

Programmable Institutions

Raez Lorgat

April 2025

Abstract. Institutions are the last analog layer between contemporary technology and the real economy. Every industry rebuilt by competition was rebuilt because its infrastructure became reproducible. The institutional layer, who can form a company, under what rules, with what obligations, in which jurisdictions, has never crossed that threshold. This essay advances five claims. First, programmability, distinguished from digitization, produces verifiable rules, tamper-evident action chains, and reproducible instances. Second, making the compliance floor structural, determined by the algebra rather than by discretion, removes a class of institutional capture that no legal remedy can reach. Third, cheaper exit across jurisdictions, through formal compliance corridors, creates a sharper jurisdictional-quality signal under named conditions; proving its equilibrium effect remains a political-economy obligation. Fourth, voice, built on the same interface, serves the immobile in ways exit cannot. Fifth, four structural properties, offline-verifiable proofs, node autonomy, append-only history, and a typed human-judgment boundary, are necessary conditions for infrastructure *designed for the governed*. We state the propositions; five companion papers supply the constructions and proofs. The deliberation over which properties must hold by construction is the political question the coming decade will either answer or fail to.

I The last analog layer

Institutions fail in recognizable patterns: across parties, across countries, across decades. The failures do not correlate with who holds office. They correlate with the architecture of the institutions themselves, an architecture designed for a world where every participant was human, every process was manual, and every jurisdiction was a closed system. That world is gone. The institutions designed for it remain, increasingly disconnected from the reality they claim to govern. A constitutional crisis produces statements, not corrections. A regulatory agency captured by the industry it regulates continues to issue rulings. A sovereign-debt restructuring proceeds without the participation of the holders most affected. A pandemic arrives and the international coordination machinery produces exhortations. Each of these is a failure mode of an institutional form; each recurs.

The literature on hollowed-out democratic institutions (Mair 2013; Levitsky & Ziblatt 2018; Runciman 2018) describes the symptoms. The cause is older and more architectural. Checking compliance once a year was reasonable when paper files crossed desks on a cycle of days. Holding elections every four years was reasonable when information moved at the speed of a printing press. Running institutional infrastructure

as a single instance, controlled by a single authority, impossible to reproduce, was reasonable when the alternative was no institution at all. None of these premises holds now.

Every industry rebuilt by competition was rebuilt because its infrastructure became reproducible. When anyone could publish a web page, media stopped being a monopoly. When anyone could process a payment, banking opened. The one layer whose infrastructure has never become reproducible is the institutional layer: who can form a company, under what rules, with what obligations, in which jurisdictions. This is why software-operated entities cannot participate in the economy as principals, why digital assets sit uneasily across jurisdictions whose frameworks remain emergent and incomplete, and why a company that succeeds in one jurisdiction starts from zero in another.

What happens when that layer becomes programmable?

2 What programmability means

Digitization changes nothing that matters. Moving a paper form to a screen still means quarterly compliance checks, annual audits, a single authority running a single instance. The feedback is still episodic. The monopoly is intact. Programmability is the different thing. The distinction is not new: Lessig (1999, 2006) and Reidenberg (1998) argued that code itself acts as law when code mediates the transaction; Hildebrandt (2015) developed the case for *legal protection by design*; De Filippi & Wright (2018) trace the technical lineage. Our argument is narrower. Programmability of the *institutional interface*, the rules that admit or refuse an actor's action, has three consequences that digitization does not produce.

First, the rules become verifiable. They are published as programs, identified by their contents via cryptographic hashing in the tradition of Merkle (1980, 1987), Haber & Stornetta (1991), and Nakamoto (2008). Change a single provision and the program's fingerprint changes. The chain of institutional actions becomes tamper-evident: any alteration leaves a detectable trace in the hash-linked log, provided some party re-verifies. Verification is not automatic. It is available.

Second, compliance shifts from episodic to continuous. Every formation, every transfer, every payment is evaluated and recorded as an event in a hash-linked chain. Each link is sealed so no entry can be altered without detection. Each link satisfies every authority that needs to know, simultaneously. Periodic documents are photographs. Continuous attestation is a live feed.

Third, because the software component of the institution is reproducible, the software monopoly dissolves. The marginal cost of another instance approaches zero. Software reproducibility is not institutional reproducibility. A jurisdiction's kernel can be forked at near-zero cost; its legal authority cannot. A fork of the Delaware Division of Corporations is not the Delaware Division of Corporations. What the construction buys is narrower and precise: when an operator corrupts the software, a participant can *verify* the corruption (the rules and the action chain are published and cryptographically bound), and can *choose* to transact with another instance, whether another sovereign jurisdiction or a different operator under the same sovereign, whose kernel has not been corrupted. The construction converts silent, unverifiable capture into visible, verifiable capture, and lowers the software-compliance component of switching cost

from months to seconds. It neither eliminates switching cost entirely nor reproduces sovereign authority. The claim is bounded, not anarchist.

A compact set of *system calls* defines the interface between any actor and any jurisdiction's institutional infrastructure. A working set of five covers the operational surface most entities touch: entity formation, ownership management, fiscal operations, identity verification, and consent. The set is extensional, not principled. Adjudication, registered property transfer, and cross-border reporting are plausible additions. What matters is that the set be finite, typed, and published. The claim is not that five is the right number. The claim is that the institutional layer needs a published interface at all, and does not presently have one. Machine reasoning can parse large portions of published law; it still cannot incorporate a company in any jurisdiction.

A single entity in a single jurisdiction is the first-order case. The same interface supports an entity present in multiple jurisdictions simultaneously, with compliance state composed pointwise across all of them. The corridor architecture that makes this possible is the subject of the companion paper *The Multi-Harbored Institution*.

3 Structural safety and exit from below

If institutions are software, what prevents the software from being corrupted?

The compliance architecture evaluates every actor across a working taxonomy of twenty-three regulatory domains: anti-money laundering, know-your-customer, sanctions, tax, securities, corporate governance, custody, data privacy, licensing, banking, payments, clearing, settlement, digital assets, employment, immigration, intellectual property, consumer protection, arbitration, trade, insurance, anti-bribery, and Sharia compliance. The taxonomy is pragmatic, not principled: it is fixed at evaluation time and extended by schema evolution as new domains (climate, biodiversity, model governance) enter the institutional agenda. The algebra is independent of the taxonomy. The taxonomy is an empirical choice.

The verdicts in each domain form a totally ordered chain: *NonCompliant* < *Pending* < *NotApplicable* < *Exempt* < *Compliant*. Composed pointwise across domains, the evaluator takes the *meet*: the lowest verdict. The overall verdict is the most restrictive component. An actor begins at Pending everywhere, and Pending does not pass. No action proceeds until the actor has been affirmatively cleared in every applicable domain. A sanctions license, humanitarian carveout, sovereign exemption, or shared-authority recognition must enter before the sanctions evaluator returns its verdict, as rule-scope or applicability evidence. Once the applicable Sanctions coordinate returns `NonCompliant`, the aggregate is blocked by the meet. This is a structural property of the algebra, not a policy choice.

Proposition 3.1 (Structural non-overrideability). *Let V be a totally ordered finite verdict lattice with bottom \perp (`NonCompliant`) and a fixed clearance threshold $v^* \succ \perp$. Let \mathcal{D} be the applicable domain set and $T: \mathcal{D} \rightarrow V$ the actor's verdict tensor. Let a kernel admit an action a only if $T_d(a) \succeq v^*$ for every $d \in \mathcal{D}$. Then (i) no single-domain verdict of \perp can be compensated by verdicts in other domains; (ii) no evaluator operator ω distinct from*

the published rule function can reverse the verdict without changing the rule function's cryptographic fingerprint. The conclusion is mechanical: a property of the algebra, not of the policy.

Sketch. (i) Let $d^* \in \mathcal{D}$ with $T_{d^*}(a) = \perp \prec v^*$. Admissibility requires $T_{d^*}(a) \succeq v^*$, which fails. The pointwise universal quantifier admits no cross-domain compensation by construction. (ii) Any ω that maps $(T, \perp) \mapsto \top$ is not a function of the published rule module; it is a second rule module. Content-addressing binds the module's hash to its bytes (Merkle 1980; Haber & Stornetta 1991); substituting ω changes the hash. A code change whose fingerprint differs from the declared fingerprint is detectable by any party that re-verifies. \square

The Applicable-fragment verdicts therefore form a bounded distributive lattice. The structure extends to a Heyting algebra on each per-domain Applicable chain: for any current state a and bound b , the residual $a \rightarrow b$ is the weakest additional constraint that can be composed with a while remaining within b . It is Heyting implication, not remediation. A distinct operator on the composed tensor, the remediation operator, identifies the authorized state transitions, fresh attestations, proof refreshes, or discretion-hole fills required to change the underlying facts and then re-evaluate. This essay uses only those scoped consequences.

The domains are not independent. A sanctions finding forces re-evaluation of banking, payments, clearing, and settlement. An anti-money-laundering failure propagates to KYC and data privacy. Tax changes reach licensing. These causal dependencies form a *propagation graph*: a directed map of which domains trigger re-evaluation of which others. The algebra computes the correct verdict at any snapshot. The propagation graph determines when the snapshot must be recomputed. Together they yield static correctness (the answer is right given current information) and dynamic completeness (the answer is recomputed when relevant information changes).

A subsequent administration cannot reverse Theorem 3.1 by executive order. It is a property of the design, in the way a bridge's load capacity is a property of its engineering. The tradition of *security by design*, Saltzer & Schroeder (1975) on protection of information, Bell & LaPadula (1973) on mandatory access control, is the nearest analog in systems research. Making safety structural rather than discretionary, on the institutional layer, is the contribution.

Structural safety is necessary, not sufficient. Someone still writes the rules, selects the domains, runs the infrastructure. In an analog system, that someone holds unchecked power: one instance, one operator, no alternative. Corruption has no remedy. The argument for *exit* as accountability is Hirschman's (1970): when exit is available, organizations face a mechanical pressure to correct. The argument for inter-jurisdictional exit as a specific form of that pressure is older still: Tiebout (1956) on local public goods, Oates (1972) and Fischel (1975) on fiscal federalism, Romano (1985, 1993) on charter competition in corporate law, with Cary (1974) and Bebchuk & Hamdani (2002) supplying the canonical critique. Frey & Eichenberger's (1996) functional, overlapping, competing jurisdictions (FOCJ) extend the argument to overlapping functional authorities, anticipating the corridor geometry of Section 4.

Programmable institutions extend the argument. Because the infrastructure is software, the cost to deploy a competing instance approaches zero. If an operator corrupts the rules, a competing instance can fork the infrastructure and run it honestly. If a node uses the system for surveillance or political exclusion,

participants with portable compliance move to nodes that do not. The network routes around bad actors through participant exit, *under conditions*.

Proposition 3.2 (Conditional self-correction). *Let K be a captured institutional-infrastructure instance and K' an honest fork of K . Suppose:*

- (1) Software reproducibility. *K' can be deployed at marginal cost $\varepsilon \rightarrow 0$.*
- (2) Legal recognition. *At least one sovereign authority recognizes outputs of K' as valid within its jurisdiction.*
- (3) Passport portability. *A participant's prior compliance state T is cryptographically verifiable by K' without re-evaluation, via the corridor guarantee formalized in The Multi-Harbored Institution.*
- (4) Discoverability. *The existence of K' is published to a channel not controlled by K .*

Then the rate of participant outflow from K to K' strictly exceeds the baseline outflow rate under a captured analog institution with the same participant pool and the same capture.

Conditions (1) and (3) are properties of the construction. Conditions (2) and (4) are institutional conditions the framework does not itself produce. They are the open political questions the framework creates, not answers it supplies.

The claim is weaker than *structurally self-correcting*. It is also the true claim. Hirschman (1970, chapters 4 and 5) modeled the case where exit *degrades* accountability: when the most demanding participants leave first, the organization faces a less exacting remainder, accelerates decline, and the immobile are left stuck. When conditions (2) and (4) fail, exit accelerates rather than repairs the decline of the captured instance. This is the Hirschman residual. The voice mechanisms of Section 5 are complement, not substitute.

In analog institutions, capture is invisible and frequently permanent. Participants cannot verify what the institution is doing, and they cannot build an alternative. In programmable institutions, capture is visible (the rules are published, the chain of actions is tamper-evident) and impermanent under the conditions above. Trust rests on verifiability and reproducibility: the ability to check whether an institution follows its own rules, and the ability to transact with an alternative if it does not.

4 Corridors

Consider running a company in one jurisdiction and wishing to operate in another. In current practice, one starts over. New compliance, new identity verification, new regulatory relationships. Years of operation count for nothing. The compliance history is locked to the jurisdiction that issued it.

A corridor between two jurisdictions changes this. It is a formal bridge specifying which compliance domains the destination re-evaluates from scratch and which it accepts from the origin, enabling an entity to operate in both simultaneously or to migrate entirely. Sanctions are re-evaluated at every general-purpose crossing. Identity verification, filed tax-status evidence, accepted tax facts, and corporate governance attes-

tations may carry forward where the corridor recognizes them. Liability allocation, treaty relief, transfer pricing, withholding, and top-up-tax interactions do not carry by meet; they are separate tax-fiber obligations. Portability reduces duplication without solving tax interaction.

Proposition 4.1 (Exit cost decomposition). *The cost to an entity of moving from jurisdiction J_i to J_j decomposes as*

$$C = C_{\text{compliance}} + C_{\text{legal}} + C_{\text{operational}} + C_{\text{relational}}.$$

The construction reduces $C_{\text{compliance}}$ to the cost of fresh evaluation over the re-evaluation set $R \subseteq \mathcal{D}$, typically a proper subset of the twenty-three domains. For small R , $C_{\text{compliance}}$ approaches API-call economics. The construction does not reduce C_{legal} (local counsel), $C_{\text{operational}}$ (physical presence, local hiring), or $C_{\text{relational}}$ (banking, counterparty trust). The claim that exit becomes fast applies to $C_{\text{compliance}}$, not to C as a whole.

One hundred jurisdictions connected by directed corridors produce up to $100 \times 99 = 9,900$ asymmetric edges. Each corridor is asymmetric (what Singapore accepts from ADGM is not what ADGM accepts from Singapore) because each jurisdiction retains the legal authority to decide what it recognizes. The structure matches bilateral investment treaties, double-taxation agreements, and the EU passporting regime (Investment Services Directive 1993; MiFID II 2014; see also the *Centros* ruling, C-212/97, 1999), but it is computational rather than documentary. Estonia's e-Residency (since 2014) and the OECD BEPS framework (2015) are the nearest living analogs.

The construction preserves sovereignty in a precise sense: every jurisdiction's kernel is the sole writer of its own record, and no corridor can force a verdict on a jurisdiction whose kernel has not agreed to recognize it. It *disciplines* sovereignty in another sense: a jurisdiction that governs badly loses participants through corridors at the switching cost the network has established. Whether the net effect is gain or loss of sovereignty is a political question, live since Slaughter (2004) argued for networked sovereignty and Krasner (1999) described sovereignty as organized hypocrisy. The construction raises the question without answering it.

What the construction provides is structural transparency about what recognition means. Each corridor specifies recognition as a machine-readable mapping, which documentary arrangements do not. The network is not a monopoly because the infrastructure is reproducible. When a dominant node imposes coercive terms, other nodes can implement the same protocol. The 2022 exclusion of designated banks from SWIFT illustrated a concentrated-infrastructure vulnerability. The open-protocol pattern, following the TCP/IP lineage (Cerf & Kahn 1974; RFC 1771, 1995), is a specific alternative.

Once a non-trivial number of jurisdictions participate, a connected jurisdiction offers what an isolated one cannot: portability of accumulated compliance history, regulatory relationships, and institutional credibility. A participant choosing between a connected and an isolated jurisdiction has an asymmetric choice.

Portability creates accountability. If a jurisdiction governs badly, participants with portable compliance can leave, carrying their history to other nodes. Institutional quality faces continuous competitive pressure rather than purely electoral pressure. The race-to-the-bottom concern, classic since Cary (1974) and refined

by Bebchuk & Hamdani (2002), is real above the structural floor. The floor itself is not subject to it. What enters the floor is a political question, addressed in Section 5.

5 Exit, voice, and the double monopoly

We adopt Hirschman's (1970) vocabulary of *exit*, *voice*, and *loyalty*, and argue that programmable institutions are the first infrastructure on which both exit and voice can operate continuously on the same substrate.

Before competitive markets, guilds controlled production. If the bread was bad, there was no recourse. The guild held a double monopoly: over the product, and over the means of production (the ovens, the flour, the market stalls). Buying from a competing baker was impossible because there was no competing bakery. The historical interpretation of guilds is contested: Ogilvie (2019) treats them as rent-extracting cartels; Epstein & Prak (2008) as information-sharing coordinators. Our argument does not turn on which view is correct. It turns on a structural point, that a double monopoly can be dissolved. Markets dissolved this one when feedback became continuous and the infrastructure became reproducible.

Governance today holds a structurally similar double monopoly. A jurisdiction controls the rules. The institutional infrastructure exists as a single, irreproducible instance. Even when exit from the jurisdiction is legally available, exit from the institutional technology is not. There is no competing instance to exit to.

Programmable institutions dissolve both. Corridors enable exit across jurisdictions. Reproducible infrastructure enables exit across operators. Exit alone is insufficient. Exit serves those who can leave. It does little for the poor, the immobile, the rooted, the undocumented. Governance that responds only to exit is governance for the mobile.

Voice matters as much as exit. The same five-call interface serves natural persons and legal entities alike. A participant registered in a single jurisdiction uses the same interface a multi-harbored entity uses: proposals are entity-formation events; votes are consent events; delegation is a transfer of a typed right that is revocable. The verifiability that makes exit meaningful (the chain of actions is tamper-evident) makes voice meaningful in the same way. A vote is attested and cannot later be erased. A delegation is recorded and cannot later be denied. A public allocation is published and can be audited by a party that never leaves the jurisdiction. Privacy-preserving eligibility proofs allow participation without surveillance. Liquid delegation lets a participant entrust a voice on a specific issue to someone they trust, revocable at any time.

Voice is not a substitute for exit. It is the mechanism by which the immobile exercise agency *as the system acknowledges them*. Exit without voice is libertarianism. Voice without exit is monopoly. Both, operating continuously on infrastructure no single actor can monopolize, is what we propose.

Markets without floors degrade through competitive undercutting. Our structural safety properties constrain the most dangerous forms of competition. Above those constraints, competition is open. The *floor* is non-overridable by construction, not by policy. A jurisdiction that participates in the network accepts the floor as a precondition of corridor formation. A jurisdiction that does not participate is not forced to accept the floor. It is outside the network. The floor is an entry condition, not a conquest. What

enters the floor (financial integrity today; labor, environment, human rights as candidates) is the subject of ongoing negotiation between participating jurisdictions, conducted through the same five-call interface. We do not specify the negotiation. We make the negotiation possible by giving it a venue, a vocabulary, and a machine-readable outcome.

Some problems this framework cannot address. Climate, biodiversity, and shared resources in the sense of Ostrom (1990) require collective commitment, and collective commitment requires binding constraints on exit. The planet cannot migrate. The poor cannot migrate. Any serious deployment must pair competitive governance with collective commitments that constrain competition where competition would be destructive.

The framework will not stop armed force or constrain a state that has decided to use it. What it does is make institutional accountability continuous and institutional claims verifiable, on infrastructure anyone can run and no one can monopolize, under the conditions named in Theorem 3.2. What enters the structural floor, and what collective commitments constrain competition above it, are political questions the infrastructure enables but cannot answer.

6 Arrival

Machine-speed institutional operators need institutional interfaces. They cannot call law firms or walk into government offices. A delegated program that can read the compliance state, compute what the algebra requires, and have the system verify it can operate an institution as a principal. Where the computation reaches a boundary that requires human judgment (a risk assessment, a suitability determination, a principle-balancing step), the system marks it explicitly: a *typed proofhole* naming the kind of decision needed and the authority that may supply it. Digital assets accumulate institutional weight while legal frameworks evolve. The compliance model of the analog era has three mechanical failures: periodic evaluations miss events that occur between evaluations; adversarial audits incentivize concealment rather than honest reporting; report-based compliance produces artifacts no downstream consumer can verify. Jurisdictional competition accelerates as each jurisdiction that upgrades raises the bar for every other.

Institutional infrastructure consolidates. Fragmented payment systems became Visa and Mastercard. Regional clearing became DTCC. Interbank messaging became SWIFT. Those systems consolidated around operators: single organizations controlling the infrastructure. The alternative is the open-protocol pattern of consolidation, argued most carefully by Masnick (2019) as *protocols, not platforms*, and traceable through Benkler (2006) and Zittrain (2008). Programmable institutional infrastructure consolidates around open specifications anyone can implement. One pattern is a company controlling all phone calls. The other is an open standard anyone can build a phone for. Early position in defining the protocol matters. Once the standard sets, it is set.

Programmable institutions will arrive. The gap between institutional capacity and what the world demands guarantees it. Who they are designed for is what matters. Infrastructure designed for the powerful will serve the powerful: extraction, surveillance, capture. Infrastructure designed for the governed will serve

the governed: verifiable institutions, reproducible alternatives, continuous accountability, and the ability to participate in and shape the governance that affects one's life.

Designed for the powerful or *designed for the governed* can be made concrete. Four properties either hold by construction or they do not.

- (1) *Offline-verifiable proofs.* Decisions carry proofs any party with standing can check without relying on operator testimony. This is the content-addressed, tamper-evident lineage of Merkle (1980, 1987), Haber & Stornetta (1991), Nakamoto (2008), and Certificate Transparency (RFC 6962, 2013).
- (2) *Node autonomy.* No node has authority over another. A jurisdiction can disconnect from the network without permission from it. This is the design pattern of TCP/IP routing (Cerf & Kahn 1974; RFC 1771, 1995), the X-Road federation used by Estonia and Finland (Estonian Information System Authority, since 2001), and the Cosmos Inter-Blockchain Communication protocol (Kwon & Buchman 2019).
- (3) *Append-only history.* The record admits no silent deletion; a court subpoena returns history intact. This is the tamper-evident log of Haber & Stornetta (1991), now standard in distributed-ledger systems and in the Git content-addressable object store (Torvalds 2005).
- (4) *Typed human-machine boundary.* The boundary between machine judgment and human judgment is named in the rule language itself, so that where computation stops the infrastructure names the authority that must decide.

Proposition 6.1 (Necessary conditions). *For an institutional-infrastructure protocol Π to exhibit the safety floor we call designed for the governed, each of properties (1) through (4) is necessary. Sufficiency is open.*

Sketch. A failure of (1) permits an operator to forge decisions without detection; a failure of (2) permits a coordinator to compel a jurisdiction to recognize a verdict it did not issue; a failure of (3) permits silent rewriting of the past; a failure of (4) reduces the judgment boundary to operational discretion, which is the analog-era failure mode the construction is designed to avoid. Each property is the negation of a specific failure the construction must exclude. \square

Properties (1), (2), and (3) are understood in adjacent systems. Property (4), a typed boundary between machine and human judgment carried in the rule language itself, is the innovation our research programme introduces. Its construction is the subject of the companion paper *Lex: A Logic for Jurisdictional Rules*. The four are derivable from the rest of this essay. Operator-independent programmability (Section 2) requires (1). Sovereignty preserved across corridors (Section 4) requires (2). Continuous accountability (Section 3) requires (3). The typed judgment boundary is what distinguishes an institution from a machine that pretends not to need human judgment.

Existing institutional infrastructures satisfy subsets. SWIFT satisfies (3), partially (1), and neither (2) nor (4). EU eIDAS satisfies (1) and (3), partially (2), with (4) absent. Public blockchains typically satisfy (1), (2), and (3), with (4) absent. No existing infrastructure satisfies all four by construction. A protocol that commits to all four is designed for the governed. A protocol that commits to three, especially one that

leaves (4) unspecified, defers the judgment boundary to operational discretion, which is the failure mode analog-era infrastructures share.

Whether the democratic deliberation that determines the answer occurs before the infrastructure consolidates, or after, is the question that remains.

7 Related thinking

Six literatures underwrite the arguments of this essay.

Institutions as infrastructure. North (1990), Ostrom (1990), and Williamson's transaction-cost tradition establish that institutions are the architecture that governs economic life. Acemoglu & Robinson (2012) synthesize the empirical case. Jensen & Meckling (1976) anchor the agency-cost account of the firm on which much of the regulatory-competition literature rests. Institutions-as-software restates the argument at the level of a specific substrate: rules that admit or refuse actions, encoded in a form a machine can evaluate. We do not argue that all institutional content is computational. We argue that the *interface layer* is.

Exit, voice, and regulatory competition. Hirschman (1970) frames exit and voice as the two mechanisms of organizational correction. Tiebout (1956), Oates (1972), and Fischel (1975) develop the fiscal-federalism argument. Romano (1985, 1993) applies it to corporate charters; Cary (1974) and Bebchuk & Hamdani (2002) press the critique that regulatory competition produces a race to the bottom absent a floor. Frey & Eichenberger (1996) generalize from territorial to functional overlap. We take the conditional view: exit creates a jurisdictional-quality signal above a floor under adoption, mobility, and information assumptions; below the floor, exit can be a harm. Structural non-overrideability (Theorem 3.1) is the construction that distinguishes the two regimes.

Code as law. Lessig (1999, 2006) and Reidenberg (1998) argued that technical architecture has the force of law where it mediates the transaction. Hildebrandt (2015) developed the case for *legal protection by design*; De Filippi & Wright (2018) examine the implications for blockchain-mediated transactions. Computational law as a research agenda traces through Catala (Merigoux, Chataing, Protzenko 2021), L4 (CCLAW Singapore), and the defeasible-logic tradition of Governatori (2005) and Prakken & Sartor (2015). Our programme formalizes these arguments in the specific setting of institutional interfaces.

Open texture. Hart (1961) argues that legal rules have an *open texture*: a core of settled application and a penumbra where human judgment is required. Dworkin (1977) presses the principles-based alternative. The typed human-judgment boundary is an engineering acknowledgment of the Hart position. The computation is bounded, and the bound is named.

Protocols, not platforms. Masnick (2019) formalizes the distinction between open interoperating systems and operator-controlled platforms. Benkler (2006) and Zittrain (2008) supply the earlier conceptual ground. The argument that the institutional layer consolidates around open protocols rather than operator-controlled platforms is a specific application of this line.

Tamper-evident logs and content-addressed systems. Merkle (1980, 1987), Haber & Stornetta (1991), and Nakamoto (2008) establish the technical machinery for cryptographic evidence that an entry has not been silently altered. Certificate Transparency (RFC 6962, 2013) is the canonical production deployment. Git is the canonical content-addressable object store (Torvalds 2005). The institutional-action chain in the construction follows this lineage.

Four claims position our programme above these literatures: (i) conditional self-correction (Theorem 3.2); (ii) exit-cost decomposition (Theorem 4.1); (iii) the structural floor as a negotiated entry condition, rather than a universal conquest; and (iv) the typed human-machine judgment boundary as a rule-language feature, developed in the companion paper *Lex*. Each is either a refinement of an existing argument or a new construction introduced by this programme.

8 The programme

This essay is the founding thesis of a six-paper programme. The companion papers supply the constructions and proofs this essay cites.

The Multi-Harbored Institution defines the computational object, a legal entity existing simultaneously in multiple jurisdictions with compliance state composed across all of them. *Lex: A Logic for Jurisdictional Rules* specifies the rule language, the typed discretion holes, defeasibility, temporal stratification, and the proof obligations the compiler discharges. *Op: A Typed Bytecode for Compliance-Carrying Operations* specifies the workflow substrate. *The Sovereign Jurisdiction Network* specifies the proof-producing kernel, corridor protocol, and proof bundles that pass between jurisdictions. *Intelligent Assets* describes the asset-side consequence: self-executing financial instruments whose compliance rules travel with them across corridors.

The programme states propositions the present essay introduces informally. Proofs are the responsibility of the papers that follow. Our ambition is to identify the problem correctly, to name the conditions under which the proposed solution works, and to say plainly what remains open.

References

- Acemoglu, D. & Robinson, J. A. (2012). *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown.
- Bebchuk, L. A. & Hamdani, A. (2002). “Vigorous Race or Leisurely Walk: Reconsidering the Competition over Corporate Charters.” *Yale Law Journal*, 112(3), 553-615.
- Bell, D. E. & LaPadula, L. J. (1973). “Secure Computer Systems: Mathematical Foundations.” MITRE Technical Report 2547.
- Benkler, Y. (2006). *The Wealth of Networks*. Yale University Press.
- Cary, W. L. (1974). “Federalism and Corporate Law: Reflections upon Delaware.” *Yale Law Journal*, 83(4), 663-705.

- Cerf, V. G. & Kahn, R. E. (1974). "A Protocol for Packet Network Intercommunication." *IEEE Transactions on Communications*, 22(5), 637-648.
- De Filippi, P. & Wright, A. (2018). *Blockchain and the Law*. Harvard University Press.
- Dworkin, R. (1977). *Taking Rights Seriously*. Harvard University Press.
- Epstein, S. R. & Prak, M., eds. (2008). *Guilds, Innovation, and the European Economy, 1400-1800*. Cambridge University Press.
- European Commission. (2014). *Directive 2014/65/EU (MiFID II): Markets in Financial Instruments Directive*. Official Journal of the European Union.
- Fischel, W. A. (1975). *The Economics of Zoning Laws: A Property Rights Approach to American Land Use Controls*. Johns Hopkins University Press.
- Frey, B. S. & Eichenberger, R. (1996). "FOCJ: Creating a Single European Market for Governments." In *Europe's Constitutional Future*, pp. 195-215. Institute of Economic Affairs.
- Governatori, G. (2005). "Representing Business Contracts in RuleML." *International Journal of Cooperative Information Systems*, 14(2-3).
- Haber, S. & Stornetta, W. S. (1991). "How to Time-Stamp a Digital Document." *Journal of Cryptology*, 3(2), 99-111.
- Hart, H. L. A. (1961). *The Concept of Law*. Oxford University Press.
- Hildebrandt, M. (2015). *Smart Technologies and the End(s) of Law*. Edward Elgar.
- Hirschman, A. O. (1970). *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Harvard University Press.
- Jensen, M. C. & Meckling, W. H. (1976). "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure." *Journal of Financial Economics*, 3(4), 305-360.
- Krasner, S. D. (1999). *Sovereignty: Organized Hypocrisy*. Princeton University Press.
- Kwon, J. & Buchman, E. (2019). "Cosmos Whitepaper: A Network of Distributed Ledgers." Interchain Foundation.
- Lessig, L. (1999, 2006). *Code and Other Laws of Cyberspace (v2, 2006)*. Basic Books.
- Levitsky, S. & Ziblatt, D. (2018). *How Democracies Die*. Crown.
- Mair, P. (2013). *Ruling the Void: The Hollowing of Western Democracy*. Verso.
- Masnack, M. (2019). "Protocols, Not Platforms: A Technological Approach to Free Speech." Knight First Amendment Institute, Columbia University.
- Merigoux, D., Chataing, N. & Protzenko, J. (2021). "Catala: A Programming Language for the Law." *Proceedings of the ACM on Programming Languages*, 5(ICFP).
- Merkle, R. C. (1980). "Protocols for Public Key Cryptosystems." *Proceedings of the 1980 IEEE Symposium on Security and Privacy*, pp. 122-134.
- Merkle, R. C. (1987). "A Digital Signature Based on a Conventional Encryption Function." In *Advances in Cryptology, CRYPTO '87*, pp. 369-378. Springer LNCS 293.
- Nakamoto, S. (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System."
- North, D. C. (1990). *Institutions, Institutional Change, and Economic Performance*. Cambridge University Press.

- Oates, W. E. (1972). *Fiscal Federalism*. Harcourt Brace Jovanovich.
- OECD. (2015). *Addressing Base Erosion and Profit Shifting: Final Reports*. OECD Publishing.
- Ogilvie, S. (2019). *The European Guilds: An Economic Analysis*. Princeton University Press.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.
- Prakken, H. & Sartor, G. (2015). "Law and Logic: A Review from an Argumentation Perspective." *Artif. Intell.*, 227, 214-245.
- Rekhter, Y. & Li, T. (1995). *A Border Gateway Protocol 4 (BGP-4)*. RFC 1771.
- Reidenberg, J. R. (1998). "Lex Informatica: The Formulation of Information Policy Rules through Technology." *Texas Law Review*, 76(3), 553-593.
- Republic of Estonia. (2014). *e-Residency of Estonia*. Government of the Republic of Estonia.
- Romano, R. (1985). "Law as a Product: Some Pieces of the Incorporation Puzzle." *Journal of Law, Economics, and Organization*, 1(2), 225-283.
- Romano, R. (1993). *The Genius of American Corporate Law*. AEI Press.
- Runciman, D. (2018). *How Democracy Ends*. Profile Books.
- Saltzer, J. H. & Schroeder, M. D. (1975). "The Protection of Information in Computer Systems." *Proceedings of the IEEE*, 63(9), 1278-1308.
- Slaughter, A.-M. (2004). *A New World Order*. Princeton University Press.
- Tiebout, C. M. (1956). "A Pure Theory of Local Expenditures." *Journal of Political Economy*, 64(5), 416-424.
- Torvalds, L. (2005). Git: a distributed version-control system.
- Zittrain, J. (2008). *The Future of the Internet, and How to Stop It*. Yale University Press.