

Trustworthy Electronic Business Collaboration in Supply Chain Finance: A Convergent Architecture of Distributed Ledgers, Privacy-Preserving Computation, Trusted Execution, and Digital Credentials

Guang Chen¹ Yu Song²

¹University of International Business and Economics, Beijing, China ²North China University of Technology, Beijing, China

¹University of International Business and Economics, No. 10 Huixin Dongjie, Chaoyang District, Beijing 100029, China

²North China University of Technology, No. 5 Jinyuanzhuang Road, Shijingshan District, Beijing 100144, China

Abstract—Supply chain finance (SCF) promises to convert the creditworthiness of large anchor buyers into liquidity for their smaller suppliers, yet in practice the same information asymmetries that create the financing gap also prevent its closure. This paper argues that the binding constraint is no longer the availability of digital technology but the absence of an assurance architecture capable of making a receivable verifiable, singular, and enforceable across organisational and jurisdictional boundaries. We develop and defend a convergent architecture—here termed TCCE (Trust Anchor, Credential, Confidential Computing, and Execution layers)—that composes four technological families rather than treating them as substitutes: permissioned distributed ledgers and smart contracts as a shared registry and deterministic state machine; privacy-preserving computation (secure multi-party computation, federated learning, homomorphic encryption, zero-knowledge proofs) as the mechanism that keeps raw commercial and personal data inside its originating domain; trusted computing (hardware roots of trust, remote attestation, trusted execution environments) as the mechanism that makes computation itself auditable; and digital credentials (decentralised identifiers, verifiable credentials, reliable electronic signatures, and electronic transferable records) as the bridge between cryptographic state and legal fact. Using accounts-receivable electronic certificates (ARECs) in China as the empirical setting—an instrument whose outstanding balance reached RMB 2.75 trillion and whose financing balance reached RMB 2.08 trillion in the first quarter of 2026 under a newly tightened regulatory regime—we show that duplicate financing, unverifiable trade backgrounds, failed credit pass-through, and cross-platform non-interoperability are four manifestations of a single underlying defect: the absence of a continuously verifiable link between a business event, a legal right, and a payment obligation. We formalize this defect, derive conditions under which verifiable disclosure strictly dominates both full concealment and full disclosure, specify an implementation blueprint with lifecycle controls, and evaluate the architecture against a stated adversary model, Chinese and international compliance requirements, and unit-economics constraints. The paper contributes a design theory for trustworthy inter-organisational collaboration that is simultaneously technically executable, legally grounded, and economically sustainable, and it identifies why a substantial share of prior trade-finance consortia failed on governance and adoption rather than on cryptography.

Index Terms—trustworthy electronic business collaboration; supply chain finance; accounts-receivable electronic certificates; blockchain and smart contracts; privacy-preserving computation; trusted execution; verifiable credentials; electronic transferable records; duplicate financing; inter-organisational governance

I. INTRODUCTION

A. The paradox of supply chain finance

Supply chain finance rests on a simple and compelling proposition. A small supplier sells to a large, creditworthy buyer and waits sixty or ninety days for payment. The supplier's cost of capital is high; the buyer's is low. If the

buyer's payment obligation can be made visible and verifiable to a financier, the supplier should be able to discount its receivable at a rate close to the buyer's cost of funds rather than its own, and the entire chain should become cheaper to finance. In the language of operations-finance interface research, transparency about a supplier's order book and payment record can redistribute credit along the chain and relax the supplier's financial constraint [11].

The proposition is elegant. It is also, in practice, frequently unrealised. Three decades of research on SCF document that adoption concentrates in the first tier of the supply chain, that the smallest suppliers are the least likely to benefit, and that the instrument most often used to signal creditworthiness—the anchor buyer's confirmation of a payable—is precisely the artefact most easily withheld, delayed, fabricated, or duplicated [55], [23], [34]. The literature on factoring has long observed that the value of receivables finance depends less on the arithmetic of discounting than on the institutional quality of verification and notification [29].

This paper starts from the observation that the gap between promise and practice is not primarily a technology-availability gap. Distributed ledgers, digital signatures, secure computation, and hardware enclaves are all commercially available and have been deployed, with varying success, in trade finance since at least 2016. The gap is architectural: deployments tend to bolt a single technology onto an unchanged institutional process and then discover that the bottleneck has moved rather than disappeared. A blockchain makes a record tamper-evident but cannot establish that the record corresponds to a real shipment. Privacy-preserving computation prevents data leakage but cannot establish who is legally bound. An electronic signature binds a signatory but says nothing about whether the underlying receivable has already been financed elsewhere.

B. The setting: accounts-receivable electronic certificates in China

We ground the argument in a setting where the stakes, the scale, and the regulatory response are all unusually visible: the market for accounts-receivable electronic certificates (ARECs; 应收账款电子凭证) in the People's Republic of China.

An AREC is an electronic record by which a debtor—typically the anchor buyer in a supply chain—acknowledges a payable arising from a genuine trade relationship and undertakes to pay on a stated date. It is issued through a supply-chain information service system, can be transferred

and, in many implementations, split into smaller denominations, and can be presented to a bank for discounting. As of the end of the first quarter of 2026, ARECs outstanding reached RMB 2.75 trillion with a financing balance of RMB 2.08 trillion, a quarterly issuance flow of RMB 1.08 trillion, a quarterly financing flow of RMB 0.83 trillion, and an average financing rate of 2.71 per cent, reported under the self-regulatory framework of the National Internet Finance Association of China (2026). For scale comparison, China's above-scale industrial enterprises carried RMB 26.06 trillion of accounts receivable at the end of 2024 with an average collection period of 65.6 days, up from RMB 23.72 trillion and 60.6 days a year earlier (National Bureau of Statistics, 2024, 2025), while enterprises issued RMB 38.3 trillion of commercial bills during 2024 and year-end bill financing balances stood at RMB 14.7 trillion, or 5.8 per cent of all loans (People's Bank of China, 2025).

Three features make this setting theoretically important rather than merely large.

First, it is *contested*. Whether an AREC is a negotiable instrument, an assignment of a receivable, or a factoring arrangement is not settled, and the answer determines registration, recourse, and insolvency treatment. In 2025 regulators intervened directly: the *Notice on Regulating Supply Chain Finance Business and Guiding Supply Chain Information Service Institutions to Better Serve SME Financing* (Yinfa [2025] No. 77, issued 26 April 2025, effective 15 June 2025) defined an AREC, capped its payment term in principle at six months and absolutely at one year, required registration of AREC financing in the unified movable-asset financing registration system, restricted the layering and frequency of split transfers, and prohibited information service institutions from pooling or misappropriating settlement funds. In parallel, the revised *Regulation on Ensuring Payment to Small and Medium-Sized Enterprises* [42] obliges large enterprises to pay within sixty days of delivery, forbids conditioning payment on receipt of payment from a third party, and forbids compelling suppliers to accept non-cash instruments or using such instruments to extend effective payment terms.

Second, it is *technologically heterogeneous*. Platforms differ in whether they use a distributed ledger, whether splitting is permitted, whether financing is open to multiple banks or captive to one, and whether trade documentation is verified against invoice, tax, logistics, and customs data.

Third, it is *institutionally diagnostic*. Because the regulator has now specified what must be true for an AREC to be financeable—genuine trade, bounded tenor, controlled splitting, unified registration, ring-fenced settlement—the setting allows us to ask a sharp question: which technical mechanisms can produce the specific forms of evidence that the rule demands, and which cannot?

C. Contributions and roadmap

The paper makes four contributions.

Conceptual. We decompose “trustworthy electronic business collaboration” into six properties—authenticity of

identity, singularity of rights, integrity of records, confidentiality of inputs, enforceability of obligations, and auditability of process—and show that no single technology delivers more than a subset. We then state an adversary model that makes explicit what the architecture does *not* defend against.

Architectural. We specify TCCE, a five-layer reference architecture that assigns exactly one assurance responsibility to each layer and defines the interfaces between them. The central design claim is a separation of duties: the ledger holds *commitments and state transitions*, the credential layer holds *identity and authority*, the confidential-computing layer holds *raw data*, the trusted-execution layer makes *computation attestable*, and the settlement layer completes *value transfer*.

Analytical. We formalise duplicate financing as a failure of singularity and derive conditions under which selective, verifiable disclosure weakly dominates both full concealment and full disclosure for both the financier and the borrower. We also show why credit pass-through along a multi-tier chain is limited not by the cost of computation but by the cost of *verification* of each successive tier's trade event.

Practical. We provide an implementation blueprint mapped to the actual control points demanded by Yinfa [2025] No. 77, an evaluation against security, compliance, and unit economics, and a governance analysis explaining why a number of prominent trade-finance consortia—most visibly TradeLens, discontinued at the end of the first quarter of 2023—failed on adoption and governance rather than on cryptographic correctness.

The remainder proceeds as follows. Section 2 establishes the problem space and the regulatory trajectory. Section 3 develops the trust taxonomy and adversary model. Section 4 presents the architecture. Section 5 details the mechanisms and, crucially, their composition. Section 6 formalises the economics. Section 7 gives the implementation blueprint and field evidence. Section 8 evaluates security, compliance, and cost. Section 9 addresses governance and standardisation. Section 10 concludes with limitations and a research agenda.

II. PROBLEM SPACE: ELECTRONIC RECEIVABLES AND THE FOUR STRUCTURAL FRICTIONS

A. What an accounts-receivable electronic certificate is, and what it is not

Precision about legal character is not pedantry; it determines which control mechanisms are available. Chinese law provides three distinct regimes that an AREC can fall into, and they carry different consequences.

Negotiable instrument. If an instrument qualifies as a bill of exchange under the *Law on Negotiable Instruments*, it is subject to strict formality requirements: amounts must be recorded in both Chinese capital numerals and Arabic numerals (Art. 8), material particulars such as amount, date, and payee may not be altered (Art. 9), and formats are standardised by the People's Bank of China (Art. 108). Instruments satisfying these requirements benefit from abstraction (*Wuyinxing*): the holder's right is largely severed from the underlying transaction, and recourse follows the

endorsement chain. The Shanghai Commercial Paper Exchange's supply-chain bill platform, launched on a trial basis on 24 April 2020 with the approval of the People's Bank of China, issues electronic bills within this regime and introduced *equal-denomination issuance*—the ability to split a bill into standardised fractions—precisely to accommodate supply-chain settlement patterns. First-day activity was modest: seventeen enterprises issued seventeen instruments totalling RMB 1.0442 million, and two enterprises endorsed three instruments totalling RMB 51,000.

Assignment of a receivable or factoring. If the instrument is instead evidence of a receivable that is assigned or factored, the *Civil Code* governs. Article 761 defines factoring; Article 763 protects a non-knowing factor against a debtor who colluded in fabricating the receivable; Article 764 requires the factor to identify itself and attach necessary evidence when notifying the debtor; and Article 768 supplies the priority rule that matters most for our purposes: among competing factors over the same receivable, one who has registered prevails over one who has not; among those registered, priority follows registration sequence; among those notified but unregistered, priority follows the order in which notices reach the debtor; and failing both, claims are distributed pro rata.

A distinct regulatory object. Yinfa [2025] No. 77 constructs a third category for supervisory purposes, defining an AREC functionally—an electronic record in which a debtor, on the basis of a genuine trade relationship and through a supply-chain information service system, undertakes to pay a creditor in the chain on time—without purporting to settle its private-law character. The Notice simultaneously encourages the expansion of supply-chain bills. The resultant structure is best described as *dual-track regulation with mutual constraint*: supply-chain bills enjoy unified registration and rediscount infrastructure; ARECs are closer to the procurement and payables context but are now bounded by tenor, splitting, registration, settlement, and data-protection requirements.

Two practical consequences follow, and both matter for system design. First, the English rendering of these instruments must be disciplined. A generic, splittable accounts-receivable undertaking should not be translated as an “electronic negotiable instrument” or “bill” unless the facts establish the formalities, acceptance, endorsement, and recourse structure of the instrument regime; otherwise the translation itself creates legal and regulatory misunderstanding across borders. Second, and more importantly for this paper, the *registration* requirement in Article 768 of the *Civil Code* and in Yinfa [2025] No. 77 is not a compliance footnote. It is the legally recognised mechanism for establishing priority over a receivable that may otherwise be financed repeatedly. Any technical architecture that does not produce a registration event, linked to a uniquely identified underlying receivable, leaves the central risk of the market unaddressed.

B. Friction one: duplicate financing as a failure of singularity

Duplicate financing occurs when the same receivable supports more than one extension of credit because no participant can establish that it has already been financed. It is not a software bug; it is the direct consequence of a receivable being an intangible asset whose “delivery” cannot be observed.

The mechanics are easy to describe and hard to prevent. A supplier holds a RMB 10 million receivable against an anchor buyer. It obtains a RMB 8 million factoring advance from Bank A against that receivable. It then presents the same receivable, perhaps with slightly different documentation or on a different platform, to Bank B, or it splits the receivable into ten certificates of RMB 1 million each and presents several of them to different financiers. Each financier, examining only its own records, sees a plausible, internally consistent transaction. The fraud is discovered only at maturity, when the anchor buyer pays once and two claims appear.

Yinfa [2020] No. 226 already identified the risk, requiring financial institutions to guard against inflated or fabricated receivables and inventories and against repeated pledges. Yinfa [2025] No. 77 tightened the response by requiring registration in the unified movable-asset financing registration system, by requiring institutions offering split transfers to control the number of transfer layers and transactions and to investigate, flag, and report anomalous splitting, and by prohibiting banks from financing split certificates whose debtor-creditor relationships are unclear.

Registration helps but does not close the gap by itself, for two reasons. Registration is typically performed *at the time of financing*, so it reveals a prior security interest but not a prior unregistered transfer. And registration in one system does not automatically reconcile with records held by another platform, another bank, or another core enterprise's ERP. The residual risk is therefore a *cross-platform singularity* problem: the same economic receivable must be identifiable as the same receivable in every system that could ever finance it.

C. Friction two: unverifiable trade background

The second friction is that an anchor buyer's confirmation is necessary but not sufficient. A digital signature over a payment undertaking establishes who signed, when, and that the content has not been altered; it does not establish that goods were ordered, produced, shipped, received, invoiced, and accepted.

Chinese law is explicit on this point, and the judicial treatment of blockchain evidence is correspondingly careful. Article 13 of the *Electronic Signature Law* conditions “reliable electronic signature” status on four cumulative requirements: the signature-creation data is, at the time of signing, exclusive to the signatory; it is under the signatory's sole control at the time of signing; any alteration of the signature after signing is detectable; and any alteration of the content or form of the data message after signing is detectable. Article 14 then accords a reliable electronic signature the same

legal effect as a handwritten signature or seal. The Supreme People's Court's *Rules on Online Litigation* (effective 1 August 2021) take the same structure one step further in the evidentiary domain: Article 16 permits a court to accept that electronic data stored using blockchain technology and verified as technically consistent has not been tampered with *after* being recorded on the chain, subject to rebuttal; Article 17 requires the court to examine comprehensively the compliance of the depositary platform, the interest relationship of the parties, the security of the system, the technical standard applied, and—critically—the *authenticity of the data before it was placed on the chain*.

The doctrinal point generalises well beyond China. A hash anchor proves posterior integrity, not anterior truth. An architecture that treats “on-chain” as a synonym for “true” has simply relocated the oracle problem without acknowledging it. What is required instead is a *chain of custody for business events*: order, acceptance of order, production or shipment, delivery confirmation, invoice, acceptance, payment undertaking, and payment, each signed by a different party, each timestamped, and cross-referenced so that a missing or inconsistent link is detectable before financing rather than after default.

D. Friction three: failed credit pass-through

The third friction is that SCF benefits decay sharply with distance from the anchor buyer. The mechanism is well understood in the operations literature: disclosure of a supplier's transaction history with a creditworthy buyer can substitute for the supplier's own balance sheet, but only for the tier that actually transacts with that buyer [11], [9].

Consider a four-tier chain. The anchor buyer (Tier 0) knows Tier 1 and perhaps Tier 2. Tier 3's supplier may produce a component that is functionally essential but organisationally invisible. For a financier to extend Tier 3 credit on the strength of the anchor buyer's credit, it would need to establish (a) that Tier 3's output is actually incorporated into what the anchor buyer purchased, (b) at what quantity and price, and (c) that the corresponding receivable has not already been satisfied upstream. Each of these is a verification task over data held by a party that has no contractual relationship with the financier and often a commercial interest in not sharing.

Two design mistakes recur. The first is to respond by demanding more data, which increases resistance precisely where it is highest and collides with data-minimisation obligations: Article 6 of China's *Personal Information Protection Law* and the legality-necessity-proportionality principle restated in Yinfa [2025] No. 77 both require processing to be confined to what is necessary for a specified purpose. The second is to respond by concentrating all data in a single platform, which creates a single point of failure, a single point of commercial leverage, and—as Yinfa [2020] No. 226 anticipated when it prohibited closed-loop platform circulation and captive financing arrangements—a structure that regulators will eventually constrain.

E. Friction four: non-interopability and the governance trap

The fourth friction is that even well-designed single platforms do not compose. Observers of the sector describe the data-sharing problem as a triple bind: participants are *unwilling* to share because price and customer information confers bargaining power; *unwilling to risk* sharing because of personal-information, trade-secret, and credit-reporting liability; and *unable* to share because ERP systems, tax and invoice formats, logistics identifiers, bank interfaces, and authorisation models do not align.

A substantial review literature has documented the same pattern from the organisational side: distributed-ledger initiatives in supply chains tend to be evaluated on technical criteria while the organisational, inter-organisational, and institutional conditions that determine value are treated as background [53], [12], [38], [30]. Early system-level demonstrations showed that traceability architectures are feasible [48] and that trade-finance structures such as reverse securitisation can be redesigned around them [25], while studies of enterprise adoption repeatedly find that interoperability, standards, and organisational readiness, rather than protocol performance, determine outcomes [26]. The international record is instructive and, for our purposes, more valuable than any success story. TradeLens, the shipping and documentation network jointly promoted by Maersk and IBM, was discontinued at the end of the first quarter of 2023; the stated reason was the failure to achieve the global industry collaboration required for full adoption and the failure to reach the commercial viability expected of an independent business. Several bank-consortium trade-finance networks, including we.trade and Marco Polo, contracted or restructured, and the Voltron initiative was folded into Contour. The precise dates, terms, and rationales of these consolidations vary and are not always documented in primary announcements, so we treat them as directional evidence rather than as a dataset.

What the pattern supports is a claim about causal ordering. These ventures did not fail because consensus algorithms were inadequate or because hashing was too slow. They failed because five conditions were not satisfied simultaneously: neutral multi-party governance; a large enough adoption network to amortise fixed integration and compliance costs; legal certainty about the status of the electronic records the platform produced; open interfaces that allowed participants to join and leave without rebuilding; and a unit-economics model in which each participant captured net value rather than merely absorbing cost. Technology was necessary in every case and sufficient in none.

F. Synthesis: one defect, four symptoms

The four frictions share a structure. In each case, a business event occurs in one organisation's domain; a legal right is supposed to follow from that event; a financier must act on the right; and the linkage between the three is not continuously verifiable by the party that bears the risk. Duplicate financing is the failure to verify singularity; unverifiable trade background is the failure to verify anterior truth; failed credit

pass-through is the failure to verify multi-tier linkage; non-interoperability is the failure to verify across governance domains.

If the diagnosis is right, the prescription is not “more blockchain” or “better machine learning.” It is an architecture in which each verification obligation is assigned to the mechanism that can actually discharge it, and in which the outputs of those mechanisms compose into a single auditable artefact. That is the subject of Sections 3 to 5.

III. CONCEPTUAL FOUNDATIONS: WHAT “TRUSTWORTHY” MUST MEAN

A. From institutional trust to verifiable assurance

Trust in commercial exchange has historically been produced by institutions: courts, notaries, banks, credit bureaus, and trade associations. Their function is to reduce the cost of believing a counterparty. Digital technology does not replace these institutions; it changes which assertions can be checked *without* appealing to them. Digital exchange lowers the cost of search, replication, transportation, tracking, and verification, and it is precisely the fall in verification cost that changes which institutional arrangements are viable [24]; where verification remains costly, adverse selection persists and markets for quality thin or collapse [2]. The relevant distinction for design purposes is not between “trusted” and “trustless” systems—a phrase that has done considerable damage by implying that cryptography eliminates the need for accountable parties—but between assertions whose truth can be established locally by inspection and assertions whose truth requires an institutional warrant.

We adopt a deliberately deflationary definition. An electronic business collaboration is *trustworthy to degree d* if a verifier can establish, at a cost below a stated threshold and with a residual error probability below a stated bound, that a defined set of assertions about a transaction hold. This definition has three advantages: it makes trust a matter of degree rather than a binary; it makes cost explicit, which is where most practical designs fail; and it forces the designer to enumerate the assertions.

B. Six properties

We propose six properties, each of which must be achieved by a distinct mechanism because no single mechanism achieves two of them without loss.

P1 — Authenticity of identity. The assertion “this legal person, acting through this authorised natural person, is who it claims to be.” Achieved by: enterprise identity binding to a registered legal entity, authorised-signatory verification, and credentials whose issuer authority is itself verifiable.

P2 — Singularity of rights. The assertion “this receivable is the unique, identified subject of this obligation and has not been otherwise transferred, pledged, or financed.” Achieved by: a unique business identifier, a registry of state transitions, and a legally effective registration mechanism with a defined priority rule.

P3 — Integrity of records. The assertion “this record has not been altered since the moment it was created by its stated

author.” Achieved by: digital signatures, cryptographic hashing, trusted timestamps, and append-only storage.

P4 — Confidentiality of inputs. The assertion “the verifier learned the minimum necessary to reach its conclusion and nothing more.” Achieved by: selective disclosure, zero-knowledge proofs, secure computation, and execution isolation.

P5 — Enforceability of obligations. The assertion “a court or arbitral tribunal will recognise this record as evidence of a right and grant a remedy.” Achieved by: alignment with the applicable legal regime for electronic records, signatures, and transferable records; and by evidence packaging that satisfies evidentiary rules.

P6 — Auditability of process. The assertion “an independent third party can reconstruct what happened, who did it, on what authority, and under what system configuration.” Achieved by: immutable logs, remote attestation of execution environments, and retained evidence packages.

A design that optimises one property at the expense of another is a design with a hidden risk transfer. Full on-chain disclosure maximises P3 and P6 while destroying P4 and potentially breaching data-protection law. Encrypted computation preserves P4 but, absent attestation, weakens P6. Signature-only designs strengthen P3 while leaving P2 untouched—which is exactly the configuration in which duplicate financing flourishes.

C. Adversary model

We state explicitly what the architecture is designed to resist, because unstated assumptions are the most common source of overclaiming in this literature.

A1 — The opportunistic borrower. A supplier that seeks to finance the same receivable twice, or to finance a receivable that does not exist. Capabilities: creating documents, colluding with other firms, using multiple platforms. Non-capabilities: forging signatures of non-colluding parties, breaking standard cryptographic primitives.

A2 — The conflicted anchor. A core enterprise that confirms receivables loosely because confirmation is cheap and reputational exposure is diffuse; that delays confirmation to preserve working capital; or that, in the structure prohibited by Yinfa [2020] No. 226, occupies suppliers’ funds while an affiliate earns interest on financing them. Capabilities: unilateral control over its own ERP data and over confirmation timing.

A3 — The over-reachful platform. An information service institution that selectively presents data, captures financing to an affiliated lender, or resists interoperability to preserve lock-in. Capabilities: control over its own platform’s records, interfaces, and disclosure policies.

A4 — The curious insider and the external attacker. An operator, auditor, or intruder seeking to read commercial or personal data not needed for their function. Capabilities depending on role: database read access, network position, or cloud-administrative privilege.

A5 — The compromised execution environment. An attacker with software or physical access to a compute node, including speculative-execution and side-channel techniques against trusted execution environments. We do *not* assume hardware enclaves are inviolable; the published record of transient-execution attacks against isolation boundaries [52] is precisely why attestation, defence in depth, and cryptographic fallback are required.

Explicitly out of scope. We do not claim to defend against (i) a supplier and anchor buyer who collude to fabricate a complete, internally consistent set of trade documents including third-party logistics confirmation; (ii) a legal system that declines to recognise electronic records; (iii) macroeconomic insolvency of the anchor buyer; or (iv) capture of the governance body itself. Item (i) deserves emphasis: no cryptographic mechanism can substitute for the possibility of physical inspection, and the architecture must therefore lower the cost and increase the frequency of *sampled* physical and documentary verification rather than pretend to eliminate it.

D. Design principles

Five principles follow from the properties and the adversary model.

DP1 — Assertion-based minimisation. Data should be disclosed or computed over only to the extent required to establish a stated assertion, and the assertion—not the dataset—is the unit of design.

DP2 — Separation of registry and transaction. The shared ledger records commitments, state transitions, and proofs. Raw commercial documents, personal data, and model inputs remain in the domain of their controller.

DP3 — Verifiable rather than trusted computation. Any computation performed on data that its owner cannot see must be accompanied by evidence that the agreed computation was performed by the agreed code on the agreed inputs. Attestation and zero-knowledge proofs are the two available instruments.

DP4 — Legal-technical co-design. Every state in the technical state machine must map to a recognised legal event, and every legal requirement (registration, notification, tenor, settlement) must have a technical control point. A smart contract that transfers a token is not a transfer of a receivable; it is evidence that the parties invoked a procedure whose legal effect must be separately established.

DP5 — Contestability. Every automated decision—financing approval, anomaly flag, refusal—must be reproducible from retained inputs, model version, and configuration, and must be challengeable by the affected party. This is both a fairness requirement and, as the following sections argue, an evidentiary one.

IV. TCCE: A CONVERGENT REFERENCE ARCHITECTURE

A. Overview

We propose TCCE, a five-layer reference architecture. The name denotes the four assurance responsibilities it separates—Trust anchor, Credential, Confidential computing, and

Execution and settlement—arranged over a **shared state layer**. The architecture is deliberately not a stack of products; it is an assignment of responsibility, and each layer may be implemented by different technologies in different deployments provided the interface contracts hold.

The organising intuition is a division of labour that mirrors the division of legal labour in a properly structured financing. Someone must establish *who* the parties are (credential layer). Someone must maintain an authoritative record of *what right exists and what has happened to it* (state layer). Someone must perform *judgement* over sensitive data (confidential-computing layer). Someone must vouch for *how the judgement was reached* (trust anchor). Someone must *move the money and discharge the obligation* (execution and settlement layer). Conflating any two of these roles is what produces platforms that are simultaneously overloaded with responsibility and unable to prove anything.

B. The five layers

L1 — Trust anchor layer. Comprises hardware roots of trust, trusted platform modules or trusted cryptography modules, secure boot, measured launch, and remote attestation. Its sole responsibility is to make *computational provenance* auditable: to produce signed evidence that a specific binary was executed on a specific platform in a specific configuration. It is governed in China by GB/T 29829-2022, *Information security technology — Functionality and interface specification of cryptographic support platform for trusted computing* (issued 15 April 2022, effective 1 November 2022), within the broader baseline of GB/T 22239-2019 for cybersecurity graded protection and GB/T 35273-2020 for personal information security. International counterparts include the Trusted Computing Group’s TPM 2.0 family and the Confidential Computing Consortium’s threat model for execution isolation. The layer provides no business semantics; it answers only “was this computation genuine?”

L2 — Credential layer. Comprises decentralised identifiers, verifiable credentials, enterprise public-key infrastructure, reliable electronic signatures, and trusted timestamping. Its responsibility is P1 and the authority component of P5: binding actions to legal persons and authorised signatories, and expressing attributes—creditworthiness, tax status, business registration, delivery authority—in a form that can be verified without contacting the issuer on every occasion. The W3C Verifiable Credentials Data Model 2.0 became a W3C Recommendation on 15 May 2025, standardising the issuer-holder-verifier model, tamper-evidence, and selective-disclosure-friendly exchange mechanisms; decentralised identifier specifications decouple the identifier from any single resolution authority. In the European Union, Regulation (EU) 2024/1183 of 11 April 2024 amended Regulation (EU) No 910/2014 to establish the European Digital Identity Framework, introducing the European Digital Identity Wallet and electronic attestations of attributes with explicit support for selective disclosure.

China's equivalent legal effect flows from Articles 13 and 14 of the *Electronic Signature Law*.

L3 — Confidential-computing layer. Comprises secure multi-party computation, federated learning, homomorphic encryption, differential privacy, and zero-knowledge proof systems. Its responsibility is P4 and the input-privacy component of P2: enabling joint computation—set intersection to detect duplicate financing, risk scoring to support credit decisions, aggregate monitoring for supervisors—without any participant obtaining another's raw data.

L4 — Shared state layer. Comprises a permissioned distributed ledger, an append-only log, and smart contracts implementing a defined state machine. Its responsibility is P3 and the ordering component of P2: providing a tamper-evident, chronologically ordered, replicated record of commitments and state transitions, and executing deterministic rules without discretion.

L5 — Execution and settlement layer. Comprises payment initiation, bank settlement interfaces, the unified movable-asset registration system, and—where applicable—tokenised deposits or wholesale central bank digital currency arrangements. Its responsibility is P5 and finality: converting a verified obligation into discharged value through legally recognised channels. The design space here has widened with

the arrival of wholesale settlement experiments and tokenised-deposit arrangements, and the analytical literature has begun to clarify what is genuinely new about them: the economic questions concern settlement finality, liquidity, and the redesign of monetary arrangements, not the novelty of the ledger [8], [6], [3]. Under Yinfa [2025] No. 77, settlement of AREC financing must be conducted through institutions with the requisite qualifications, typically commercial banks; information service institutions may not use their own accounts as settlement accounts and may not pool or misappropriate funds. This constraint is not merely regulatory decoration: it is what prevents a platform from becoming an unlicensed credit intermediary.

The critical structural feature is that L4 never holds raw data and L3 never holds authoritative state. Ledgers hold hashes, commitments, state identifiers, proof handles, and minimal audit anchors. Raw contracts, invoices, delivery records, and personal data remain with their controller or inside an attested enclave. This arrangement is what allows the system to satisfy simultaneously the transparency demanded by financiers and regulators and the confidentiality demanded by commercial parties and data-protection law.

Figure 1 summarises the layering and, in particular, the responsibility that each layer must refuse.




Fig. 1. TCCE reference architecture: one assurance responsibility per layer.

C. Lifecycle of an AREC under TCCE

We walk a single instrument through the architecture, since the value of the design is visible only in the composition.

Stage 0 — Onboarding. The anchor buyer and each supplier obtain verifiable credentials binding their legal identity, business registration, authorised signatories, and—where

relevant—tax and customs identifiers. Credentials are issued by authorities and counterparties with standing and are revocable. Verification of a credential does not require contacting its issuer, but revocation status is checked against a status list at each material step.

Stage 1 — Trade event capture. Purchase order, delivery note, acceptance record, and invoice are created in the parties'

own systems. Each is signed by the responsible party and timestamped, and each carries a reference to a shared business identifier: the order. Hashes of these documents, together with the business identifier, are committed to L4. The documents themselves are not.

Stage 2 — Undertaking. The anchor buyer issues the AREC: amount, tenor, payment date, payer, payee, and reference to the underlying business identifier. The undertaking is signed with a reliable electronic signature. The AREC's state is recorded in L4 as ISSUED, and a unique receivable identifier is derived deterministically from the business identifier plus the payer-payee pair so that the same economic receivable produces the same identifier in every system that adopts the convention.

Stage 3 — Transfer and splitting. If the payee transfers or splits the AREC, each split is recorded as a state transition referencing the parent identifier, so that the full split tree remains reconstructible. Constraints on the number of layers and the number of transactions—required by Yinfa [2025] No. 77—are enforced by the smart contract, and deviations are flagged as anomalies rather than silently permitted.

Stage 4 — Financing request and duplicate-financing check. A supplier requests discounting from a bank. Before any credit decision, the bank participates in a private set-intersection protocol (L3) over receivable identifiers with other participating banks and platforms. The protocol reveals whether the identifier appears elsewhere; it reveals nothing else—in particular, it does not reveal the bank's own client list or the amounts involved. Where the counterparty set is small and a trusted execution environment is acceptable to all parties, the same computation can be executed inside an attested enclave (L1 + L3) at lower latency, with attestation evidence anchored to L4 so that the computation's provenance is later auditable.

Stage 5 — Trade-background verification. Rather than receiving the supplier's full invoice and contract file, the bank receives a set of assertions with proofs: that an invoice matching the undertaking amount exists and is not cancelled;

that the invoice counterparties match the undertaking parties; that a delivery record signed by a carrier exists for the corresponding period; that the tenor does not exceed the regulatory cap. Where the underlying source is itself authoritative and machine-readable—tax-authority invoice data, customs declarations, warehouse receipts—the assertion is verified by a credential issued by that authority rather than by data transfer.

Stage 6 — Registration and notification. Upon approval, the bank registers its interest in the unified movable-asset financing registration system (L5), and the registration reference is committed to L4. The receivable's state becomes FINANCED. Notice to the debtor is given in the form required by Article 764 of the *Civil Code*, and the notice event is recorded. At this point the priority position established by Article 768 of the *Civil Code* is secured.

Stage 7 — Settlement and discharge. Funds move through the bank, not through the platform. At maturity the anchor buyer pays into a designated account; the state becomes DISCHARGED, and any residual claim is released. If payment is not made, the state becomes DEFAULTED, and the evidence package assembled across L1 to L5 is available for enforcement.

Stage 8 — Evidence assembly. At any point, and especially on default, the system can assemble a package containing: the signed trade documents; signature certificates and timestamps; the chain of state transitions; the attestation records and proof objects for computations performed; registration references; and the configuration of the systems involved. This is the artefact that satisfies the Supreme People's Court's requirement, in Article 17 of the *Rules on Online Litigation*, to examine the reliability of the depositary platform and the authenticity of the data *prior* to its being recorded on the chain. The architecture's answer to the anterior-truth problem is not to assert it but to make the anterior evidence part of the package.

Figure 2 traces a single instrument through these stages and identifies the layer responsible at each point.

Lifecycle of an accounts-receivable electronic certificate under TCCE




Fig. 2. Lifecycle of an accounts-receivable electronic certificate under TCCE, with regulatory control points.

D. Why convergence rather than substitution

Each layer addresses a distinct failure mode, and the layers are complements, not alternatives. A ledger without a credential layer cannot bind records to legal persons. A credential layer without a state layer cannot establish singularity. Confidential computing without a trust anchor produces results that cannot be audited, and a trust anchor without confidential computing protects computation but not data confidentiality. Execution without the other layers is an ordinary payment.

This complementarity has a testable implication that distinguishes our account from techno-optimistic accounts: the marginal value of adding a layer should be conditional on the presence of the others. Adding a blockchain to a process with no credential infrastructure should produce little measurable reduction in fraud; adding privacy-preserving computation to a process with no registration mechanism should produce little reduction in duplicate financing. Section 6 formalises this conditional structure.

V. CORE MECHANISMS AND THEIR COMPOSITION

A. Distributed ledgers and smart contracts: registry, ordering, and constrained discretion

The contribution of a permissioned ledger to inter-organisational collaboration is frequently overstated and then, in reaction, understated. It is worth being precise about what it delivers.

What a replicated, append-only, cryptographically linked log delivers is *ordering without a trusted orderer* and *posterior tamper-evidence*. In a multi-party setting where no participant will accept another's database as authoritative, this is genuinely valuable: it removes a class of disputes ("the record says you confirmed on the 3rd") without requiring anyone to trust anyone. The formal literature has modelled the associated trade-offs carefully, showing that distributed settlement changes the economics of verification and commitment and can reduce the cost of certain forms of contracting, while simultaneously introducing new incentive problems among validators [10], [13], [5]. Corporate-governance scholarship has argued that the effects of distributed ledgers on disclosure and intermediation are likely to be substantive but are mediated by existing governance structures rather than replacing them [57]. Operations-management scholarship has made the parallel point that distributed ledgers should be understood as governance mechanisms with a defined boundary—effective where the transaction can be made explicit, less effective where it cannot [4], [31].

What a ledger does not deliver is anterior truth, confidentiality, or legal effect. It is therefore a mistake to put business documents on a chain. The correct division, implemented in L4, is: the chain holds *identifiers, states, hashes, and proof handles*.

Smart contracts add value in a narrower band than is often claimed. Their legitimate role is to remove discretion where

discretion is costly: enforcing that a split does not exceed the outstanding amount, that a transfer respects the layering constraint, that financing cannot occur while a DISPUTED flag is set, that a discharge releases all descendants in the split tree. This is a finite state machine derived from a legal contract, not a replacement for it—a distinction the legal literature on smart contracting has insisted upon [54]. Two engineering consequences follow. First, contracts must be minimal and formally specified; complexity is the enemy of auditability. Second, every contract must include an administrative path for correction and freeze, because the real world generates erroneous states, and a system that cannot correct an error will be abandoned by practitioners.

A recurring design error deserves explicit mention: treating the ledger as a *settlement* system. In most jurisdictions, including China under Yinfa [2025] No. 77, settlement finality is a regulated activity performed by licensed institutions. A ledger can evidence that an obligation should be discharged; it cannot, by itself, discharge it.

B. Privacy-preserving computation: choosing among instruments

Five families of technique are available, and they are not interchangeable. The selection should be driven by three questions: what is the *joint* function to be computed; who must *not* learn what; and what is the *acceptable latency and trust assumption*.

Secure multi-party computation (MPC). Multiple parties compute a function of their private inputs without revealing those inputs. Protocols such as secret-sharing-based arithmetic circuits and their adversarially secure variants, and garbled-circuit or hybrid constructions, offer different positions on the trade-off between rounds of communication, preprocessing, and security against malicious versus semi-honest adversaries. The systems literature has demonstrated that privacy-preserving machine learning over secure computation is practical at meaningful scale [33], and secure aggregation protocols have made federated model updates private against an honest-but-curious server [7]. MPC is the right instrument for *low-frequency, high-sensitivity, well-defined* joint computations. Duplicate-financing detection by private set intersection is a canonical example: the function is a set intersection, the output is small, and the input—a list of receivable identifiers—is commercially sensitive.

Federated learning (FL). Parties collaboratively train a model while keeping training data local, exchanging model updates rather than records. The classical taxonomy distinguishes horizontal, vertical, and federated transfer learning according to how the sample and feature spaces overlap [56]; the comprehensive synthesis of open problems identifies non-identically-distributed data, communication cost, robustness to poisoning, and fairness as the principal obstacles [28]. In supply-chain finance the vertical case dominates: a bank holds repayment behaviour, a platform holds transaction frequency, a core enterprise holds order and acceptance data, and a logistics provider holds delivery

reliability—different features over overlapping firms. Federated learning is the right instrument when the output is a *model* and the evaluation is statistical, not when a specific transaction requires a specific yes-or-no answer.

Differential privacy (DP). A computation's output is randomised so that the presence or absence of any single record is bounded in its effect, with a cumulative privacy budget tracked across queries [16]. DP suits *statistical release*—supervisory aggregates, sector-level exposure reporting—rather than transaction-level decisions, and it composes naturally with the other instruments: a federated model can be trained under secure aggregation and its aggregate outputs released under a privacy budget.

Homomorphic encryption (HE). Computation proceeds directly on ciphertext. Partially homomorphic schemes such as Paillier support addition and are efficient enough for aggregation; levelled and fully homomorphic schemes support richer circuits at substantially greater cost [1]. HE suits *narrow, repeated, arithmetic* tasks—encrypted aggregation of exposure across institutions, encrypted scoring with a linear or low-degree model—and suits settings where one party holds the encrypted data and another holds the key.

Zero-knowledge proofs (ZKP). A prover convinces a verifier that a statement is true without revealing why. Proof systems differ in proof size, prover time, verification cost, and setup assumptions, with pairing-based succinct systems, universal-updateable systems, and transparent hash-based systems occupying distinct positions. ZKP is the right instrument when the *assertion* is simple, the *witness* is

sensitive, and verification must be public or cheap. In our architecture ZKP carries assertions such as: “the financed amount does not exceed the outstanding balance of this identifier,” “the tenor of this undertaking is within the permitted maximum,” “this entity appears on the regulator’s qualified-institution list without disclosing which entry matched.” The European framework’s emphasis on selective disclosure in electronic attestations of attributes is the regulatory counterpart of the same idea (Regulation (EU) 2024/1183).

Trusted execution environments (TEE). An isolated processing environment that provides confidentiality and integrity for code and data loaded inside it, protected from the surrounding operating system and, to a defined degree, from physical observation [39]. Hardware-isolated execution (Intel SGX and TDX, ARM TrustZone and Confidential Compute Architecture, AMD SEV-SNP) allows complex computation on plaintext inside an enclave at near-native speed. The threat model is well characterised and the honest characterisation is that an enclave protects against software-privileged and, in varying degrees, physical adversaries, but *not* against microarchitectural side channels unless specifically mitigated [15]. Transient-execution attacks have repeatedly broken enclave confidentiality in laboratory conditions [52]. The correct conclusion is not to reject enclaves but to treat them as one control among several, to require remote attestation, to keep sensitive functions small, and to pair them with cryptographic mechanisms where the trust assumption cannot be accepted.

Table 1 summarises the selection logic.

Table 1. Selecting a privacy-preserving instrument

Instrument	Best suited to	Principal cost	Trust assumption	Role in TCCE
MPC	Low-frequency joint functions over sensitive inputs (e.g. private set intersection for duplicate detection)	Communication rounds; malicious-security overhead	Honest majority or dishonest majority with cryptographic setup	Cross-institution duplicate-financing check
Federated learning	Joint model training where parties hold different features or samples	Communication; non-IID degradation; poisoning risk	Honest-but-curious aggregator; secure aggregation for update privacy	Multi-tier credit scoring across bank–platform–anchor data
Homomorphic encryption	Repeated narrow arithmetic over encrypted values	Ciphertext expansion; multiplicative depth	Key holder does not collude with compute host	Encrypted exposure aggregation; encrypted linear scoring
Zero-knowledge proofs	Proving simple assertions about hidden witnesses with cheap verification	Prover time; circuit design; setup for some systems	Soundness of the proof system; correct circuit	Tenor, balance, and eligibility assertions
TEE	Complex, low-latency computation on plaintext	Hardware dependence; side-channel exposure	Hardware, firmware, microcode, and attestation infrastructure	Attested risk engine; model execution inside L1

C. Trusted computing: making computation auditable

If computation moves into places its subjects cannot observe, auditability must be re-established by other means. Two instruments do this.

Remote attestation. An attesting platform produces a signed statement—rooted in a hardware key provisioned at manufacture—identifying the measured boot chain, the firmware and microcode versions, and the hash of the executing code. A relying party verifies the signature chain and compares measurements against known-good values before releasing sealed data. In our architecture the attestation is not merely a gate; the attestation record is committed to L4 so that, months later, an auditor can establish that a decision was produced by the code the participants had agreed to run.

This is the mechanism that converts “the platform says” into “the platform can prove.”

Chain of custody for evidence. Attestation covers code, not inputs. Input provenance requires the credential layer: every input to a computation is accompanied by a credential identifying its source and a signature over its content, so that the enclave can itself verify what it consumed. The combination—attested code plus credentialed inputs—is what allows a computation performed in a place nobody can see to yield a result everyone can check.

Chinese standardisation provides the domestic baseline for the root of trust and graded-protection requirements (GB/T 29829-2022; GB/T 22239-2019), while data-lifecycle capability assessment (GB/T 37988-2019) and personal-

information security requirements (GB/T 35273-2020) bound what may be processed at all. A pragmatic deployment note: hardware-rooted attestation is most valuable where the alternative is trusting a software operator, and least valuable where the participant set is small enough for contractual and audit controls to dominate.

D. Digital credentials and electronic transferable records

The bridge from cryptographic state to legal fact runs through two bodies of law and one body of standards.

Signatures and evidence. As noted, China's *Electronic Signature Law* conditions reliable-signature status on exclusivity, sole control, and post-hoc detectability of alteration to both signature and content, and accords such signatures the effect of a handwritten signature or seal. The Supreme People's Court's *Rules on Online Litigation* accept that chain-stored data verified as technically consistent has not been altered after recording, subject to rebuttal, while requiring examination of platform reliability and pre-chain authenticity. The practical implication for system design is direct: a platform that stores only hashes without retaining the original documents, the signature certificates, the timestamps, and the generation logs will fail evidentiary scrutiny at precisely the moment it matters.

Transferable records. Possession performs a distinctive function in commercial law: it is how a holder demonstrates entitlement to a negotiable document. Any electronic substitute must therefore supply a functional equivalent. The UNCITRAL Model Law on Electronic Transferable Records (2017) provides the dominant template, replacing possession with *control* and requiring that the method used establish singularity and exclusive control, preserve integrity, and permit the identification of the person in control and the transfer of control to another. The United Kingdom's Electronic Trade Documents Act 2023 (c. 14), which received Royal Assent on 28 June 2023 and came into force on 20 September 2023, implements the same logic by giving electronic versions of specified trade documents the same legal treatment as their paper counterparts. These instruments matter for our setting in a specific way: they establish that the legal question is *control*, not *encryption*, and that a system can be technically excellent and legally insufficient if it cannot demonstrate who has exclusive control at any instant.

Standards for expression. The W3C Verifiable Credentials Data Model 2.0, a Recommendation since 15 May 2025, standardises the issuer-holder-verifier model and the representation of tamper-evident claims; decentralised identifiers decouple identity from any single registry. The value for practitioners is not novelty but interoperability: a credential issued by a tax authority, a customs administration, or a core enterprise can be presented to any verifier without bilateral integration, and revocation can be checked against a status list rather than by calling the issuer.

A caveat is necessary. Neither a verifiable credential nor a model-law-compliant control mechanism establishes that the underlying fact is true; both establish that a stated issuer made a stated claim, and that the claim has not been altered. The

distinction—between the authenticity of an utterance and the truth of its content—is the single most important thing for system architects to internalise, and it mirrors exactly the anterior-truth limitation identified in Section 2.3.

E. Composition patterns

The architecture's practical value lies in five recurring composition patterns, each of which assigns one assertion to one mechanism.

Pattern A — Private duplicate check. *Assertion:* this receivable identifier is not already financed. *Mechanism:* private set intersection (MPC) or attested enclave comparison (TEE + L1), with the resulting negative or positive flag—and, for audit, a commitment to the query—anchored on the ledger. *Why not alternatives:* a central registry would be simpler but requires participants to disclose their books; bilateral disclosure would breach commercial confidentiality and possibly data-protection law.

Pattern B — Attested scoring. *Assertion:* this financing request passes the agreed risk model. *Mechanism:* the model executes inside an attested enclave; inputs arrive credentialed and encrypted; the enclave emits a score, a reason code set, the model version, and an attestation; all four are anchored on the ledger. *Why:* the borrower can contest a decision without seeing the model, and the supervisor can audit without seeing the data.

Pattern C — Credentialed attribute verification. *Assertion:* this invoice exists, is not cancelled, and matches the undertaking. *Mechanism:* the tax authority or an authorised service issues a verifiable credential over defined attributes; the supplier presents it; the bank verifies signature, issuer, and revocation status. *Why:* it replaces document transfer with attribute transfer, which is precisely the data-minimisation required by the principle of necessity.

Pattern D — Programmable covenant. *Assertion:* splitting has not exceeded the permitted layers and the aggregate of splits does not exceed the parent. *Mechanism:* smart contract enforcement at transfer time, with anomaly flagging rather than silent rejection. *Why:* discretion at this point is exactly what generates systemic risk, and the regulator has explicitly required control over layers and frequency.

Pattern E — Evidence package on default. *Assertion:* the facts supporting enforcement. *Mechanism:* assembly of signed trade documents, timestamps, state transitions, attestation records, proof objects, and registration references into a single exportable bundle conforming to evidentiary expectations. *Why:* the value of every other mechanism is realised only if it converts into an enforceable claim.

The patterns share a discipline: never ask a mechanism to prove more than it can, and never let a layer hold data it does not need to discharge its own responsibility.

VI. ANALYTICAL MODEL: SINGULARITY, DISCLOSURE, AND THE DEPTH OF CREDIT PASS-THROUGH

This section formalises three claims made informally above: that duplicate financing is a pricing problem as much as

a fraud problem; that detection and legal priority are complements; and that verifiable selective disclosure can dominate both concealment and full disclosure. The models are deliberately simple; their purpose is to make the conditional structure of the architecture's value explicit and falsifiable.

A. Model 1 — Duplicate financing as a detection-and-priority problem

Setup. A supplier holds a single genuine receivable of face value F against an anchor buyer that pays in full with probability one at maturity. There are $n \geq 1$ potential financiers. Financier i contacted by the supplier advances λF , $0 < \lambda < 1$, at simple rate r for term T . The supplier attempts to obtain financing from all n financiers; financier j agrees only if it fails to detect that the receivable is already financed. Let $\delta \in [0, 1]$ denote the *detection capability* of the system: the probability that a financier discovers a prior financing of the same receivable before advancing funds. Let $\rho \in [0, 1]$ denote the *priority effectiveness* of the legal-institutional layer: the share of a contested receivable that a financier recovers by virtue of registration or notice, as opposed to pro-rata sharing.

Assume, matching Article 768 of China's *Civil Code*, that if no financier has established priority, competing claims share pro rata; and if priority has been established, the prior-ranking financier recovers first up to its claim.

Expected number of advances. Financier 1 always advances. Each subsequent financier advances with probability $1 - \delta$. Hence

$$E[m] = 1 + (n-1)(1-\delta),$$

with $E[m] \rightarrow 1$ as $\delta \rightarrow 1$ and $E[m] = n$ when $\delta = 0$.

Recovery. Total claims equal $m\lambda F$. If $m\lambda F \leq F$, all claims are paid in full. Otherwise the receivable is short by $m\lambda F - F$. Under pro-rata sharing ($\rho = 0$) each financier recovers F/m ; under perfect priority ($\rho = 1$) the first-ranking financier recovers λF and the remainder, if any, cascades. For a financier that advances without knowing its rank, expected recovery per unit advanced is therefore

$$R(\delta, \rho) = \min \left(\lambda F, \rho \lambda F + (1-\rho) \frac{F}{E[m]} \right).$$

Zero-profit pricing. Competition sets r such that the advance plus interest equals expected recovery:

$$\lambda F(1+rT) = R(\delta, \rho) \Rightarrow r(\delta, \rho) = \frac{1}{T} \left(\frac{R(\delta, \rho)}{\lambda F} - 1 \right)$$

Note $r \leq 0$ is impossible in equilibrium; the meaningful range is $r \geq 0$, i.e. the financier participates only when $R \geq \lambda F$, which requires $E[m] \leq 1/\lambda$.

Proposition 1 (financing cost falls with detection, and the return to detection rises with fragmentation). $r(\delta, \rho)$ is non-increasing in δ , and $|\partial r / \partial \delta|$ is increasing in n .

Intuition. Detection reduces the expected number of competing claims, which raises expected recovery and lowers the rate required to break even. Because $\partial E[m] / \partial \delta = -(n-1)$, a given improvement in detection removes more duplicate claims when there are more potential financiers. The practical reading is important: the same technology produces larger measured benefits in a fragmented, multi-platform market than in a concentrated one, so cross-jurisdictional evaluations of “does blockchain reduce SCF fraud” are not comparable without controlling for market structure.

Proposition 2 (detection and legal priority are complements, not substitutes). $R(\delta, \rho)$ has $\partial^2 R / \partial \delta \partial \rho \leq 0$, and the cost-reducing effect of an increase in ρ is larger when δ is low, and vice versa.

Intuition. Detection prevents multiple advances; priority determines the outcome *conditional* on multiple advances having occurred. When detection is poor, the residual protection supplied by registration is what limits loss, so improving registration is valuable; when detection is good, registration matters less for loss but still matters for certainty. Neither instrument alone drives expected loss to zero, because detection is probabilistic and recovery without priority is dilutive. This is the formal content of the claim, advanced in Section 4.4, that the architecture's layers are complements: a deployment that implements a ledger ($\delta \uparrow$) but omits registration integration ($\rho = 0$) captures only part of the available benefit, and vice versa.

Proposition 3 (adoption externality). The private return to an individual financier from joining a shared detection arrangement is lower than the social return whenever the financier's own detection improvement prevents advances that would otherwise have been made by others.

Intuition. Detection is a public good among financiers: if bank i improves its checks, the supplier simply shifts to bank j , and i bears the cost while j avoids the loss. Conversely, if all join, all gain. This is a familiar collective-action structure, and it explains an empirical regularity that naive technology-centred accounts cannot: the technology was available and understood, yet adoption stalled. It also explains why a *mandated* control point—such as the registration requirement and the information-service-institution filing and reporting regime introduced by Yinfa [2025] No. 77, under which 237 institutions had filed information and reported business data by the first quarter of 2026—can move a market that voluntary coordination could not.

B. Model 2 — Concealment, disclosure, and verifiable selective disclosure

Setup. A supplier's default probability is θ , drawn from a prior with mean $\hat{\theta}$ and variance σ_0^2 . A financier prices at $r = r_0 + \gamma \sigma^2$, where σ^2 is the financier's posterior variance after observing evidence and $\gamma > 0$ captures risk pricing. Disclosing information about the supplier reduces posterior

variance but also reveals proprietary attributes—unit prices, customer identities, input costs—that cause competitive harm κ per unit of proprietary information leaked, with $\kappa > 0$.

Let evidence e be characterised by two scalars: the reduction in posterior variance it achieves, $s(e) \in [0, \sigma_0^2]$, and the quantity of proprietary information it leaks, $l(e) \geq 0$. The supplier's total cost of financing is

$$C(e) = (r_0 + \gamma(\sigma_0^2 - s(e)))\lambda FT + \kappa l(e).$$

Three evidence regimes are available.

Concealment: $s = 0, l = 0$. Cost $C_C = (r_0 + \gamma\sigma_0^2)\lambda FT$.

Full disclosure: $s = \sigma_0^2$ (the financier learns θ essentially exactly), $l = L > 0$. Cost $C_F = r_0\lambda FT + \kappa L$.

Verifiable selective disclosure: a zero-knowledge proof or a credentialed attribute reveals the truth of the specific assertion the financier needs—for instance, “the ratio of on-time deliveries exceeds the threshold associated with low default risk”—so $s = \sigma_0^2$ as well, while $l = l_Z \ll L$, because the verifier learns the truth value of a statement and not the underlying records.

Proposition 4 (dominance of verifiable selective disclosure). For any $\kappa > 0$ and any full-disclosure regime with leakage L , there exists a threshold $\hat{l}_Z$ such that if verifiable selective disclosure achieves the same variance reduction with $l_Z < \hat{l}_Z$ and verification cost per financing below $\kappa(L - l_Z)$, then $C_Z < \min\{C_C, C_F\}$; and the inequality is strict in the financier's rate relative to concealment whenever $\gamma\sigma_0^2 > 0$.

Intuition. Concealment is costly because uncertainty is priced; full disclosure is costly because proprietary information is leaked. Verifiable disclosure separates the *epistemic* function of disclosure (reducing the financier's uncertainty) from its *informational* side effect (revealing raw records), and in doing so removes the trade-off rather than balancing it. This is the economic content of data minimisation: minimisation is not a constraint reluctantly accepted for compliance reasons, but a mechanism that strictly lowers the cost of capital when proofs are available.

The same analysis reframes the “unwilling to share” problem identified in Section 2.5. Data is non-rival: its use by one financier does not diminish its use by another, so restricting access imposes a social cost that is invisible in any bilateral negotiation [27]. The reason firms restrict access anyway is that disclosure reveals proprietary attributes, which is exactly the κ term above. It follows that the policy instrument that unlocks data sharing is not a mandate to share and not a property right in raw records, but a reduction in the marginal leakage per unit of variance reduction—that is, better proofs and better credentials.

Corollary 4.1 (why privacy-preserving computation is not merely compliance). The value generated by L3 is $\gamma s(e)\lambda FT + \kappa(L - l_Z)$, which is increasing in both the risk-

pricing coefficient γ and the competitive sensitivity κ . Firms in industries with high proprietary sensitivity and heterogeneous credit quality—precisely the profile of many multi-tier manufacturing supply chains—capture the largest benefit.

C. Model 3 — The endogenous depth of credit pass-through

Setup. A chain has tiers $t = 1, \dots, T$, with tier 0 the anchor buyer. To extend credit to tier t on the strength of the anchor's credit, a financier must verify t linkage events (each tier's trade with the next). Verifying one linkage costs v , and each verification is subject to an independent error probability ε —the probability that a linkage is accepted when it does not, in fact, exist. The benefit of extending anchor-rate financing rather than supplier-rate financing to tier t is $B_t = (r_t - r_0)\lambda F_t T$, where $r_t - r_0$ grows with t .

Naive condition. Pass-through reaches tier t if $B_t \geq t v$. Reducing v through standardisation and credentials therefore extends pass-through, which is the usual argument for interoperability.

Proposition 5 (there is an endogenous limit to pass-through even when verification is free). The expected loss from extending credit to tier t is increasing in t through the compounding of verification error: the probability that at least one of the t linkages in the chain is spurious is $1 - (1 - \varepsilon)^t$, which is increasing and concave in t . Consequently, even as $v \rightarrow 0$, the financier's participation constraint binds at a finite t^i , and t^i is increasing in the accuracy of linkage verification but at a decreasing rate.

Intuition. Depth is limited not by the cost of asking but by the reliability of the answers. Each additional tier adds a link that can be fabricated, and the probability of at least one fabrication compounds. This result tempers a common claim in practitioner literature—that distributed ledgers allow anchor credit to reach arbitrarily deep tiers—and it identifies the correct research target: interventions that reduce ε (third-party corroboration, physical sampling, credentialed logistics data from authoritative sources) extend pass-through; interventions that merely reduce v do so only up to the accuracy ceiling. It also supplies a design directive: the architecture should treat verification *accuracy*, not *throughput*, as its primary performance metric, and should allocate marginal investment towards authoritative data sources near the deep tiers.

D. What the models predict and how they could be falsified

Taken together, the models yield four testable predictions.

1. *Heterogeneous treatment effects by market structure.* Detecting-duplication technology should show larger measured effects in fragmented, multi-platform markets than in concentrated ones (Prop. 1).

2. *Complementarity in adoption.* Platforms that implement a ledger but not registration integration should show smaller reductions in duplicate financing

than those implementing both, and the interaction term should be negative in a difference-in-differences specification with platform and time fixed effects (Prop. 2).

3. *Rate effects concentrated in sensitive, heterogeneous industries.* Reductions in financing cost from privacy-preserving verification should be largest where proprietary sensitivity and credit-quality dispersion are both high (Cor. 4.1).

4. *Bounded pass-through.* Credit pass-through depth should increase with interoperability investments but plateau, and the plateau should be higher where third-party corroboration of linkage events is available (Prop. 5).

These predictions are stated so that they can be tested with platform-level data when it becomes available. We return to the data problem in Section 10.

VII. IMPLEMENTATION BLUEPRINT AND FIELD EVIDENCE

A. From regulatory control points to technical controls

A design is practical only if it maps onto what a regulated institution is actually required to do. Table 2 performs that mapping for the Chinese AREC regime established by Yinfa [2020] No. 226, Yinfa [2025] No. 77, State Council Decree No. 802, and the *Civil Code*. The mapping is the substantive contribution of this section: each row converts a legal obligation into a mechanism, a layer, and an auditable artefact.

Table 2. Mapping regulatory control points to technical controls

Requirement (source)	Technical control	Layer	Auditable artefact
Genuine trade relationship (Yinfa [2025] No. 77)	Signed, timestamped chain of order → delivery → acceptance → invoice, cross-referenced by a shared business identifier; assertions verified against authoritative sources	L2, L4	Event chain with signature certificates and timestamps
Payment term ≤ 6 months in principle, ≤ 12 months absolutely (Yinfa [2025] No. 77)	Contract-level constraint at issuance; exception flagging requiring documented justification and enhanced review	L4, L5	Issuance record with tenor assertion and exception log
Control over transfer layers and number of splits; investigation of anomalous splitting (Yinfa [2025] No. 77)	Split-tree state machine with layer and count constraints; anomaly scoring on layering velocity, fragmentation, and concentration	L4	Split tree; anomaly flag with features and thresholds
Registration of AREC financing in the unified movable-asset registration system (Yinfa [2025] No. 77; Civil Code Art. 768)	Automated registration trigger on financing approval; registration reference committed to the ledger	L5, L4	Registration reference and its commitment
Settlement through qualified institutions; no pooling by information service institutions (Yinfa [2025] No. 77)	Payment instruction origination restricted to licensed institutions; platform excluded from fund flows	L5	Settlement instruction and bank confirmation
Notification to the debtor (Civil Code Arts. 546, 764)	Notice generation with identity evidence and delivery confirmation	L2, L5	Notice record and delivery evidence
No closed-loop circulation; no captive financing (Yinfa [2020] No. 226)	Open financing interface; prohibition of exclusivity enforced by contract and monitored by transaction-level concentration metrics	L4, L5	Distribution of financiers per anchor and per supplier
Legality, legitimacy, necessity in data processing (PIPL; Yinfa [2025] No. 77)	Purpose-bound field-level authorisation; attribute-based disclosure; proofs in place of records	L3, L2	Purpose register, authorisation records, proof objects
Bank retains credit decision and whole-debt monitoring (Yinfa [2025] No. 77)	Bank-side decision service consuming attested scores and its own exposure data; cross-product debt aggregation	L1, L3	Decision log with model version and attestation
Information security assessment of partner institutions (Yinfa [2025] No. 77)	Attestation-based assurance of partner compute environments; periodic re-attestation	L1	Attestation records and their commitments

The table is deliberately exhaustive about the *control points* and silent about products, because the same control point can be satisfied by different technologies. What matters is that each obligation has exactly one owner and one artefact.

B. Reference implementation

We describe a reference implementation as a sequence of services, independent of vendor choice.

Identity and authority service. Issues and verifies credentials for legal entities and authorised signatories, maintains revocation status lists, and exposes verification to all participants. Enterprise identity is anchored in official registration data; signatory authority is anchored in corporate authorisation documents and, where available, in bank-side mandate records.

Business-identifier service. Assigns and resolves the unique business identifier from which the receivable identifier is derived. The design choice that matters here is determinism: given the same order, payer, and payee, any implementation must derive the same identifier. Without determinism, cross-platform duplicate detection is impossible, because the same

economic receivable appears under two different identifiers and set intersection returns a false negative.

Event capture service. Ingests signed event records from participant systems, verifies signatures and timestamps, checks reference integrity, and commits hashes to the ledger. It rejects—rather than silently accepts—events whose references do not resolve, since a dangling reference is exactly the defect that produces unverifiable trade backgrounds.

State machine service. Maintains the AREC state graph including the split tree, enforces constraints, and emits events. States: DRAFT, ISSUED, CONFIRMED, SPLIT, TRANSFERRED, PLEDGED, FINANCED, DISPUTED, DISCHARGED, DEFAULTED. Transitions are constrained, and every transition carries the actor's credential reference.

Confidential computation service. Runs private set intersection for duplicate checks, federated or enclave-based scoring, and proof generation. Exposes only outputs and attestations. Scheduled in batch where latency permits—duplicate-financing set intersection over identifier commitments need not be synchronous with the credit decision; a pre-clearance run at onboarding plus a

confirmation run before disbursement captures most of the benefit at a fraction of the interactive cost.

Evidence service. Assembles exportable evidence packages on demand, including pre-chain source documents. This service is what makes the preceding ones legally meaningful, and it is the component most often omitted in practice.

Supervisory interface. Publishes aggregate statistics—outstanding by tenor bucket, splitting depth distribution, concentration by anchor, registration coverage, anomaly counts—without exposing transaction-level commercial data. Aggregation over a privacy budget, or computation inside an attested enclave with only aggregate release, allows supervision without bulk data collection.

C. Field evidence and its limits

The Chinese market provides several datapoints that inform, without proving, the architecture's claims.

Scale and price. As of the first quarter of 2026, ARECs outstanding stood at RMB 2.75 trillion against a financing balance of RMB 2.08 trillion, with an average financing rate of 2.71 per cent across 237 reporting institutions. The financing-rate level is notable: it is materially below typical unsecured SME lending rates, which is consistent with the claim that anchor credit is being passed through to at least some suppliers—but it is equally consistent with selection, since the suppliers that obtain AREC financing are those the anchor buyer is willing to confirm. The aggregate figure therefore cannot be used to measure pass-through depth.

Infrastructure maturity. The Shanghai Commercial Paper Exchange's supply-chain bill platform demonstrates that a unified registry with equal-denomination issuance is technically and institutionally feasible; cumulative transaction scale reached RMB 67.163 billion with over 3,000 participating enterprises by the end of 2021. Individual platforms report substantial reach: Simplehui reports having served over 100,000 micro and small enterprises with cumulative supply-chain bill issuance above RMB 66 billion and cumulative discounting above RMB 47 billion as of end-2024; Linklogis reports RMB 411.2 billion of supply-chain assets processed through its technology solutions in 2024, up 28 per cent, serving 2,533 core-enterprise and financial-institution clients and cumulatively over 330,000 micro and small enterprises; China Corporate United Cloud Chain reports 557,000 platform users at end-2024, comprising 6,601 core enterprises, 546,400 chain-affiliated enterprises, and 3,574 financial institutions.

These figures must be read with three caveats, and the caveats are as important as the numbers. First, they are self-disclosed and measured on differing bases—cumulative flow, period processing volume, and registered users—and cannot be added together or converted into market shares. Second, no regulator has published a verified statistic on the prevalence of duplicate financing, so the *baseline* against which any intervention would be measured is unknown. Third, and most importantly for research design, no institution has published verified statistics establishing that privacy-preserving computation has been deployed at production scale in this

market; the available evidence is firm-level disclosure and technical demonstration. We therefore do not claim empirical validation of the architecture. What we claim is that the architecture's control points correspond one-to-one to obligations that regulators have now made binding, and that the gap between stated practice and verified practice is itself a finding: it identifies measurement, not invention, as the field's next task.

D. Engineering realities

Three practical constraints shape deployment more than any architectural diagram.

Latency placement. Not every check must be synchronous. Splitting and tenor constraints are state-machine checks and belong on the critical path. Duplicate-financing set intersection can be pre-computed in batches. Model scoring must be synchronous but is fast relative to human credit review. Attestation verification adds a round trip and should be cached with a defined validity window, since attestation freshness requirements ought to be risk-based rather than uniform.

Key management is the actual attack surface. In practice, losses from inter-organisational systems arise far more often from key custody, signing-authority abuse, and operational error than from cryptographic breaks. The architecture therefore treats signing policy—who can sign what, under what authorisation, with what quorum, and with what revocation path—as a first-class design object rather than a configuration detail. This is also why the ledger layer must support freeze and correction: a system that cannot recover from an erroneous or coerced signing event will be worked around, and the workaround will be worse than the original problem.

Integration with legacy ERPs dominates cost. The marginal cost of the cryptographic components is small relative to the cost of extracting reliable event data from heterogeneous enterprise systems. Any serious cost-benefit analysis must therefore count data engineering and change management, not licences and nodes.

VIII. EVALUATION: SECURITY, COMPLIANCE, AND ECONOMICS

A. Security against the stated adversary model

Table 3. Adversary-by-adversary assessment

Adversary	Mechanism	Residual risk
A1 Opportunistic borrower (duplicate financing)	Deterministic receivable identifiers; private set intersection across institutions; mandatory registration with priority effect; split-tree constraints	Collusion with a second nominal supplier to present a different payer-payee pair over the same goods; identifier convention not adopted by all platforms
A2 Conflicted or dilatory anchor	Signed, timestamped event chain; regulatory tenor caps; mandatory confirmation within thirty days under Decree No. 802; trade-background corroboration from third parties	Anchor simply refuses to confirm; acceptance milestones manipulated; confirmation obtained but shipment never occurs
A3 Over-reachful platform	Platform excluded from fund flows; open financing interface; anomaly metrics on financier concentration; supervisory reporting	Platform controls the identifier convention or the reference data and can ration access; mitigated only by governance, not by cryptography
A4 Curious insider or external attacker	Confidential computing; purpose-bound authorisation; attested enclaves; minimised on-chain data	Insider with legitimate functional access; metadata inference from query patterns even when payloads are protected
A5 Compromised execution environment	Remote attestation; defence in depth; cryptographic fallback to MPC where the hardware trust assumption is unacceptable	Microarchitectural side channels; supply-chain compromise of the attestation root; mitigated, not eliminated

Two observations follow. First, the architecture's strongest contribution is to A1 and A4—precisely the two adversaries that ordinary platform designs handle least well. Second, its weakest contribution is to A2 and A3, and in both cases the residual risk is *governance* risk. This is not a defect of the design; it is a statement about where the remaining risk actually lives.

B. Compliance mapping

China. The architecture satisfies, by construction, the obligations that matter: genuine-trade verification through the event chain; tenor and splitting control through the state machine; priority through registration integration; ring-fenced settlement by excluding the platform from fund flows; data minimisation through attribute-based disclosure and proofs; and information-security assurance through attestation-based partner assessment. Evidentiary compliance is addressed by the evidence service, which retains pre-chain source documents and generation logs—the specific deficiency that Article 17 of the *Rules on Online Litigation* instructs courts to examine. Baseline system requirements are addressed under GB/T 22239-2019, with personal information handled under GB/T 35273-2020 and data-lifecycle capability assessed under GB/T 37988-2019.

Cross-border. For transactions touching other jurisdictions, the architecture is compatible with, but not automatically compliant under, three regimes. The UNCITRAL Model Law on Electronic Transferable Records (2017) requires that the method used establish singularity and exclusive control; the architecture's deterministic identifier and control-transfer state machine are designed to that standard, but adoption must be verified jurisdiction by jurisdiction, and the Model Law is a model, not a treaty. The United Kingdom's Electronic Trade Documents Act 2023 supplies a domestic path for specified trade documents. In the European Union, Regulation (EU) 2024/1183 establishes the wallet and electronic-attestation framework within which selective disclosure operates, while Regulation (EU) 2022/858 (the DLT pilot regime), Regulation (EU) 2023/1114 (MiCA), and Regulation (EU) 2023/1113 (transfer-of-funds information requirements) may apply depending on the activity actually performed. The operative principle is that regulatory classification follows *activity and*

Table 3 evaluates the architecture adversary by adversary. The value of the exercise lies in the right-hand column: in three of five cases the architecture reduces rather than eliminates risk, and stating this is more useful than claiming otherwise.

legal character, not the fact that distributed ledger technology is used; a permissioned network carrying internal receivable states is a different legal object from a network issuing transferable tokens to the public, and system designers must resist the assumption that a shared technical substrate implies a shared regulatory treatment.

Implementation caveat. Implementing acts, national transpositions, and supervisory guidance continue to evolve, and several technical annexes—including the detailed architecture and reference framework for European digital identity wallets—should be verified against current official sources before any cross-border deployment. The architecture is designed to accommodate change because the credential and settlement layers are substitutable; the state and confidential-computing layers are not jurisdiction-specific.

C. Economics

A deployment must be justified net of cost. The benefit side comprises four components: (i) reduction in credit losses from duplicate and fictitious financing, equal to the reduction in expected loss derived in Section 6.1; (ii) reduction in the cost of capital for confirmed suppliers, derived in Section 6.2; (iii) reduction in manual verification and reconciliation effort; and (iv) reduction in settlement and dispute-resolution time. The cost side comprises five: (i) integration with participant ERPs, which typically dominates; (ii) cryptographic computation and attestation infrastructure; (iii) key management and operational security; (iv) governance and coordination; and (v) ongoing compliance and audit.

Three properties of this cost structure determine viability, and they explain the adoption failures surveyed in Section 2.5.

Costs are front-loaded and concentrated; benefits are back-loaded and diffuse. Integration cost falls on each participant before any transaction flows. Financing-cost benefits accrue to suppliers; loss-reduction benefits accrue to banks; working-capital-discipline benefits accrue to regulators and to suppliers collectively. Unless the allocation is explicitly designed—for instance, by having the anchor buyer bear onboarding cost in exchange for supply-chain stability, which is itself a form of the cost-sharing that Yinfa [2025] No. 77 requires—the arrangement will undersupply.

Value scales superlinearly in participation and sublinearly in spend. Duplicate detection is more valuable the more institutions participate (Prop. 1), while additional spending on a single participant's own infrastructure yields diminishing returns. This asymmetry predicts both a coordination failure and a policy remedy: mandates that force simultaneous participation, such as registration requirements and institution filing regimes, can move the market where voluntary coordination cannot.

Unit economics must be positive per transaction, not per platform. The TradeLens discontinuation is the canonical illustration: technical capability and major-industry sponsorship were insufficient because network adoption and commercial viability were not achieved. A design principle follows directly: before building, specify the per-transaction charge, the party that bears it, and the benefit that party receives. If that calculation cannot be written down, the platform should not be built.

Two cautions from the market-microstructure literature reinforce this restraint. First, token-based networks exhibit strong adoption externalities and highly non-linear valuation dynamics, so early participation figures are poor predictors of steady-state value [14]. Second, evidence from open cryptocurrency markets shows both that illicit activity can be substantial and concentrated and that transparency enables its measurement—a useful reminder that traceability is an instrument of supervision, not a substitute for it [22], [32].

We decline to publish cost or performance benchmarks. Publicly available figures for secure computation, homomorphic encryption, and enclave performance vary by orders of magnitude with circuit size, security parameter, batch size, network conditions, and hardware generation, and we were unable to obtain a reproducible benchmark under stated conditions. Reporting a point estimate would give a spurious precision that the underlying evidence does not support. Instead we recommend that any deployment publish its own benchmark with the full parameter set—circuit or model specification, security parameter, batch size, network characteristics, hardware model and microcode version, and attestation configuration—as an appendix, and that the research community treat such appendices as a condition of publication.

IX. GOVERNANCE, STANDARDISATION, AND THE ORGANISATIONAL PROBLEM

A. Governance is the binding constraint

The argument so far establishes that the technical architecture is feasible and that its control points map onto binding obligations. It does not establish that anyone will build it, because the obstacles that stopped prior consortia were organisational. Four are decisive.

Neutrality. A network operated by a participant with a commercial interest in the outcome—the anchor buyer, a platform that also arranges financing, a bank that wishes to capture the lending—will be distrusted by the others regardless of its cryptography. The requirement under Yinfa [2020] No. 226 that platforms return to information services

and refrain from 变相 engaging in financial business is, in this light, a governance rule expressed as a licensing rule. The structural implication is that the state layer and the identifier convention should be operated by an entity with no lending interest, and that the identifier convention itself should be governed as a commons.

Exit and migration. Participants will not contribute data to a network from which they cannot leave with their records. Portability must be designed in: credentials are held by their holders, not by the platform; evidence packages are exportable in a standard format; and the identifier convention survives any operator's departure.

Benefit allocation. Where the anchor buyer captures most of the benefit through extended payables discipline and suppliers bear onboarding cost, adoption will be coerced rather than voluntary, and will be fragile. Yinfa [2025] No. 77's requirement of reasonable cost-sharing is the regulatory expression of this constraint.

Contestability and error correction. Every automated decision must be challengeable, and every state must be correctable through a defined, logged, and authorised process. Research on platform governance has long argued that the coevolution of architecture, governance, and environment determines whether an ecosystem thrives [49]; in a regulated financial setting, contestability is not merely good governance, it is the mechanism by which the system remains lawful when it is wrong.

B. Standardisation: what helps and what does not

Standardisation reduces interface ambiguity but cannot resolve conflicts of law or rights. The standards that matter for this architecture fall into four groups.

Credential and identity. W3C Verifiable Credentials Data Model 2.0 and the decentralised identifier specifications supply the expression layer; in the European context, the electronic-attestation and selective-disclosure mechanisms established by Regulation (EU) 2024/1183 supply the legal-institutional counterpart. The practical benefit is bilateral-integration avoidance: one credential format, many verifiers.

Business semantics. Financial messaging standards, notably ISO 20022, and trade-document data standards supply the event and field semantics without which two platforms will describe the same event differently. Semantic alignment at the event layer is more valuable than interface alignment at the API layer: an interface that connects two systems with incompatible definitions of “accepted” produces agreement at the wrong level of abstraction.

Ledger vocabulary and architecture. ISO/TC 307 terminology provides a shared vocabulary—note that ISO 22739:2020 has been superseded by ISO 22739:2024, and citations must specify the version. Terminology matters less than it appears; what matters is the state-transition semantics that a specific network defines and documents.

Security baseline. Trusted-computing interface and cryptographic-support-platform specifications (GB/T 29829-2022 in China; the TPM 2.0 family internationally) and

graded-protection requirements supply the assurance baseline for L1. Several other standards in this space are at varying stages of development, and precise designations should be verified against current official registers before being cited in compliance documentation.

The limit of standardisation is worth stating plainly: a standard can define how to express that a party has control of a record; it cannot determine whether control of that record transfers a property right in a given jurisdiction. Technical semantics and legal semantics must be mapped in two separate layers, and the mapping must be maintained as law changes.

C. Organisational adoption

Three conditions predict adoption in practice. First, the *anchor* must participate, because without its confirmation no amount of cryptography creates a verifiable receivable; regulatory pressure on confirmation timing, such as the thirty-day confirmation requirement in Decree No. 802, is often the decisive instrument. Second, *banks must retain decision rights*, because the alternative—platform-originated, platform-underwritten, bank-funded lending—has been identified by regulators as a source of hollowed-out risk management and is now constrained. Third, *suppliers must see a price benefit*, because they bear the onboarding burden; where financing rates do not fall, participation becomes a compliance exercise with no economic content.

X. CONCLUSION

A. Summary

This paper began from a paradox: supply chain finance should transfer credit from strong buyers to weak suppliers, and the technologies needed to make that transfer verifiable are available, yet the transfer remains limited and the instruments used to effect it remain contested. We argued that the obstacle is architectural rather than technological, and that the four most visible symptoms of failure—duplicate financing, unverifiable trade backgrounds, limited credit pass-through, and non-interoperability—are manifestations of one defect: the absence of a continuously verifiable link between a business event, a legal right, and a payment obligation.

We proposed a decomposition of trustworthiness into six properties and showed that each requires a distinct mechanism. We specified TCCE, a five-layer architecture that assigns exactly one assurance responsibility per layer and keeps raw commercial and personal data out of the shared state layer. We argued, and then formalised, that the layers are complements rather than substitutes: detection without legal priority leaves losses dilutive, and priority without detection leaves losses frequent. We showed that verifiable selective disclosure can strictly dominate both concealment and full disclosure, which reframes data minimisation from a compliance burden into a source of economic value. We showed that credit pass-through has an endogenous depth limit set by verification accuracy rather than verification cost, which redirects design effort from throughput to authoritative corroboration. And we mapped the architecture onto the actual control points that Chinese regulators have made binding,

while noting candidly that no verified statistic yet exists on the prevalence of the frauds the architecture is designed to prevent.

B. Limitations

Four limitations bound the claims. First, the models are stylised; they identify comparative statics and complementarity structure, not magnitudes. Second, the empirical content is institutional and descriptive: we report verified regulatory requirements and disclosed platform figures, but we have no access to transaction-level data and therefore offer no causal estimates. Third, the cross-border compliance discussion is indicative; implementing acts and national transpositions must be verified at the time of deployment. Fourth, we deliberately decline to report performance benchmarks, because available figures are not reproducible under stated conditions.

C. Research agenda

The field's scarcest resource is not technology or theory but verifiable production evidence. Five studies would materially change what can be claimed.

1. *Duplicate-financing incidence*. A multi-institution study using private set intersection over receivable identifiers, reporting the incidence of overlapping claims and the reduction attributable to registration integration. The protocol itself is the contribution: it is the only way to measure the baseline without any institution disclosing its books.

2. *Complementarity test*. A difference-in-differences design exploiting the staggered adoption of ledger infrastructure and registration integration across platforms, testing the negative interaction predicted by Proposition 2.

3. *Pass-through depth*. Measurement of the realised distribution of financing rates by tier, and its response to interoperability and third-party corroboration, testing Proposition 5's plateau prediction.

4. *Reproducible privacy-computation benchmarks*. Publication of full parameter sets—circuit or model, security parameter, batch size, network, hardware and microcode, attestation configuration—so that cost claims become comparable across studies.

5. *Governance and adoption*. Comparative institutional analysis of why some multi-party networks achieve participation and others, including well-funded ones, do not; the discontinuation of TradeLens at the end of the first quarter of 2023 and the restructuring of several bank trade-finance consortia are natural cases, provided primary announcements and terms are verified rather than reconstructed from secondary reporting.

A final word on framing. The literature has oscillated between enthusiasm for distributed ledgers as a substitute for institutions and scepticism about them as an expensive database. Both framings are unhelpful. The architecture advanced here treats cryptography as a way of making *specific assertions* checkable at *specific costs*, and treats institutions as

the source of the authority that makes those assertions legally operative. Trustworthy electronic business collaboration is not the elimination of trust in institutions; it is the engineering of the interface between what machines can verify and what only institutions can warrant.

REFERENCES

- [1] Acar, A., Aksu, H., Uluagac, A. S., and Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys*, 51(4), Article 79.
- [2] Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.
- [3] Auer, R., and Böhme, R. (2020). The technology of retail central bank digital currency. *BIS Quarterly Review*, March, 85–100.
- [4] Babich, V., and Hilary, G. (2020). Distributed ledgers and operations: What operations research can teach us. *Management Science*, 66(10), 4350–4359.
- [5] Biais, B., Bisière, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *Review of Financial Studies*, 32(5), 1662–1715.
- [6] Bindseil, U. (2020). Tiered CBDC and the financial system. *ECB Working Paper Series*, No. 2351.
- [7] Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., and Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 1175–1191.
- [8] Bordo, M. D., and Levin, A. T. (2017). Central bank digital currency and the future of monetary policy. *NBER Working Paper*, No. 23711.
- [9] Chen, X., Wang, X., and Jiang, X. (2019). The impact of power structure on supply chain financing. *International Journal of Production Economics*, 208, 1–13.
- [10] Chiu, J., and Koeppl, T. V. (2019). Blockchain-based settlement for asset trading. *Review of Financial Studies*, 32(5), 1716–1753.
- [11] Chod, J., Trichakis, N., Tsoukalas, G., Aspegren, H., and Weber, M. (2020). On the financing benefits of supply chain transparency and blockchain adoption. *Management Science*, 66(10), 4378–4396.
- [12] Cole, R., Stevenson, M., and Aitken, J. (2019). Blockchain technology: Implications for operations and supply chain management. *Supply Chain Management: An International Journal*, 24(4), 469–483.
- [13] Cong, L. W., and He, Z. (2019). Blockchain disruption and smart contracts. *Review of Financial Studies*, 32(5), 1754–1797.
- [14] Cong, L. W., Li, Y., and Wang, N. (2021). Tokenomics: Dynamic adoption and valuation. *Review of Financial Studies*, 34(3), 1105–1155.
- [15] Costan, V., and Devadas, S. (2016). Intel SGX explained. *IACR Cryptology ePrint Archive*, Report 2016/086.
- [16] Dwork, C., and Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407.
- [17] European Union. (2014). Regulation (EU) No 910/2014 on electronic identification and trust services for electronic transactions in the internal market (eIDAS).
- [18] European Union. (2022). Regulation (EU) 2022/858 on a pilot regime for market infrastructures based on distributed ledger technology.
- [19] European Union. (2023a). Regulation (EU) 2023/1114 on markets in crypto-assets (MiCA).
- [20] European Union. (2023b). Regulation (EU) 2023/1113 on information accompanying transfers of funds and certain crypto-assets.
- [21] European Union. (2024). Regulation (EU) 2024/1183 of 11 April 2024 amending Regulation (EU) No 910/2014 as regards establishing the European Digital Identity Framework.
- [22] Foley, S., Karlsen, J. R., and Putnăș, T. J. (2019). Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies? *Review of Financial Studies*, 32(5), 1798–1853.
- [23] Gelsomino, L. M., Mangiaracina, R., Perego, A., and Tumino, A. (2016). Supply chain finance: A literature review. *International Journal of Physical Distribution & Logistics Management*, 46(4), 348–366.
- [24] Goldfarb, A., and Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1), 3–43.
- [25] Hofmann, E., Strewe, U. M., and Bosia, N. (2018). *Supply Chain Finance and Blockchain Technology: The Case of Reverse Securitisation*. Springer.
- [26] Hughes, A., Park, A., Kietzmann, J., and Archer-Brown, C. (2019). Beyond Bitcoin: What blockchain and distributed ledger technologies mean for the enterprise. *Business Horizons*, 62(3), 311–321.
- [27] Jones, C. I., and Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, 110(9), 2819–2858.
- [28] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... and Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.
- [29] Klapper, L. (2006). The role of factoring for financing small and medium enterprises. *Journal of Banking & Finance*, 30(11), 3111–3130.
- [30] Kshetri, N. (2021). Blockchain and the economics of supply chain finance. *International Journal of Information Management*, 61, Article 102361.
- [31] Lumineau, F., Wang, W., and Schilke, O. (2021). Blockchain governance—A new way of organizing collaborations? *Organization Science*, 32(2), 500–521.
- [32] Makarov, I., and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2), 293–319.
- [33] Mohassel, P., and Zhang, Y. (2017). SecureML: A system for scalable privacy-preserving machine learning. *Proceedings of the 2017 IEEE Symposium on Security and Privacy (S&P)*, 19–38.
- [34] Moretto, A., Grassi, L., Caniato, F., Giorgino, M., and Ronchi, S. (2019). Supply chain finance: From traditional to supply chain credit rating. *Journal of Purchasing and Supply Management*, 25(2), 197–217.
- [35] National Bureau of Statistics of China. (2024, 2025). Statistical releases on accounts receivable and average collection period of above-scale industrial enterprises, 2023 and 2024.
- [36] National Internet Finance Association of China. (2026). Self-regulatory report on accounts-receivable electronic certificate business and supply chain information service institutions, first quarter of 2026.
- [37] People’s Bank of China. (2025). *China Monetary Policy Report*, Quarter Four, 2024.
- [38] Saberi, S., Kouhizadeh, M., Sarkis, J., and Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Economics*, 207, 21–35.
- [39] Sabt, M., Achemlal, M., and Bouabdallah, A. (2015). Trusted execution environment: What it is, and what it is not. *Proceedings of the 14th IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, 57–64.
- [40] Shanghai Commercial Paper Exchange. (2020, 2022). Announcements and public materials on the supply-chain bill platform and its access rules.
- [41] State Council of the People’s Republic of China. (2020). Regulation on Ensuring Payment to Small and Medium-Sized Enterprises (Decree No. 728).
- [42] State Council of the People’s Republic of China. (2025). Regulation on Ensuring Payment to Small and Medium-Sized Enterprises (revised, Decree No. 802), published 17 March 2025, effective 1 June 2025.
- [43] Supreme People’s Court of the People’s Republic of China. (2021). Rules on Online Litigation of the People’s Courts, effective 1 August 2021.
- [44] The People’s Republic of China. (1995, as amended). Law on Negotiable Instruments.
- [45] The People’s Republic of China. (2004, as amended). Electronic Signature Law.
- [46] The People’s Republic of China. (2021). Personal Information Protection Law.
- [47] The People’s Republic of China. (2021). Civil Code (Book on Contracts, factoring provisions, Arts. 761–768).
- [48] Tian, F. (2016). An agri-food supply chain traceability system for China based on RFID & blockchain technology. *Proceedings of the 13th*

- International Conference on Service Systems and Service Management (ICSSSM), 1–6.
- [49] Tiwana, A., Konsynski, B., and Bush, A. A. (2010). Research commentary—Platform evolution: Coevolution of platform architecture, governance, and environmental dynamics. *Information Systems Research*, 21(4), 675–687.
- [50] UNCITRAL. (2017). Model Law on Electronic Transferable Records.
- [51] United Kingdom. (2023). Electronic Trade Documents Act 2023, c. 14 (Royal Assent 28 June 2023; in force 20 September 2023).
- [52] Van Bulck, J., Minkin, M., Weisse, O., Genkin, D., Kasikci, B., Piessens, F., Silberstein, M., Wenisch, T. F., Yarom, Y., and Strackx, R. (2018). Foreshadow: Extracting the keys to the Intel SGX kingdom with transient out-of-order execution. *Proceedings of the 27th USENIX Security Symposium*, 991–1008.
- [53] Wang, Y., Han, J. H., and Beynon-Davies, P. (2019). Understanding blockchain technology for future supply chains: A socio-technical perspective. *Supply Chain Management: An International Journal*, 24(1), 80–96.
- [54] Werbach, K., and Cornell, N. (2017). Contracts ex machina. *Duke Law Journal*, 67(2), 313–382.
- [55] Wuttke, D. A., Blome, C., and Henke, M. (2013). Focusing the financial flow of supply chains: An empirical investigation of financial supply chain management. *International Journal of Production Economics*, 145(2), 773–789.
- [56] Yang, Q., Liu, Y., Chen, T., and Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), Article 12.
- [57] Yermack, D. (2017). Corporate governance and blockchains. *Review of Finance*, 21(1), 7–31.
- [58] People’s Bank of China et al. (2020). Opinions on Regulating the Development of Supply Chain Finance to Support the Stable Circulation and Optimisation and Upgrading of Supply Chains and Industrial Chains (Yinfa [2020] No. 226), issued 18 September 2020, published 22 September 2020.
- [59] People’s Bank of China, National Financial Regulatory Administration, Supreme People’s Court, National Development and Reform Commission, Ministry of Commerce, and State Administration for Market Regulation. (2025). Notice on Regulating Supply Chain Finance Business and Guiding Supply Chain Information Service Institutions to Better Serve SME Financing (Yinfa [2025] No. 77), issued 26 April 2025, effective 15 June 2025.
- [60] Standardisation Administration of China. (2019). GB/T 22239-2019: Information security technology — Baseline for classified protection of cybersecurity.
- [61] Standardisation Administration of China. (2019). GB/T 37988-2019: Information security technology — Data security capability maturity model.
- [62] Standardisation Administration of China. (2020). GB/T 35273-2020: Information security technology — Personal information security specification.
- [63] Standardisation Administration of China. (2022). GB/T 29829-2022: Information security technology — Functionality and interface specification of cryptographic support platform for trusted computing (issued 15 April 2022, effective 1 November 2022).
- [64] International Organization for Standardization. (2024). ISO 22739:2024 Blockchain and distributed ledger technologies — Vocabulary (supersedes ISO 22739:2020).
- [65] W3C. (2025). Verifiable Credentials Data Model v2.0, W3C Recommendation, 15 May 2025.

APPENDIX A. NOTATION

Symbol	Meaning
F	Face value of the receivable
λ	Advance rate ($0 < \lambda < 1$)
n	Number of potential financiers
T	Term of the financing
r	Financing rate
δ	Detection capability: probability that a prior financing is discovered before advancing
ρ	Priority effectiveness: share recovered by virtue of registration or notice
m	Number of advances actually made
$R(\delta, \rho)$	Expected recovery per financier
σ_0^2, σ^2	Prior and posterior variance of the supplier’s default risk
γ	Risk-pricing coefficient
κ	Marginal competitive harm of disclosing proprietary information
$s(e)$	Variance reduction achieved by evidence e
$l(e)$	Proprietary information leaked by evidence e
v	Cost of verifying one supply-chain linkage
ε	Probability that a linkage verification accepts a spurious linkage
t^c	Endogenous depth limit of credit pass-through

APPENDIX B. ARCHITECTURE LAYERS AND THEIR SOLE RESPONSIBILITIES

Layer	Sole responsibility	Must never hold
L1 Trust anchor	Evidence of computational provenance	Business data
L2 Credential	Identity, authority, and attribute assertions	Underlying source documents
L3 Confidential computing	Computation over data it cannot disclose	Authoritative state
L4 Shared state	Ordering, commitments, state transitions, proof handles	Raw commercial or personal data
L5 Execution and settlement	Registration, payment, discharge, finality	Discretion over credit decisions