Programmable Money in Three Layers

Value, purpose, and evidence in stablecoin settlement

Jace Shim · November 2026 · https://doi.org/10.5281/zenodo.24580931

What makes a stablecoin payment programmable?

A stablecoin payment becomes programmable when a specified rule can govern the release of funds based on evidence that a system can evaluate. The token provides the means of settlement. The payment agreement defines the condition. A separate evidence process determines whether the condition has been met. This paper develops that distinction from Chapter 4 of Target Money (Shim, 2026) and asks when it is worth keeping those three responsibilities separate.

Abstract

A donor, public agency, or business may need to commit funds before the event that justifies payment. A stable-value token can make funds available for settlement, but its transfer record provides limited information about what happened outside the payment system. The value–purpose–evidence model identifies three responsibilities in conditional settlement and proposes a comparison of integrated and modular designs. The principal hypothesis is that modularity is most valuable when evidence requirements change frequently while the underlying settlement obligation remains stable. Integration can have lower costs when a program has few exceptions and its evidence sources rarely change.

1. A deposit is easy to count. A completed service is harder.

Consider a municipal agency that funds an hour of respite care. A transfer to a provider can show when and how much the agency paid. The agency still needs a way to decide whether the agreed care took place. A provider's attendance record may be useful, but a missing entry might reflect a failed terminal, a late submission, or a service that never happened. A system that treats every missing entry as nonperformance will produce a clean ledger and may leave legitimate invoices unpaid.

The same difficulty appears in software contracts. A payment system can settle an invoice immediately after authorization, while the buyer and developer may disagree about whether the submitted software meets the milestone. Both examples motivate the same design question: which part of a payment system should change when parties revise what counts as evidence?

2. A three-layer model of conditional settlement

The value layer determines the unit, amount, funding and redemption route. A fiat-referenced stablecoin could serve as the settlement asset where counterparties can use and redeem it under acceptable conditions. Custody, liquidity, settlement finality and conversion costs belong to this layer.

The purpose layer defines what the parties agreed to fund. It identifies the payer and permitted payee, specifies an amount or budget, states the event that triggers payment, and assigns authority to grant exceptions or resolve disputes. The model treats authorization as a distinct decision even if one institution operates multiple layers.

The evidence layer translates events into information the decision process can use. Each claim has a source, an observation time and a status. The status can be sufficient, incomplete or disputed. A signed claim establishes who made a statement and whether it changed in transit. Whether the claim accurately describes the service depends on the evidence process and its oversight.

An implementation can use one ledger or several systems. Functional separation means that the settlement asset, rule and evidence source can be identified and changed under documented procedures. For an individual claim, let V describe the availability of committed funds, R the rule in force, and E the evidence status. Where funds are available and the evaluator finds the rule satisfied, the claim can be authorized for release. Incomplete or contested evidence sends the claim to review. An agreed timeout governs cases that remain unresolved. A payment reference ties the decision to the invoice while sensitive supporting records can remain with the responsible institution.

3. When does separation pay off?

Compare two ways to process the same contracts: an integrated system in which the payment provider controls the rule and verification process, and a modular system with documented boundaries between those functions. Let A denote the cost of adapting to changes in contracts or evidence sources, H the expected harm from mistaken payment decisions, and K the cost of coordinating and operating the system. The modular design improves measured performance if its reductions in A and H together exceed any additional K. Report components separately because some design choices shift costs from the payer to the provider.

Suppose two programs use the same currency and payment rail. In one, a licensed provider submits a standard attendance record. In the other, multiple sources and frequent policy changes make evidence irregular. An integrated system might handle the first program at low cost. The second could benefit from replacing a verification process without migrating balances or modifying the underlying payment asset. This advantage depends on having usable interfaces, a versioned rulebook and sufficient review capacity.

4. Three propositions and a test

P1. Evidence changes. Holding the settlement rail fixed, programs that change evidence sources more often should see larger reductions in reconfiguration time from a modular design. Measure staff hours, interruption time and the number of affected integrations.

P2. Review and error. Moving incomplete claims to a time-limited review queue should reduce mistaken denials relative to immediate rejection. Compare the reduction with delays, administrative cost and possible improper releases. Independently adjudicate a sample of claims to estimate error in each direction.

P3. Settlement asset. Holding the policy and evidence process constant, stablecoin settlement should improve end-to-end performance only where gains such as more reliable timing or lower transfer costs exceed custody, conversion, reconciliation and exception costs. Compare against conventional transfers using the same business rules.

A pilot could assign comparable contracts to different combinations of rule architecture and settlement rail. Record time from service completion to usable funds, authorization errors, rule-update costs and review outcomes. Where a randomized comparison is impractical, compare phased changes with similar contracts processed under the original arrangements. A pilot should track merchants who decline participation as well as those who complete a transaction.

5. How this paper relates to the earlier working papers

When Do Programmable Vouchers Improve Public Spending? evaluates checkout enforcement using eligible completion, improper spending and total cost (Shim, 2026). Its question is whether a particular enforcement policy improves benefit delivery. The present paper asks which system owns the rule, the evidence and the payout, and whether each can change without forcing a change in the others. It can apply before checkout, after service delivery, or during a dispute. The Reserve–Distribution Feedback Model addresses how issuer reserve income may finance payment-network growth. The unit here is an individual conditional claim and the institutions responsible for resolving it.

The Monetary Authority of Singapore's Purpose Bound Money whitepaper supplies a relevant precedent for attaching usage conditions to existing forms of money. The present model emphasizes the treatment and replacement of external evidence sources and sets out observable measures for comparing designs. The propositions require empirical testing and the examples are hypothetical.

Keywords

Programmable money · stablecoin settlement · purpose-bound money · payment authorization · conditional payments · off-chain evidence

References

Monetary Authority of Singapore. (2023). Purpose Bound Money Technical Whitepaper. https://www.mas.gov.sg/publications/monographs-or-information-paper/2023/purpose-bound-money-whitepaper

Shim, J. (2026). Target Money (Chapter 4). Raon Book.

Shim, J. (2026). When Do Programmable Vouchers Improve Public Spending? Zenodo. https://doi.org/10.5281/zenodo.24578261

Shim, J. (2026). The Stablecoin Reserve–Distribution Feedback Model. Zenodo. https://doi.org/10.5281/zenodo.24577526