the official record, with the chain as the operational surface. The answer is legal-architectural; the token standard merely implements it. An AI asked who owns this share when chain and register diverge is not being asked to interpret; it is being asked to adjudicate, and adjudication is an office.

Tokenized bonds. Digital bond issuances since 2021, including those of supranational issuers under successively adapted French and Luxembourg law, demonstrate both the feasibility and the current shape of the semantic gap: the instruments work, but each issuance is a bespoke assembly of law, documentation, platform, and cash leg, which is precisely the condition, artisanal semantics, that prevents scale. The standardisation now underway, bond data taxonomies, common lifecycle event models, DLT-adapted market infrastructure regimes (ICMA, 2023; European Union, 2022), targets the assembly, not the messages.

Tokenized loans and private credit. Loans add servicing: payment application, covenant monitoring, modifications, participations versus assignments, each with distinct legal consequences that the token must respect. Here AI has enormous operational purchase, credit agreements are long, heterogeneous, and machine-readable only with effort, and Appendix A credits it fully; but the transferability of a loan interest and the rights of participants against the agent are fixed by the credit documentation and by doctrine, not recoverable from token state.

Derivatives and collateral. Tokenized collateral is among the most promising near-term applications, because collateral movement is a domain of pure operational friction wrapped around crisp legal semantics: eligibility criteria, haircuts, transfer-of-title versus security-interest constructions, close-out netting. The crisp semantics is why it can be tokenized safely, and it is crisp because decades of documentation standards made it so (ISDA, 2023). The lesson runs opposite to the abolitionist reading: tokenization succeeds first where standardisation already did the semantic work.

Tokenized real-world assets. At the far end sit assets, real estate, commodities, receivables, whose value and enforceability live entirely off-chain. Here the fifth binding fails hardest: the token moves at the speed of consensus and the deed moves at the speed of a registry, and the gap between them is not an interoperability gap but an *enforcement* gap. The token can only be worth its off-chain linkage, which is an attestation problem, who certifies that the warehouse holds the metal, that the title is unencumbered, that the SPV owns the receivable, and attestation at scale is a standards problem: standardised claims, standardised attestors, standardised liability when attestations fail. Proof, not hope, is the design criterion, and proof requires a specification of what is being proved.

7.4 Controls: token standards, transfer agents, wallets, identity, permissions

The control plane of tokenized finance is where code, contract, and regulation meet, and it is instructive that every element of it is currently being standardised, by different communities, at once. Transfer restrictions are moving from legal boilerplate to executable eligibility logic bound to identity credentials (Lebrun et al., 2021); the identity credentials themselves are being standardised as verifiable credentials and decentralised identifiers (W3C, 2022a,b), with government-anchored identity wallets arriving through the European digital identity framework (European Union, 2024b); originator and beneficiary information for asset transfers has a dedicated inter-VASP messaging standard because the travel rule demanded one (interVASP JWG, 2020); and custody is being reshaped by the accounting and prudential treatment of safeguarded assets, where the rescission of restrictive accounting guidance in early 2025 materially changed the institutional calculus (SEC, 2025). What AI contributes to this plane is monitoring, screening, and anomaly detection at scale. What it cannot contribute is the thing every element above supplies: a determination, valid against third parties, that this holder is eligible, this transfer is restricted, this key controls this asset. Eligibility is a legal status; control is a legal test (Uniform Law Commission, 2022); neither is an inference.

7.5 Cross-chain, cross-platform, cross-jurisdiction

Interoperability across ledgers is the domain where the cost of substituting cleverness for standards has been paid most publicly. Cross-chain bridges, which are precisely dynamic adapters between heterogeneous systems, produced several of the largest thefts in the history of finance in 2022, with individual incidents in the hundreds of millions of dollars, because they coupled ledgers with different security models and different finality semantics without a shared definition of either. The institutional alternatives under construction are instructive in form: synchronised settlement across ledgers with an RTGS system acting as the coordinating leg (Bank of England and BIS, 2023); wholesale-CBDC-based cross-currency settlement with standardised venue logic (BIS Innovation Hub et al., 2023); and, in the euro area, a dual-track programme to settle DLT transactions in central bank money, a near-term bridge, Pontes, linking DLT platforms to existing TARGET services, and a longer-term exploration, Appia, of a native integrated architecture (ECB, 2025). Each of these is an exercise in standardising finality across heterogeneous systems: defining the conditional states, the moment of unconditionality, and the failure procedures, so that atomicity is a property guaranteed by design rather than approximated by inference. Cross-jurisdiction interoperability raises the same structure one level up, in conflict of laws: which law governs a token, where a digital asset is situated, whose insolvency rules apply to an intermediated holding. The international instruments addressing this (UNIDROIT, 2023; UNCITRAL, 2017) are, again, semantic standards in legal form. An agent can route around a missing bridge. It cannot route around a missing rule of priority, because priority is not in the state; it is in the law that ranks claims to the state.

7.6 The DAO lesson, and the ranking restated

One episode compresses the whole section. In 2016, the largest smart contract of its era was drained through behaviour that its code permitted, and the ecosystem most ideologically committed to the proposition that code is law responded by forking the ledger to reverse the outcome, a governance act that overrode execution in the name of intent. Whatever one thinks of the decision, its lesson is structural: when execution and meaning diverge at sufficient stakes, meaning wins, and the mechanism by which it wins is governance, not interpretation. The pragmatic layer is sovereign over the syntactic layer, in tokenized finance as everywhere else; tokenization merely shortens the distance between them until the hierarchy is impossible to miss.

The section's conclusion, then, is the ranking already stated, now earned: legal semantics binds tightest, then control and lifecycle semantics, then process, then syntax. AI enters this stack from the bottom, and its effect is to strip away the layers at which the industry has historically located its interoperability problem, revealing that the problem was always higher up. In tokenized finance, machine-mediated interoperability is not the alternative to shared institutional meaning. It is the searchlight that shows how little of it yet exists.

8 What AI Can Solve, and What It Cannot Settle

This section applies the report's central discipline to sixteen operational domains, from data mapping to model governance. For each domain, five questions are asked that are habitually conflated: what AI can plausibly *automate*; what still requires *standards*, meaning fixed reference points; what requires *governance*, meaning an accountable process for decisions and exceptions; what requires *legal authority*, meaning a determination that binds; and what is best left to *market convention*, meaning defaults that fill the space specification leaves open. The verb in the section title is chosen deliberately: AI solves problems; institutions settle questions. Because the sixteen analyses repay reading but share a stable structure, the full domain-by-domain treatment is carried in Appendix A. This section exhibits the method on three domains chosen to span the spectrum and then distils the pattern that holds across all sixteen. Throughout, claims about automation potential are inference from the observed capability trajectory, and claims about where authority sits are statements of current institutional fact, in the registers of Section 2.4.

8.1 Three exhibits

Data mapping anchors the automation end. It is the strongest case for machine substitution in the report, and the concession is full: expect near-total automation of mapping *proposal*. What survives is small and load-bearing. A canonical target model is still needed so that n systems require n mappings rather than n^2; golden datasets are still needed to define what a correct mapping even is, since without an agreed answer key accuracy is unmeasurable; and approval and change control under model-risk discipline are still needed wherever mappings feed regulated processes (Federal Reserve and OCC, 2011). Legal authority is rarely engaged, which is precisely why the domain automates so well.

Default determination and corporate actions anchor the authority end. AI can monitor every covenant, parse every issuer announcement, and flag the approach of the cliff with superhuman vigilance; it cannot pronounce the fall, because a default declaration is a legal act with formalities, and markets have deliberately institutionalised the systemic cases in determinations bodies whose single governed answer binds thousands of contracts identically. A machine's parse of a corporate action announcement, however accurate, is likewise not the entitlement; the issuer's corporate act and the registrar's record are. Detection is a capability; declaration is an office.

Model governance is the reflexive exhibit, and the report's argument in miniature. Governing the AI that performs everything above is a domain that barely existed before machine interpretation, and it now consists almost entirely of new demand for standard forms: model-risk doctrine extended from quantitative models to learned ones (Federal Reserve and OCC, 2011), statutory obligations of logging, transparency, and human oversight (European Union, 2024a), management-system and risk-management frameworks (ISO/IEC, 2023; NIST, 2023), and, ahead, the certification and provenance regimes of Section 9. The technology that was supposed to dissolve standards became their fastest-growing subject.

8.2 The pattern

Across the sixteen domains analysed in Appendix A, the division of labour is stable enough to state as a rule, and Table 2 compresses it. *AI dominates perception and drafting*: reading, extracting, matching, monitoring, reconstructing, proposing. *Standards dominate reference*: what counts as what, against which key accuracy is even defined, in which form commitments are expressed at the boundary. *Governance dominates change and exception*: who approves, who decides the hard case, how the rules themselves evolve. *Legal authority dominates bindingness*: whose statement is final, who bears the loss, what survives insolvency. *Convention dominates the residue*: tolerances, defaults, and practice, now discoverable by machines at unprecedented speed and therefore likely to crystallise into explicit standards faster than before. In imperative form, the rule of the AI era in finance reads: automate the inference, standardise the reference, govern the exception, and legislate the finality.

9 Future Forms of Standards

If the analysis so far is right, the interesting question is not whether standards survive but what they will physically and institutionally *be*. This section projects the forms, organised into six families, and marks

Domain	AI increasingly automates	Binding constraint
Data mapping	Mapping proposal, transformation code, anomaly flags	Canonical models; golden datasets; approval control
Process mapping	Process discovery from logs; executable process drafts	Shared event taxonomy; coordinated change
Legal interpretation	Retrieval, comparison, ambiguity detection, drafting	Binding interpretation by an authorised body
Contractual reconciliation	Matching, break classification, resolution drafting	Common trade representation; dispute protocol
Risk classification	Feature extraction; classification proposals	Category boundaries fixed in regulation
Regulatory reporting	Report assembly, anomaly detection, mapping upkeep	Identifiers; executable rules; statutory rulebook
AML, KYC, sanctions, tax	Entity resolution, alert triage, narrative drafting	Lists and thresholds as legal acts; compelled data
Settlement disputes	Evidence assembly; timeline reconstruction	Finality regime; infrastructure rulebook
Corporate actions	Announcement parsing; entitlement computation	Issuer's act and registrar's record; loss allocation
Transfer restrictions	Continuous eligibility monitoring	Eligibility as legal status, evidenced by attestation
Beneficial ownership	Graph analysis across ownership structures	Compelled declaration under statutory definitions
Off-chain collateral	Schedule extraction; monitoring; call drafting	Executable eligibility criteria; insolvency law
Default events	Covenant surveillance; early warning	Declaration formalities; determination governance
Waterfalls and NAV	Waterfall encoding; scenario testing	Governing agreement; NAV as an authoritative act
Loan servicing	Tape reading; payment application; reporting	Servicing data model; borrower-protection law
Model governance	Evaluation generation; drift monitoring	Deployment governance; deployer liability

Table 2: Sixteen operational domains: what automates and what binds.

each as evidence, inference, or speculation. The families are cumulative: each presupposes the ones before it, and together they describe a standards stack for machine-mediated finance.

9.1 The semantic family

Shared ontologies with formal semantics. The successor to the data dictionary is an ontology whose concepts carry formal definitions, constraints, and relationships, machine-consumable and human-governed. The lineage exists as evidence: the financial industry business ontology, the CDM's product and event model, ACTUS's algorithmic contract types (EDM Council, n.d.; FINOS and ISDA, n.d.; Brammertz et al., 2009). The inference is that these become the verification substrate for machine translation, the answer key of Section 6.1, and that their governance, who may change what a coupon *is*, becomes correspondingly consequential. The speculative extension is domain ontologies for the properties of financial smart contracts themselves, of which the Blockchain Property Ontology programme is an early instance (IEI, 2026), giving names and formal content to properties like conservation of value, authorisation monotonicity, and reentrancy safety so that they can be required by reference.

Semantic standards for legal concepts. Model-readable legal and regulatory rules extend the family into law: statutory and contractual definitions published in a canonical, machine-consumable representation alongside the authoritative prose, with the isomorphism between them maintained as a first-class obligation. The Ricardian lineage, one document, dual readability, legal text and code bound by hash, is the design pattern (Grigg, 2004; Clack et al., 2016); digital regulatory reporting is the production evidence (ISDA, 2024); rules-as-code is the public-sector generalisation (Mohun and Roberts, 2020). AI's role is translation and consistency-checking between the readings; authority's role is unchanged, since only the enacting institution can declare which reading governs when they diverge.

9.2 The executable family

Reference implementations as normative artifacts. The internet's lesson, running code beats paper, becomes doctrine: the standard ships as an implementation whose behaviour is the specification of record, with the prose demoted to commentary. Evidence: this is already how the CDM and DRR operate, and how de facto standards in open source have always worked. The inference is that standards bodies become maintainers of software, with everything that implies about release engineering, security response, and long-term stewardship.

Conformance test suites as the definition of compliance. The strongest single prediction of this report: the centre of gravity of standardisation moves from the document to the test. A standard becomes the set of behaviours its suite accepts; certification becomes a passing run under controlled conditions; adoption becomes measurable as continuous conformance telemetry rather than declared intent. AI is the enabling technology, generating adversarial test cases at a density committees never achieved, and the suite is simultaneously the discipline for AI-generated components: a machine-written adapter is acceptable exactly insofar as it passes the certified suite, which resolves the adapters-versus-certified-interfaces tension of Section 5.5 by composition rather than by choice. The standard of the future is a test you can fail; a standard you cannot execute against will not survive contact with AI-speed integration.

Formal specifications and control-property standards. Above tests sit proofs. For the commitment core, settlement engines, custody controls, token contracts holding systemic value, the appropriate standard form is a set of formally stated properties with machine-checkable proofs of conformance, the level-of-rigor logic that safety-critical software certification has applied for decades, graded by consequence. Historically this was unaffordable outside aerospace and a few infrastructure providers (Newcombe et al., 2015); AI-assisted specification and proof changes the economics, which is the concrete meaning of the arc this report has elsewhere summarised as the movement from programmable finance to provable finance. The token-native evidence that interface-only standards are insufficient, the ERC-20 and ERC-4626 pathologies of Section 7.2, is the demand signal; smart-contract package standards, versioned, signed, property-certified contract components with declared storage layouts and upgrade discipline (Mudge, 2020; Boudi, 2026), are the supply form, the lockfile-and-semver pattern of Section 3.6 transplanted to code that holds money.

9.3 The runtime family

Runtime compliance policies. Compliance moves from periodic attestation to policy-as-code evaluated at transaction time: eligibility, restriction, limit, and sanction logic executed in the transfer path, with the policy itself versioned, signed, and standard in form so that supervisors can read one policy language across many venues. Evidence exists in embryo in permissioned token standards (Lebrun et al., 2021); the inference is a common policy expression layer; the speculation is supervisory nodes that verify policy presence and version in real time.

Machine-verifiable audit logs. The evidentiary counterpart: hash-chained, signed decision records capturing inputs, model identity and version, parameters, outputs, and human overrides, sufficient to replay or at least to reconstruct any consequential automated decision. This is the audit-trail remedy of Section 5.4 made into an artifact, and its standardisation is what turns a vendor log into evidence: common schema, common integrity mechanism, common retention semantics. Statutory logging duties for high-risk systems (European Union, 2024a) are the demand; the report's inference is that evidence-grade logging becomes a procurement default in regulated finance within the decade.

Settlement and finality standards for heterogeneous ledgers. The runtime family's keystone: explicit, per-system definitions of conditional and unconditional states, the designated moment of finality, and failure procedures, expressed so that cross-system atomicity can be engineered rather than assumed, and mapped onto the existing legal machinery of designation and protection (European Union, 1998; CPMI-IOSCO, 2012, 2022). The synchronisation experiments and the euro-area dual-track programme (Bank of England and BIS, 2023; BIS Innovation Hub et al., 2023; ECB, 2025) are, in substance, drafts of this standard. Its content is the answer to the bridge failures of Section 7.5: interoperability of commitments, not of messages.

9.4 The agentic family

Agent identity, mandate, and authority standards. Before agents can transact, the market must be able to answer: whose agent, with what authority, within what limits, and who is bound by its acts. The legal frame is old, agency, and the technical instantiation is new: verifiable credentials expressing mandates, spending and scope limits as machine-checkable fields, revocation as a first-class operation, an identifier regime for agents with the same seriousness the LEI brought to entities (W3C, 2022a,b; European Union, 2024b; FSB, 2012). This is inference bordering on the obvious; the speculation is only in the timeline.

Interface-discovery protocols with attestable descriptions. Agents will discover capabilities dynamically; the early protocols already exist (Anthropic, 2024; Google, 2025). The financial-grade requirement is that discovery metadata be signed and attestable, because an interface description an agent reads is an input an adversary can forge, and the prompt-injection literature transfers directly: unauthenticated capability descriptions are the phishing surface of the agentic era. Discovery without attestation is how the next generation of bridge losses happens off-chain.

Agent negotiation protocols: negotiate, freeze, execute. The report's resolution of negotiated compatibility versus enforceable finality, stated architecturally. Agents negotiate within a certified envelope, the option space a standard defines and certification has bounded; the negotiated outcome compiles to a deterministic profile; both principals' systems sign and pin the profile's hash; execution runs only the pinned artifact; the transcript and profile enter the evidence-grade log. Negotiation upstream, finality downstream, a cryptographic commitment bridging them. Nothing in this pattern is exotic, its components are FIX rules-of-engagement, lockfiles, and code signing, recombined, and that ordinariness is the point: the institutional technology for taming adaptive behaviour already exists in fragments and awaits assembly into a named standard.

9.5 The assurance family

Certification regimes for models and machine-generated components. Type approval comes to interpretation: models certified for defined tasks under defined conditions, adapters certified against suites, with certificates that are themselves machine-verifiable and revocable. The graded-assurance pattern of existing safety-critical regimes supplies the template; management-system and risk-framework standards (ISO/IEC, 2023; NIST, 2023) supply the organisational layer. Continuous certification, conformance as a property monitored in production rather than examined annually, is the inference from governed velocity. Provenance standards, a bill of materials for models and adapters answering what generated this, from what, when, mirror the software supply-chain response to compromise and are, speculatively, the artifact insurers will demand first.

9.6 The change-management family

Finally, standards for how standards change: versioning contracts with machine-actionable compatibility semantics, deprecation and sunset procedures, migration test kits, and the governance metadata, who may propose, who ratifies, on what cadence, published as part of the standard itself. This family is unglamorous and decisive, because governed velocity (Section 6.3) is only as credible as its change machinery, and because the alternative on offer, semantic change arriving silently through model updates, is the precise failure mode the whole stack exists to prevent. Across all six families, one design constant should be visible: each form takes something AI does cheaply, drafting, testing, monitoring, translating, and couples it to something only institutions do, ratifying, certifying, binding, bearing liability. That coupling, not any particular artifact, is the future form of standards.

10 Strategic Implications for Builders, Regulators, and Institutions

The analysis converts into posture. This section states the implications bluntly for each of the four audiences named in the report's remit, as recommendations grounded in the preceding sections rather than as exhortation.

10.1 For founders and infrastructure builders

Do not build translation moats. Translation is the layer AI commoditises by definition, and any business whose defensibility is I connect A to B is selling an asset with a falling replacement cost. The durable positions are one level up and one level down: up, in the semantic layer, owning or stewarding the reference model against which translations are verified; down, in the commitment layer, operating the deterministic, certified core that probabilistic systems surround. If the venture is a protocol, ship the conformance suite with the protocol and treat the suite as the product's constitution; the ERC record shows that an interface without properties composes losses (Vogelsteller and Buterin, 2015; Santoro et al., 2022), and the certified-suite pattern is how machine-generated integrations become procurable (Section 5.5). Design for evidence from the first commit: signed states, replayable transitions, evidence-grade logs are not compliance overhead but the feature institutions will pay for, because they are what makes the system a legal object. For tokenization specifically, bind the token to its legal architecture explicitly, name the authoritative register, encode the controls as attestation-consuming logic, and treat the cash leg's finality as a design input, not an integration detail. And adopt the emerging semantic standards early even where they are imperfect: in a machine-mediated market, conformance is distribution, because agents route to what they can verify.

10.2 For regulators and central banks

Publish rules as both prose and code, and own the isomorphism. The digital-regulatory-reporting evidence (ISDA, 2024; Bank of England and FCA, 2019) shows that executable rules eliminate interpretive divergence at the source; the rules-as-code practice (Mohun and Roberts, 2020) shows the method generalises. Extend the identifier doctrine that fixed derivatives reporting (CPMI-IOSCO, 2018; FSB, 2012) to the new objects: tokens already have one (ISO, 2021); agents and mandates need theirs. For AI in regulated processes, three requirements follow directly from Sections 5 and 8: deterministic commitment paths (probabilistic components advise, certified components commit); evidence-grade logging sufficient for reconstruction (European Union, 2024a; Federal Reserve and OCC, 2011); and model-concentration supervision, because correlated interpretation across a market monoculture is a stability exposure that no existing return captures, and stress-testing it is feasible now, before it matters. Use sandboxes and pilot regimes (European Union, 2022) for their highest purpose, which is not to relax standards but to *discover* them: let adaptive, AI-mediated practice run under observation, and codify what converges. Finally, treat finality guidance for heterogeneous ledger arrangements (CPMI-IOSCO, 2022; ECB, 2025) as urgent rather than eventual; every month of ambiguity is being priced somewhere, usually as a haircut on innovation.

10.3 For financial institutions

Rewrite procurement before rewriting architecture. The operative questions for any AI-mediated or machine-generated component are the ones Section 5.5 identified: what specification does it warrant, what suite certifies it, what happens on drift, who bears interpretive loss, and can its decisions be reconstructed in five years. Demand property conformance, not just interface conformance. Treat every consequential mapping, classification, and interpretation as a model under governance (Federal Reserve and OCC, 2011), because the supervisor will, and because the discipline is protective rather than performative: the institution that can replay its decisions owns its disputes. Build the internal ontology as a strategic asset; institutions that maintain one canonical semantic layer will deploy AI integration an order of magnitude faster than those that let every project negotiate meaning afresh, and the same layer is the map that makes exposure aggregable on demand, the capability regulators have demanded since 2013 (BCBS, 2013). Allocate interpretation risk contractually and explicitly in every AI-mediated service relationship, because silence allocates it too, just unpredictably. And participate in the semantic consortia early, not for citizenship but for interest: the definitions being fixed now are where the next decade's rents and constraints are being written.

10.4 For standards bodies

The product changes or the institution fades. The deliverable migrates from documents to maintained artifacts: data models, reference implementations, conformance suites, machine-readable rule text, with the prose as commentary (Section 9.2). Adopt AI internally without apology, for drafting, formalisation, test generation, and consistency checking, and measure the gain where it counts, cycle time from proposal to certified release; the residual, consensus and authority, is the comparative advantage, and it always was (Section 6.2). Restructure around governed velocity: living standards under explicit change control, releases in weeks, conformance measured continuously; a body whose cadence is quinquennial is not participating in this market. Co-draft with legislators and regulators where the standard's concepts must carry legal force, because semantic standards without legal anchoring reproduce the isomorphism gap in the other direction. Measure success as executed conformance, suites passing in production, not as document sales or citations. And accept, without institutional anxiety, the role the analysis assigns: when machines write the drafts, the certifier of meaning is not diminished; it is all that is left, and it is the part that was always sovereign.

11 Final Synthesis

11.1 The argument in one movement

The question was whether standards remain necessary for software, and especially for tokenization and finance, in an era of exceptionally advanced AI. The answer this report has defended is that the question conflates two functions that history fused and AI separates. Standards compress the cost of compatibility, and standards fix the locus of commitment. Machine-mediated interoperability substantially dissolves the first function, and the dissolution should be welcomed: the long tail of integration, the reading of legacy formats, the bespoke glue that consumed the industry's engineering budgets, all of it passes to machines, and the standards that existed only to make that work tractable will quietly retire. The second function does not dissolve, because it was never an information-processing problem. Commitment, finality, liability, eligibility, and authority are relations among accountable persons, constituted by governance and law, and no improvement in interpretation supplies them, for the same reason that no improvement in cartography supplies territory.

From that separation, the four scenarios of the report's remit resolve into one differentiated forecast rather than four rivals. Standards become less important exactly where their function was compression: syntactic formats, bilateral integration conventions, the documentation of the peripheral. Standards become more important exactly where automation raises speed, coupling, and adversarial pressure on the commitment core: finality, canonical validation at boundaries, identifiers, evidence, model governance. Standards change role pervasively, migrating from prose to executable and testable artifacts whose production AI subsidises and whose ratification institutions retain. And the hybrid scenario is not a compromise among these but their composition: a periphery of fluid, machine-negotiated compatibility around a core of deepening, machine-verifiable commitment, with the boundary between them, the point where meaning is pinned, becoming the most carefully engineered location in the architecture.

Tokenized finance, the report's central case, is where this composition will be built first and tested hardest, because tokenization removes the buffers of time and manual mediation that let conventional markets survive semantic vagueness. Its bottleneck, the report has argued from institutional behaviour, from programme history, and from the token ecosystem's own losses, is shared institutional meaning, not syntactic interoperability, and the arrival of machines that eliminate the syntactic excuse will make the semantic deficit impossible to ignore. That is, on balance, the most optimistic reading of the era: AI does not solve the hard problem of tokenized finance, but it strips away everything that allowed the industry to avoid it.

11.2 The allocation of labour

Table 3 states the synthesis operationally, layer by layer.

11.3 Ten conclusions, provocative but defensible

1. AI commoditises translation and thereby raises the premium on what cannot be translated. The scarce resource in interoperability shifts from compatibility, which machines manufacture, to commitment, which only institutions confer.
2. The first act of the agentic-AI ecosystem was to write standards for itself. Before autonomous agents were widely deployed, their developers published fixed protocols so that adaptive systems would have firm interfaces to stand on: a common protocol for connecting models to tools and data in 2024, and a common protocol for exchange between agents in 2025 (Anthropic, 2024; Google, 2025). The lesson generalises: the more behaviour adapts, the more valuable the points that do not. Machine autonomy manufactures demand for fixed points; it does not dissolve it.
3. In finance, the bottleneck was never syntax. A market can survive bad message formats indefinitely; it cannot survive contested finality for an afternoon.
4. Two highly capable models with adverse mandates will disagree more precisely, not less. Capability arms interpretation; only authority converges it. Ambiguity detection is a capability; ambiguity resolution is an office.
5. Probabilistic systems can advise on state transitions but cannot constitute them. Determinism at the settlement boundary is not a limitation to be engineered away; it is the product.

Layer	AI can help with	Standards still needed for	Governance or law needed for
Wire syntax and transport	Parsing, translation, adapter generation, legacy format recovery	Canonical serialisation, strict validation and signing at commitment boundaries	Change control at regulated interfaces
Data semantics	Mapping proposal, anomaly detection, ontology drafting and population	Shared ontologies, identifiers, golden datasets that define correctness	Ratifying semantic change; stewardship of reference models
Contractual meaning	Extraction, comparison, ambiguity detection, drafting of computable clauses	Definition booklets, computable-contract forms, legal-code isomorphism	Binding interpretation: courts, determinations bodies, rulebooks
Process and lifecycle	Process discovery, event parsing, entitlement computation, exception triage	Event taxonomies, executable lifecycle models, market practice	Deadlines, loss allocation for failures, change coordination
Compliance and reporting	Assembly, screening triage, entity resolution, anomaly detection	Identifiers, schemas, executable rules, attestation formats	Rulebooks, sanctions acts, thresholds, cross-border recognition
Settlement and finality	Pre-settlement checking, evidence assembly, dispute reconstruction	Finality definitions, DvP and atomicity models, fail procedures	Designation, insolvency protection, default management, unwinds
Identity and mandates	Verification workflows, fraud and anomaly detection	Credential formats, agent identifiers, mandate and limit fields	Attestor accreditation, revocation authority, agency liability
Assurance and audit	Test generation, conformance monitoring, drift detection, log analysis	Conformance suites, property specifications, evidence-grade log formats	Certification regimes, auditor and supervisor authority, sanctions for breach

Table 3: Allocation of labour across the interoperability stack in the AI era.

6. When drafting becomes free, agreement becomes the entire cost of a standard, and the comparative advantage of standards institutions shifts from producing documents to conferring bindingness.
7. Interface standards without property standards compose bugs as efficiently as they compose features. The token ecosystem paid to establish this, and the receipt is public.
8. The successor to the standards document is the conformance suite: an executable battery of tests that defines compliance by acceptance rather than by reading. When components are machine-generated in minutes, a standard that can only be read cannot keep pace with what is built against it, while a standard that can be run disciplines every generation instantly. A machine-generated component is procurable exactly insofar as a certified suite accepts it, adoption becomes measurable as passing runs in production rather than declared intent, and a standard you cannot execute against will not survive contact with AI-speed integration.
9. Model monoculture is the new systemic concentration. Correlated interpretation error across a market is a financial-stability exposure that no current supervisory return measures, and it is accumulating now.
10. Tokenization will scale when control is standardised in law and lifecycle semantics are standardised in code. AI accelerates every step of that programme except the one that matters most: the institutional act of making a definition binding.

11.4 Final thesis

In an era of exceptionally advanced AI, standards survive because their deepest function was never to help machines or people understand each other but to fix what shall count as understanding when interests diverge and value is at stake. Artificial intelligence collapses the cost of translation, adaptation, and integration, and in doing so retires the syntactic standards that existed only as workarounds for expensive interpretation; but the same force raises the speed, coupling, and adversarial sophistication of markets, and therefore raises the value of the reference points, semantic, executable, testable, and legally anchored, against which machine behaviour can be verified, certified, bounded, and adjudicated. Standards accordingly migrate rather than disappear: up the stack from syntax to semantics to obligation, from prose documents to conformance suites and formal properties, from the periphery of integration to the commitment core of settlement, custody, compliance, and identity, and outward onto AI itself, which becomes the newest and fastest-growing object of standardisation. Tokenized finance, whose binding constraint is shared institutional meaning rather than message compatibility, is where this migration will be completed first, because it is the domain in which interpretation must terminate, deterministically and accountably, in an answer that binds. Translation has become cheap. Commitment remains scarce. The institutions that understand which of the two they are in the business of supplying will define the next era of financial infrastructure.

References

Abbate, J. (1999). *Inventing the Internet*. MIT Press.

ACTUS Financial Research Foundation (n.d.). *Algorithmic Contract Types Unified Standards: Technical Specification*. ACTUS FRF.

Anthropic (2024). *Introducing the Model Context Protocol*. Anthropic.

Basel Committee on Banking Supervision (2013). *Principles for Effective Risk Data Aggregation and Risk Reporting*. Bank for International Settlements.

Bank for International Settlements (2023). Blueprint for the future monetary system: improving the old, enabling the new. *BIS Annual Economic Report 2023*, Chapter III.

Bank for International Settlements Innovation Hub (2024). *Project Agorá: Central Banks and Banking Sector Explore Tokenisation of Cross-Border Payments*. BIS.

Bank of England and BIS Innovation Hub London Centre (2023). *Project Meridian: Innovating Transactions with Synchronisation*. Bank for International Settlements.

Bank of England and Financial Conduct Authority (2019). *Digital Regulatory Reporting: Pilot Phase 2 Viability Assessment*. London.

BIS Innovation Hub, Banque de France, Monetary Authority of Singapore, and Swiss National Bank (2023). *Project Mariana: Cross-Border Exchange of Wholesale CBDCs Using Automated Market-Makers. Final Report*. Bank for International Settlements.

Boudi, Z. (2026). *ERC-8155: Kernel-Orchestrated Modular Smart Contracts*. Intelligence Economy Institute Research Note, Paris.

Brammertz, W., Akkizidis, I., Breymann, W., Entin, R., and Rustmann, M. (2009). *Unified Financial Analysis: The Missing Links of Finance*. Wiley.

Bratus, S., Locasto, M. E., Patterson, M. L., Sassaman, L., and Shubina, A. (2011). Exploit programming: from buffer overflows to weird machines and theory of computation. *;login:*, USENIX Association.

Commodity Futures Trading Commission and Securities and Exchange Commission (2010). *Findings Regarding the Market Events of May 6, 2010*. Report of the Staffs of the CFTC and SEC.

Clack, C. D., Bakshi, V. A., and Braine, L. (2016). Smart contract templates: foundations, design landscape and research directions. *arXiv preprint* arXiv:1608.00771.

Committee on Payment and Settlement Systems and Technical Committee of IOSCO (2012). *Principles* *for Financial Market Infrastructures*. Bank for International Settlements.

Committee on Payments and Market Infrastructures and IOSCO (2018). *Harmonisation of Critical Data Elements (Other Than UTI and UPI): Technical Guidance*. Bank for International Settlements.

Committee on Payments and Market Infrastructures and IOSCO (2022). *Application of the Principles for Financial Market Infrastructures to Stablecoin Arrangements*. Bank for International Settlements.

David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75(2).

David, P. A. and Greenstein, S. (1990). The economics of compatibility standards: an introduction to recent research. *Economics of Innovation and New Technology*, 1.

EDM Council (n.d.). *Financial Industry Business Ontology (FIBO)*. EDM Council.

European Central Bank (2025). *Eurosystem Initiative to Settle DLT-Based Transactions in Central Bank Money: The Pontes and Appia Tracks*. ECB, Frankfurt.

European Central Bank, Advisory Group on Market Infrastructures for Securities and Collateral (2020, as updated). *Single Collateral Management Rulebook for Europe (SCoRE) Standards*. ECB, Frankfurt.

European Union (1998). Directive 98/26/EC on settlement finality in payment and securities settlement systems. *Official Journal of the European Communities*.

European Union (2022). Regulation (EU) 2022/858 on a pilot regime for market infrastructures based on distributed ledger technology. *Official Journal of the European Union*.

European Union (2023). Regulation (EU) 2023/1114 on markets in crypto-assets (MiCA). *Official Journal of the European Union*.

European Union (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*.

European Union (2024). Regulation (EU) 2024/1183 establishing the European Digital Identity Framework. *Official Journal of the European Union*.

Farrell, J. and Saloner, G. (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics*, 16(1).

Farrell, J. and Saloner, G. (1986). Installed base and compatibility: innovation, product preannouncements, and predation. *American Economic Review*, 76(5).

Financial Action Task Force (2012, as amended). *International Standards on Combating Money Laundering and the Financing of Terrorism and Proliferation: The FATF Recommendations*. FATF, Paris.

Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency (2011). *Supervisory Guidance on Model Risk Management*. SR Letter 11-7.

Fintech Open Source Foundation and International Swaps and Derivatives Association (n.d.). *Common Domain Model (CDM)*. FINOS.

Financial Stability Board (2012). *A Global Legal Entity Identifier for Financial Markets*. FSB, Basel.

Galati, G. (2002). Settlement risk in foreign exchange markets and CLS Bank. *BIS Quarterly Review*, December 2002.

Google (2025). *Announcing the Agent2Agent (A2A) Protocol*. Google Developers Blog.

Grigg, I. (2004). The Ricardian contract. *Proceedings of the First IEEE International Workshop on Electronic Contracting*. IEEE.

Hart, H. L. A. (1961). *The Concept of Law*. Oxford University Press.

International Capital Market Association (2023). *Bond Data Taxonomy*. ICMA, Zurich.

Intelligence Economy Institute (2026). *Blockchain Property Ontology: A Formal Property Framework for Financial Smart Contracts*. IEI Working Programme, Paris.

International Swaps and Derivatives Association (2019). *Legal Guidelines for Smart Derivatives Contracts: Introduction*. ISDA.

International Swaps and Derivatives Association (2023). *ISDA Digital Asset Derivatives Definitions*. ISDA.

International Swaps and Derivatives Association (2024). *Digital Regulatory Reporting (DRR): Programme Documentation*. ISDA.

International Organization for Standardization (2013). *ISO 20022: Universal Financial Industry Message Scheme*. ISO, Geneva.

International Organization for Standardization (2021). *ISO 24165-1: Digital Token Identifier (DTI)*. ISO, Geneva.

International Organization for Standardization and International Electrotechnical Commission (2023). *ISO/IEC 42001: Information Technology, Artificial Intelligence, Management System*. ISO, Geneva.

Joint Working Group on interVASP Messaging Standards (2020). *IVMS 101: interVASP Messaging Standard*.

Katz, M. L. and Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, 75(3).

Law Commission of England and Wales (2023). *Digital Assets: Final Report*. London.

Lebrun, J., Malghem, T., Thizy, K., Falempin, L., and Boudjemaa, A. (2021). *ERC-3643: T-REX, Token for Regulated EXchanges*. Ethereum Improvement Proposals.

Mohun, J. and Roberts, A. (2020). *Cracking the Code: Rulemaking for Humans and Machines*. OECD Working Papers on Public Governance No. 42. OECD Publishing, Paris.

Mudge, N. (2020). *EIP-2535: Diamonds, Multi-Facet Proxy*. Ethereum Improvement Proposals.

Newcombe, C., Rath, T., Zhang, F., Munteanu, B., Brooker, M., and Deardeuff, M. (2015). How Amazon Web Services uses formal methods. *Communications of the ACM*, 58(4).

National Institute of Standards and Technology (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST, Gaithersburg.

Postel, J. (1981). *Transmission Control Protocol*. RFC 793. Internet Engineering Task Force.

Russell, A. L. (2014). *Open Standards and the Digital Age: History, Ideology, and Networks*. Cambridge University Press.

Santoro, J., Jadeja, J., Cuesta Cañada, A., and others (2022). *EIP-4626: Tokenized Vaults*. Ethereum Improvement Proposals.

Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press.

U.S. Securities and Exchange Commission (2025). *Staff Accounting Bulletin No. 122*. SEC, Washington.

Surden, H. (2012). Computable contracts. *UC Davis Law Review*, 46.

Thomson, M. and Schinazi, D. (2023). *Maintaining Robust Protocols*. RFC 9413. Internet Engineering Task Force.

United Nations Commission on International Trade Law (2017). *Model Law on Electronic Transferable Records*. United Nations, Vienna.

International Institute for the Unification of Private Law (2023). *Principles on Digital Assets and Private Law*. UNIDROIT, Rome.

Uniform Law Commission and American Law Institute (2022). *Uniform Commercial Code Amendments (2022), Article 12: Controllable Electronic Records*.

United Kingdom (2023). *Electronic Trade Documents Act 2023*. London.

Vogelsteller, F. and Buterin, V. (2015). *EIP-20: Token Standard*. Ethereum Improvement Proposals.

World Wide Web Consortium (2022). *Verifiable Credentials Data Model v1.1*. W3C Recommendation.

World Wide Web Consortium (2022). *Decentralized Identifiers (DIDs) v1.0*. W3C Recommendation.

A Domain Analyses

This appendix carries the full domain-by-domain analysis summarised in Section 8. Each operational domain is examined through the same five questions: what AI can plausibly automate; what requires standards; what requires governance; what requires legal authority; and what is best left to market convention. The registers of Section 2.4 apply throughout: capability claims are evidence where current systems demonstrably perform the task and inference where they extrapolate from the trajectory, while statements about where authority sits are institutional fact. Several entries condense two adjacent domains into one discussion where the analysis is shared.

A.1 Data mapping

Data mapping is the strongest case for automation in this report, and the concession should be full. Given examples, documentation, or merely the data itself, current models propose field-level mappings between heterogeneous schemas with high accuracy, flag unmapped residues, and generate transformation code with tests. By inference from trajectory, expect near-total automation of mapping *proposal*. What still requires standards is the target and the truth: a canonical model to map into, so that n systems require n mappings rather than n^2, and golden datasets, curated input-output pairs with agreed answers, against which any proposed mapping is scored, because without an agreed answer key, accuracy is unmeasurable and approval reduces to impression. Governance is required for the approval and change control of mappings that feed regulated processes: a mapping into a regulatory report is a model in the supervisory sense and belongs under model-risk discipline (Federal Reserve and OCC, 2011). Legal authority is rarely engaged, which is exactly why this domain automates well. Market convention supplies the canonical models themselves, and the report's expectation is that conventions will consolidate around the ontologies of Section 9 precisely because AI mapping makes a good target model more valuable, not less.

A.2 Process mapping

Processes are harder than data because they embed timing, conditionality, and obligation. AI can plausibly automate the *discovery* and *drafting* of process maps: mining logs and communications to reconstruct how a workflow actually runs, expressing it in executable form, and diffing it against how the rulebook says it runs, a capability with immediate audit value. What requires standards is the event vocabulary: a shared taxonomy of lifecycle events and states, of the kind the Common Domain Model supplies for derivatives (FINOS and ISDA, n.d.), without which two firms' process maps are incommensurable even when both are accurate. Governance is required for change: deadlines, cut-offs, and sequencing in market processes are collective goods, and altering them is a coordination event, not a local optimisation. Legal authority fixes the points where process steps have legal effect, notice periods, record dates, finality points. Convention fills the rest, and much of operations is convention; the reasonable expectation is that AI-discovered practice will feed convention formation faster than committees ever did, the discovery dynamic of Section 4.6 operating in its legitimate habitat.

A.3 Legal interpretation

Models read legal text with real competence: they retrieve, summarise, compare drafts, surface tensions between clauses, and answer doctrinal questions with useful accuracy. All of this is assistance, and its value is large. What none of it produces is an interpretation that *binds*. Legal language has open texture; the hard cases are hard not because readers lack capability but because the text underdetermines the answer and interests diverge over which answer obtains (Hart, 1961). The gap is closed, in every legal system, by authority: a court, a tribunal, a determinations body. Standards reduce the space in which authority must operate, which is their quiet contribution: standardised definitions and clause libraries, and beyond them computable contracts whose performance conditions are evaluable by machine (Surden, 2012; Clack et al., 2016; Grigg, 2004; ISDA, 2019), shrink the residue of interpretive discretion. Governance allocates the residue: who determines, under what procedure, with what appeal. Legal authority is the domain itself. Market convention, meaning the shared understandings of practitioners, remains the interpretive background that both models and courts consult, and it is worth noting that convention is a corpus, which is why models absorb it so well, and also why models freeze it at training time while convention moves on: model drift and convention drift are not synchronised, a small but genuinely novel source of interpretive risk.

A.4 Contractual reconciliation

Reconciling two parties' records of the same trade, position, or portfolio is among AI's highest-value financial applications, because today it consumes armies and its failures cause margin disputes. Models can match records across formats, classify breaks by probable cause, draft the resolution proposal, and learn each counterparty's systematic quirks. What requires standards is decidability: a break is only *objectively* a break relative to a common representation of the trade, which is what shared domain models provide, and a discrepancy in an uncommon representation is merely a disagreement with extra steps. Governance is required for the dispute path: collateral and margin disputes run under agreed protocols with timelines and escalation precisely because reconciliation sometimes fails and the failure must have a procedure. Legal authority stands at the end of the path. Convention supplies tolerances, the thresholds beneath which differences are absorbed rather than pursued, and tolerance-setting in an AI-reconciled market deserves more attention than it gets: when detection becomes free, every immaterial discrepancy becomes visible, and the market must decide anew what it will agree to ignore. Materiality is a convention, and machines will force it to be an explicit one.

A.5 Risk classification

Classifying instruments, exposures, and counterparties into risk categories looks like a natural language task and is substantially a legal one. AI can plausibly automate the extraction of risk-relevant features from documents and data, the proposal of classifications with reasons, and the wholesale re-screening of portfolios when rules change, a task that today makes regulatory transitions multi-year programmes. What requires standards is the schema of categories itself, asset classes, seniority ladders, product taxonomies, because a classification is a coordinate in a shared space and an idiosyncratic space defeats aggregation. What requires governance is the model that assigns coordinates: classification models are the canonical objects of model-risk management, with validation, monitoring, and challenge (Federal Reserve and OCC, 2011). What requires legal authority is the category boundary, because boundaries carry consequences, capital, eligibility, permission, and are therefore set by regulation, not induced from data; whether a token is an e-money token is decided by MiCA's definitions, not by clustering (European Union, 2023). Convention handles the residual judgement calls, and supervisors have learned to standardise even those where they matter enough, which is what regulatory technical standards are.

A.6 Regulatory reporting

Reporting is the domain where the executable-standards future is already operational, and the division of labour is unusually clean. AI can plausibly automate the assembly of reports from source systems, the detection of anomalies before submission, the reconciliation of rejected reports, and the maintenance of mappings as source systems change. What requires standards is nearly everything else: the identifiers that make reports matchable (CPMI-IOSCO, 2018), the schemas and validation rules that make them processable, and increasingly the reporting logic itself, distributed as executable code so that a thousand firms implement one interpretation rather than a thousand (ISDA, 2024; Bank of England and FCA, 2019). Governance is required for versioning the rules and for the exception process when the rules and reality diverge. Legal authority is the rulebook: the obligation to report, the content required, and the liability for misreporting are statutory. Convention has historically filled the gaps in ambiguous fields, which is precisely the pathology, divergent conventional fillings are why matching failed, that executable rules exist to eliminate. The forward-looking observation is that AI plus executable rules changes the supervisor's production function: when rules are code and reports are data, supervision can move from sampling documents to continuously evaluating conformance, but only if the rules were standardised into evaluable form first. AI makes machine-executable regulation more valuable; it does not make it exist.

A.7 Cross-border compliance: AML, KYC, sanctions, and tax

Financial crime compliance is an adversarial screening problem at planetary scale, and AI's contributions are real: entity resolution across transliterations and shell structures, network analysis, narrative generation for suspicious activity reports, triage of the false-positive floods that consume compliance staffing. Two structural facts nonetheless anchor the domain in standards and authority. First, the reference objects are authority artifacts: a sanctions list is not a fact about the world that inference can recover; it is a legal act, and matching against it is deliberately deterministic because strict-liability regimes make interpretive creativity a hazard rather than a virtue. AI triages around the deterministic match; it must not replace it. Second, the data that screening consumes exists only because standards compel its existence and format: the travel rule's originator and beneficiary information, standardised through a dedicated inter-VASP message format (FATF, 2012; interVASP JWG, 2020); beneficial-ownership registers populated under statutory definitions with numeric thresholds; tax-residency data exchanged under common reporting schemas. In each case the legal authority defines the category (who *is* a beneficial owner), the standard fixes the representation, governance runs the registries and update cycles, AI works the graph, and convention decides how much friction legitimate customers will be asked to bear. Cross-border, the structure repeats with a twist: the same conduct is differently categorised across jurisdictions, and AI can map the divergences exquisitely without being able to reconcile them, because reconciliation of legal categories across borders is treaty-making, the slowest standards process there is.

A.8 Settlement disputes and corporate actions

Settlement disputes test the report's finality thesis in miniature. AI can plausibly automate the assembly of the evidentiary record, the reconstruction of the timeline across systems, the identification of where instruction and outcome diverged, and the drafting of claims. What it cannot do is declare which state of the world is the legally operative one, because that is fixed by the infrastructure's rulebook and the finality regime behind it (European Union, 1998; CPMI-IOSCO, 2012): the register's answer is the answer, not because it is more likely to be accurate but because the system is constructed so that its answer is constitutive. Standards define the fail-management procedures around the edges, buy-ins, penalties, recycling conventions; governance runs the default-management and error procedures; legal authority backs the rulebook; convention prices the frictions.

Corporate actions deserve separate notice because they are simultaneously the industry's most notorious data-quality problem and a perfect specimen of the five-way split. The pain begins upstream, where issuers announce events in unstructured prose that intermediaries re-key divergently; AI attacks exactly this, parsing announcements into structured events, and the value is real and immediate. But the parsed event then meets the parts inference cannot supply: the standardised event taxonomy and message flow that let a custody chain process one representation (ISO, 2013; ECB AMI-SeCo, 2020); the market practice that fixes entitlement conventions and deadlines; the governance question of who bears the loss when an election deadline is missed because a machine-parsed announcement was wrong, a question custody agreements answer in advance because answering it afterwards is litigation; and the legal fact that the entitlement itself flows from the issuer's corporate act and the registrar's record, not from anyone's parse of the press release. The end state worth aiming for, and a genuinely available one, is that issuers publish corporate actions natively in structured, standardised, signed form, so that the parsing problem is not solved but abolished. Note the shape of that end state: it is a standard, and AI's role in reaching it is to make the interim bearable, which slightly and perversely delays it. Palliatives compete with cures.

A.9 Transfer restrictions and beneficial ownership

Transfer restrictions convert securities law into access control, and the token world has made the conversion literal: eligibility logic bound to identity credentials, enforced at transfer time (Lebrun et al., 2021). AI can plausibly automate the monitoring side, screening holders against eligibility criteria continuously, flagging structures that appear designed to defeat restrictions, and reading the restriction language out of legacy documents into executable form. What it cannot manufacture is the eligibility fact itself: whether a holder is a qualified purchaser, an accredited investor, or a permitted transferee is a legal status, evidenced by attestation, not induced from behaviour. The architecture that scales is therefore attestation-based: standardised, verifiable claims issued by accountable attestors (W3C, 2022a; European Union, 2024b), consumed by executable restriction logic, with governance over attestor accreditation and liability when attestations are false. Beneficial ownership is the same structure under adversarial pressure: the concealment strategies, nominees, omnibus accounts, layered vehicles, are designed precisely so that ownership cannot be inferred from records, which is why every serious regime pairs graph analysis with compelled declaration under legal penalty. AI walks the graph brilliantly; the graph's load-bearing edges are declarations; declarations exist because law compels them in standard form. Inference is the complement of attestation, never its substitute, and the ratio between them is set by how adversarial the environment is.

A.10 Off-chain collateral and default events

Collateral management outside the ledger is a domain of schedules, eligibility criteria, haircuts, concentration limits, valuation, and calls, and its tokenized future depends on making those semantics executable. AI can plausibly automate the extraction of eligibility schedules from documentation, continuous monitoring of posted collateral against criteria, valuation ingestion, and the drafting of calls and substitutions. What requires standards is the criteria language itself, eligibility expressed in a form both parties' systems evaluate identically, and the valuation and haircut conventions that make one party's call another party's expected number; central bank eligibility frameworks are the template, being precisely published, versioned, machine-consumable criteria. Governance runs disputes and the update of schedules; legal authority determines the character of the collateral arrangement, title transfer or security interest, and its behaviour in insolvency, which no runtime can alter; convention sets thresholds and timings.

Default events sharpen the point to its finest edge. Whether a default has occurred is the most consequential classification in credit, and it is frequently, irreducibly, a judgement: grace periods, materiality qualifiers, cure rights, cross-default webs. AI can watch every covenant, every payment, every disclosure, and flag the approach of the cliff with superhuman vigilance. It cannot pronounce the fall, because the pronouncement is a legal act with formalities, notices, declarations, and because markets have deliberately institutionalised the judgement where it is systemic: the determinations-committee model exists so that one binding answer, produced by a governed process, settles thousands of contracts identically. Detection is a capability; declaration is an office; and the interval between them, where a machine believes a default has occurred and no authorised person has yet said so, is a new and interesting risk window that documentation will have to learn to govern.

A.11 Waterfalls, fund accounting, and NAV

Distribution waterfalls are algorithms trapped in prose, and liberating them is a genuine AI contribution: models can translate a limited partnership agreement's waterfall into executable form, test it against scenarios, and surface the ambiguities, of which real LPAs contain many, catch-up mechanics and clawback interactions being notorious, that only surfaced historically at the worst possible time, distribution. The five-way split then runs: standards are needed for the reporting and data templates that make one fund's numbers commensurable with another's, work the institutional-investor community has pursued for years through standardised reporting templates, and, prospectively, for reference waterfall semantics against which an executable encoding can be validated; governance is needed for the sign-off of the encoding as the operative one, because between two readings of an ambiguous clause a fiduciary must choose and document; legal authority sits in the LPA itself and in the auditor's attestation; and convention supplies the accounting policies where GAAP and IFRS leave choices. The deep point generalises beyond funds: a NAV is not a computation output, it is an *authoritative act*, struck by an administrator who takes responsibility for it. AI will compute NAVs flawlessly and will still not be able to strike one, for the same reason a clock cannot convene a meeting.

A.12 Tokenized loan servicing

Servicing is where tokenized credit meets the physical world of borrowers, and it composes several prior domains: payment application (reconciliation), covenant monitoring (collateral and default), investor reporting (regulatory reporting's private-market cousin), modifications (legal interpretation). AI can plausibly automate the reading of servicing tapes, the application of payments under documented rules, borrower communications, and the propagation of loan events to token holders. What requires standards is the servicing data model, the disclosure-template lineage from securitisation regulation shows both the feasibility and the value of standardised loan-level data, and the event taxonomy by which a modification, a delinquency, or a payoff is represented identically on and off chain. Governance is required for servicer transitions, the moment the whole apparatus must be legible to a stranger, which is the strongest practical argument for standardised servicing data that exists. Legal authority governs the borrower relationship, consumer protection where applicable, and the enforceability of the token holder's claim through the servicing chain. Convention fills the operational pores. The tokenization-specific observation is that a loan token is a claim on a *process*, and a process is only as transferable as its documentation is standard.

A.13 Model governance

The final domain is reflexive: governing the AI that performs everything above. Here the report's inversion completes itself, because this is a domain that barely existed before machine interpretation and consists almost entirely of new standards demand. AI can plausibly automate parts of its own oversight, generating evaluation suites, monitoring drift, red-teaming other models, but the frame within which that happens is being built, visibly and rapidly, from standard forms: model risk management doctrine extended from quantitative models to learned ones (Federal Reserve and OCC, 2011); statutory obligations of logging, transparency, and human oversight for high-risk systems (European Union, 2024a); management-system standards for AI organisations (ISO/IEC, 2023); risk-management frameworks (NIST, 2023); and, ahead, the certification, provenance, and disclosure regimes of Section 9. What requires governance is the deployment decision itself, which models, for which tasks, under which controls; what requires legal authority is the allocation of responsibility when governed models fail, deployer obligations being the current instrument; what falls to convention is evaluation practice, which is consolidating in real time. The domain is the report's argument in miniature: the technology that was supposed to dissolve standards arrived and immediately generated the fastest-growing standardisation agenda in the industry.