

Supervisory Verifiability Across Crypto-Fiat Rails: Evidence Sufficiency, Cross-Rail Linkage, and AI-Assisted Supervision in Da Nang

Pham Quoc Trung¹, Alfred Taudes²

¹ Da Nang Fintech Lab (DFL24), Da Nang, Vietnam; ² Research Institute for Cryptoeconomics, WU Vienna, Austria

Abstract- Crypto-fiat conversion spans independently governed systems, including public blockchains, operators, identity providers, and bank or payment rails. A valid record on each rail does not by itself establish that the records refer to the same economic transaction. This study defines supervisory verifiability in terms of evidence access, source authority, cross-rail linkage, and decision timing, and uses that criterion to constrain AI-assisted evidence orchestration. AI components may extract approved controls, request records, propose linkages, and route unresolved cases, whereas compliance determination remains rule-based and auditable. Empirical grounding combines double-coded regulatory mapping with an initial coded incident-reporting obligation from the governing instruments of a Da Nang controlled trial. The synthetic evaluation design uses known-ground-truth cases with missing evidence, delay, batching, false attestations, transaction density, and linkage ambiguity. The analysis distinguishes configuration-level verifiability from case-level outcomes, derives obligation-specific inclusion-minimal evidence configurations, and specifies a comparison of operator-only reporting, deterministic multi-source verification, and AI-assisted orchestration. The study therefore provides a control-level supervisory data contract, explicit conditions for observational impossibility, and run-derived measures of evidence burden and unresolved-case routing without treating model output as evidence.

Keywords- *SupTech; AI-assisted supervision; crypto-fiat conversion; supervisory verifiability; cross-rail linkage*

1. Introduction

Crypto-fiat conversion creates a supervisory problem when the asset leg, identity controls, and fiat settlement are recorded by different institutions. Observability of each component does not imply verifiability of the composite transaction. A blockchain record may establish that a digital-asset transfer occurred and a bank record may establish that fiat settlement occurred, yet neither record establishes that the two observations belong to the same conversion unless an authoritative linkage is available. The research problem is therefore to identify the evidence conditions under which a regulatory obligation is decidable within a defined supervisory time window.

The study addresses two questions. RQ1 asks under what evidence-access, trust, cross-rail linkage, and timing conditions a supervisory obligation is independently verifiable, and which evidence configurations are minimally sufficient. RQ2 asks whether constrained AI-assisted orchestration can reduce evidence-request burden and routine unresolved-case routing while preserving the same deterministic and auditable verification criterion. RQ1 concerns what can be established from an institutional evidence regime; RQ2 concerns how that regime can be executed at scale without treating model output as evidence.

The contribution is threefold. The formal contribution is a trust-, linkage-, and delay-aware criterion for obligation-level verifiability. The systems contribution is an AI-assisted supervisory workflow in which evidence acquisition and exception handling may be automated while the compliance result remains deterministic. The applied contribution is a document-grounded coding scheme for Da Nang's expanding sandbox, instantiated on an initial coded obligation, together with a synthetic evaluation design for linkage ambiguity and evidence-handling burden. The analysis concerns the architecture of supervision rather than the legality or economic desirability of crypto-assets themselves.

2. Related Work and Research Setting

Prior research addresses separate parts of the problem. Auer [1] develops embedded supervision, in which trusted ledger data support regulatory monitoring. Taudes et al. [2] examine DLT-based securities settlement in an integrated regulated infrastructure with approved workflows for matching, settlement, and oversight. Ly et al. [3] formalise compliance monitoring over heterogeneous and partially observable business-process data. The present setting differs because the facts required for one regulatory obligation are distributed across systems with distinct evidence owners, access rights, and trust properties. Integrated settlement can make transaction linkage part of the infrastructure; fragmented crypto-fiat rails make the same linkage an evidentiary condition that must be established separately. The novelty is therefore not observational equivalence itself, but its instantiation under an observation model whose mask is determined by legal access rights, source authority, and cross-rail linkage across separately governed systems.

Da Nang is used as the empirical setting because several controlled trials operate under one supervisory jurisdiction. Decisions 1181/QD-UBND and 2895/QD-UBND provide two earlier documentary cases [4,5], the

latter covering non-custodial USDT-VND conversion. Decisions 3809-3812/QD-UBND added four further controlled trials in August 2026 [6]. The portfolio spans crypto-fiat conversion, tourism/payment use cases, gateway integration, and cross-border settlement infrastructure. These trials concern controlled conversion and settlement rather than domestic stablecoin issuance.

The earlier decisions provide the longest documentary record; the later approvals test whether the same obligation-evidence representation remains applicable as the portfolio expands. The unit of analysis is the supervisory obligation and its evidence regime, not the operator. Architecture is coded from documented evidence ownership, settlement roles, and approved controls. Where primary instruments do not establish custody, settlement topology, or blockchain network, the corresponding field is recorded as not publicly established.

3. Supervisory Verifiability

Formal criterion. Let $O=\{o_1,\dots,o_m\}$ denote supervisory obligations and $D=\{C,P,F,I\}$ the evidence domains: public chain, platform/operator, fiat bank or payment rail, and identity provider. For obligation o , S denotes the evidence domains available to the supervisor, T the sources authorised to establish the relevant facts, L the authoritative relation linking records across rails, and k the latest admissible decision time. Let $Obs_k(x; S,T,L)$ denote the authoritative observations about execution x that are accessible from S , accepted under T , linked by L , and available no later than k . $V(o,S,T,L,k)=1$ iff no compliant execution x^+ and non-compliant execution x^- yield the same $Obs_k(x; S,T,L)$; otherwise $V=0$. Thus verifiability is a property of the evidence regime, not of data volume alone.

The definition is related to diagnosability and decentralised diagnosis [7,8], while opacity gives the dual interpretation of observational equivalence as a protection objective [9]. These concepts provide formal grounding without claiming a complete discrete-event-system diagnoser in the present abstract. Trust is represented explicitly because verifiability under partially trusted parties is weaker than full accountability [10]. The time bound k retains the distinction between eventual observability and observability early enough to support supervisory action.

Configuration and case outcomes. V is a property of an evidence regime rather than a label for a particular transaction. For execution x , the operational result is $R(o,x)$ in $\{\text{verified compliance, verified non-compliance, inconclusive}\}$. A regime may be verifiable in principle while a particular case remains inconclusive because a required record has not yet arrived or because several admissible linkages remain. Conversely, $V=0$ means that at least one compliant and one non-compliant execution remain observationally indistinguishable under the specified regime. An inconclusive result is not evidence of non-compliance.

Minimal evidence configurations. For each obligation, the relevant object is the family $M(o;T,L,k)$ of inclusion-minimal source configurations S for which $V=1$. A single obligation may have several incomparable minimal configurations, for example when one authoritative source can substitute for another. This representation supports a supervisory data contract that specifies the smallest admissible evidence combinations rather than a universal request for all records generated by all participants.

Worked impossibility case. Assume the supervisor observes an authentic chain transfer c and an authentic fiat settlement f . In execution x^+ , c and f are the two legs of the same compliant conversion. In execution x^- , both records are genuine but belong to different economic transactions. If the only assertion pairing c with f is supplied by an institution that is not authoritative for linkage L , the authoritative observation is identical in x^+ and x^- . The supervisor cannot distinguish the executions and $V=0$ for the obligation requiring correspondence between the two legs. Additional authentic records restore verifiability only if they establish the missing relation; a second truthful record that is unrelated to the first does not do so. The example separates evidence authenticity from evidence correspondence, which are often conflated in cross-rail reporting.

4. AI-Assisted Supervisory Architecture and Evaluation

AI components are limited to evidence and workflow functions. Following human validation of machine-readable controls, they may select among prevalidated admissible evidence configurations, request the corresponding records, propose candidate cross-rail linkages, identify the source of an inconclusive result, and route cases requiring additional evidence or legal review. The deterministic verifier evaluates V and R from authoritative evidence. AI output is non-evidentiary: it cannot create an authoritative fact, admit an unvalidated evidence configuration, lower the verification threshold, or determine a legal sanction. This separation is consistent with machine-executable reporting and integrated analytics in BIS Project Ellipse [11] and with the constrained supervisory use of AI being tested in Project AISE [12].

Approved obligations Human-validated controls	AI orchestration Select prevalidated evidence set; propose linkage	Deterministic verification Evaluate V and case result R	Supervisory action Escalate material or legal cases
--	--	--	--

Figure 1. Functional separation in the AI-assisted supervisory architecture. AI operates only within prevalidated evidence configurations; deterministic rules determine case status; material or legal decisions remain human-authorised.

Evaluation baselines. B0 is the operator-only evidence regime: the supervisor receives the operator's report without independently integrated evidence. B1 adds deterministic verification over multiple evidence domains but uses fixed evidence requests and rule-based exception routing. B2 adds AI-assisted selection among prevalidated evidence configurations, evidence requests, candidate linkage, and routing of unresolved cases. Synthetic cases with known ground truth are processed under all three regimes. The comparison uses run-derived outputs only; it does not treat a synthetic run as a measurement of human handling time.

Evaluation measures. The principal measures are evidence items requested per resolved case, additional-evidence routing rate, determinate-resolution rate, inconclusive rate, reconciliation error, and unsupported AI assertion rate. Tests vary missing or false attestations, reporting delay, reversals, 1:N and N:1 batching, and conflicts between evidence sources. Human case-handling time and escalation duration are reserved for a later handler study rather than inferred from synthetic runs. Public-chain attribution is evaluated only when the settlement network and an independently anchored reference set are available; otherwise the case is recorded as not evaluable rather than as an attribution failure. This restriction follows from the distinction between public-chain visibility and reliable attribution [13].

Scale test. Cross-rail correspondence can appear reliable at pilot scale even when it is based on weak identifiers such as amount and a time window. Let $A(n, \Delta t)$ denote the mean number of admissible fiat candidates for a chain event at transaction density n and matching window Δt . When A is close to one, heuristic pairing may appear unambiguous; as n or Δt increases, several valid candidates can satisfy the same rule and the case becomes inconclusive unless an additional authoritative identifier is available. Synthetic traffic therefore tests whether a supervisory design that works at low volume remains verifiable as transaction density and batching increase.

5. Empirical Grounding, Supervisory Data Contract, and Limitations

Empirical analysis begins with documentary coding. Two independent coders use the Vietnamese primary instruments to identify each obligation, the facts required to establish it, the authoritative evidence owner, any required cross-rail linkage, and the decision deadline. Signed or formally issued instruments are prioritised; portal summaries and vendor descriptions are used only for attributes not established in the primary text and are recorded with their source tier. Disagreements are reconciled before a control is admitted to the coding scheme. The initial coded result below is instantiated from one approved crypto-fiat conversion trial whose licence and governing resolution state both the reporting duty and its timing [5]; it is used only as a documentary instance of the coding method.

Table 1. Initial coded supervisory obligation: evidence conditions for incident-reporting timeliness [5].

Coded obligation	Evidence owner(s) and deadline	Verification implication
Serious-incident / risk reporting timeliness	Duty: the trial operator must report to the licensing authority immediately upon detection of an operational interruption, serious risk, or uncontrollable abuse risk. Evidence owners: trial-operator incident record for detection time; licensing-authority receipt/submission record for report time.	Receipt proves report time, not promptness. Timeliness is independently verifiable only if detection time is authoritative or independently constrained; otherwise the case remains inconclusive.

This row is an initial coded result, not a performance estimate. The authority receipt establishes report time; promptness also requires an authoritative or independently constrained detection time. If detection time is controlled only by the reporting party, the timing obligation remains inconclusive. This instance therefore exercises T and k, not L; L-dependent conversion and settlement obligations are the next coding target.

FATF Recommendation 16 and IVMS101 standardise originator and beneficiary information, while ISO 20022 standardises payment and reconciliation messages [14,15]. These standards are complementary to the framework. They do not, by themselves, determine which public-chain event corresponds to a particular fiat settlement, which institution is authoritative for that relation, or whether the relevant evidence is accessible within k . The minimum supervisory data contract is therefore derived at obligation level from the inclusion-minimal evidence configurations, subject to privacy, reporting burden, and institutional responsibility.

S , T , L , and k are institutional evidence conditions rather than Vietnamese legal categories; another jurisdiction changes the applicable obligations, authorised sources, and time limits, not the verification question. The same distinction separates asset control from evidence control: non-custodial execution can remain supervisorily dependent when one participant controls the authoritative linkage or the information needed to reconstruct the fiat leg.

Regulatory implications. S , T , L , and k map respectively to access rights, authoritative evidence, transaction-linkage rules or identifiers, and reporting/finality deadlines. At scale, increasing transaction density can make authoritative identifiers more important even when each rail remains individually observable.

AI governance and privacy. Automation changes evidence handling, not evidentiary authority. AI-proposed requests or linkages retain provenance to the approved obligation, source records, and deterministic rule; adverse or legally material findings remain human-authorised. Data minimisation follows because the system requests one admissible evidence configuration rather than all available records. Signed or zero-knowledge proofs may later support attribute or threshold controls.

The study is ongoing and this abstract does not report quantitative performance estimates or causal effects of the Da Nang trials. A false original attestation is not detectable unless constrained by an independent authoritative source, and authentic records remain insufficient when their correspondence is not authoritative. Synthetic evaluation can test implementation behaviour and scale-related ambiguity but cannot substitute for external validation on approved transaction data. The policy claim is limited to supervisory process design: evidence acquisition and routine exception handling may be automated, while evidentiary authority, adverse findings, and legal judgment remain auditable and institutionally assigned.

References

- [1] R. Auer, "Embedded supervision: how to build regulation into decentralised finance," BIS Working Papers No. 811, 2019.
- [2] A. Taudes, J. Hackel, H. Hermanky, and W. Haunold, "Distributed ledger technologies for securities settlement - the case for running T2S on DLT," Monetary Policy & the Economy, Q2/21, pp. 13-33, 2021.
- [3] L. T. Ly, F. M. Maggi, M. Montali, S. Rinderle-Ma, and W. M. P. van der Aalst, "Compliance monitoring in business processes: functionalities, application, and tool-support," Information Systems, vol. 54, pp. 209-234, 2015.
- [4] Da Nang People's Committee, Decision 1181/QD-UBND, controlled trial, 26 Aug. 2025.
- [5] Da Nang People's Committee, Decision 2895/QD-UBND, 6 Dec. 2025; Da Nang Department of Science and Technology, License 1724/GP-SKHCN, 18 Dec. 2025; Da Nang People's Council, Resolution 55/2024/NQ-HDND, Art. 7(3). Available: <https://dragonlab.com.vn/cert/1724.GP.SKHCN.signed.signed.pdf>; <https://vbpl.vn/TW/Pages/vbpq-toanvan.aspx?ItemID=173630>.
- [6] Da Nang People's Committee, Decisions 3809-3812/QD-UBND, four controlled trials, Aug. 2026. Portal: <https://cttdt.danangportal.gov.vn/>.
- [7] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis, "Diagnosability of Discrete-Event Systems," IEEE Trans. Automatic Control, vol. 40, no. 9, pp. 1555-1575, 1995.
- [8] R. Debouk, S. Lafortune, and D. Teneketzis, "Coordinated Decentralized Protocols for Failure Diagnosis of Discrete Event Systems," Discrete Event Dynamic Systems, vol. 10, no. 1, pp. 33-86, 2000.
- [9] S. Lafortune, F. Lin, and C. N. Hadjicostis, "On the History of Diagnosability and Opacity in Discrete Event Systems," Annual Reviews in Control, vol. 45, pp. 257-266, 2018.
- [10] R. Küsters, T. Truderung, and A. Vogt, "Accountability: Definition and Relationship to Verifiability," Proc. ACM CCS, pp. 526-535, 2010.
- [11] BIS Innovation Hub, "Project Ellipse: an integrated regulatory data and analytics platform," 2022. Available: <https://www.bis.org/project/ellipse>.
- [12] BIS Innovation Hub, Toronto Innovation Centre, "Project AISE: AI-enabled supervisory analytics for retail payments supervision," 15 Apr. 2026. Available: https://www.bis.org/about/bisih/topics/suptech_regtech/aise.htm.
- [13] A. Wahrstätter, A. Taudes, and D. Svetinovic, "Reducing Privacy of CoinJoin Transactions: Quantitative Bitcoin Network Analysis," IEEE Trans. Dependable and Secure Computing, vol. 21, no. 5, pp. 4543-4558, 2024.
- [14] FATF, Recommendation 16 and Interpretive Note, amended June 2026; interVASP, IVMS101.2023, Issue 1 Final, 2024. Available: <https://www.fatf-gafi.org/>; <https://intervasp.org/>.
- [15] ISO 20022 Registration Authority, Catalogue of Messages. Available: <https://www.iso20022.org/catalogue-messages> (accessed Aug. 2026).