Who Should Authorize Payment for Care?

The Agent–Purpose Binding Model for stablecoin settlement

Jace Shim · November 2026 · https://doi.org/10.5281/zenodo.24581092

Abstract

A town pays for a home-care visit, a provider submits the claim, and an older resident receives the service. Each party knows something that the others cannot immediately verify. This paper develops the Agent–Purpose Binding Model from Chapter 8 of Target Money. It proposes that a payment agreement link each care installment to evidence of delivery and assign responsibility for decisions when accounts differ. Incomplete records lead to time-limited review, and the undisputed portion of a claim can be settled separately where the program permits it. Stablecoins can carry the approved transfer, while the decision to release funds remains with the parties specified in the agreement. The paper proposes a pilot that separates the effect of the approval design from the effect of the payment rail.

1. The visit that went missing

Imagine a home-care worker who completed a scheduled visit but could not check out on a shared device. The provider submits an invoice. The recipient remembers the visit but does not use a smartphone. The municipal payer sees an incomplete record. Paying immediately risks accepting an unsupported claim. Rejecting immediately makes the provider bear the cost of a missing record and can discourage visits to recipients who need help using digital tools.

There is a modest precedent for collecting factual visit records. Electronic Visit Verification applies to certain in-home Medicaid services in the United States (Centers for Medicare & Medicaid Services, n.d.). A visit record can help establish attendance. It may leave questions about the scope of work or a disputed change in the care plan. Holmström's (1979) work on observability and incentives frames the question of what information is worth using when a principal cannot observe an agent's actions directly. The care setting adds a further concern: the beneficiary can have direct knowledge of the visit and still lack practical means to challenge a decision.

2. What is agent–purpose binding?

Agent–purpose binding is a conditional payment arrangement in which the funder specifies the service to be purchased, the provider submits a claim against that purpose, and a designated reviewer can resolve conflicting evidence before the disputed amount is released. The agreement gives the recipient or an authorized representative a practical way to report an error. In this paper, “agent” refers to the service provider in a principal–agent relationship.

At enrollment, the payer and provider agree on the funded service unit, the installment amount and acceptable evidence. They also identify who can review exceptions and when a decision is due. An ordinary claim can be approved on the basis of an authenticated

record that meets the program's evidence standard. A missed check-out can be documented through an assisted channel. Conflicting accounts move to review, with reasons recorded and a route to challenge the decision. The reviewer should be independent of the provider's financial interest in the disputed amount.

For an installment with an agreed ceiling M, let A be the amount already supported by sufficient evidence. The payment rule allows release of A and reserves up to $M - A$ for adjudication, with $0 \le A \le M$. Where the visit is indivisible, A is zero until the visit is approved or the reviewer authorizes a correction. If several visits appear on one invoice, the design can release the verified visits while the others are reviewed. A separate authorization check precedes settlement so a decision can be amended before funds are sent.

For example, suppose four visits are billed at 40 units each. Three records meet the evidence standard and the last has a missing checkout. Where partial settlement is permitted, 120 units can be paid while a reviewer checks the remaining 40. If the reviewer confirms that the visit occurred, the remainder can be paid. This example illustrates an allocation of decision rights and delay. It supplies no evidence about the prevalence of disputed claims.

3. When does a stablecoin help?

A fiat-referenced stablecoin can be used to settle an approved claim where the payer and provider can access a reliable payment and redemption route. The practical gain depends on the time between approval and funds becoming usable. For a small provider, that interval includes receiving funds, converting them if necessary, and reconciling the transfer against an invoice. Where bank transfers already settle promptly and cheaply, introducing a token could add steps and costs.

The care evidence belongs in a controlled record system with an assisted alternative for beneficiaries who cannot use a digital interface. The payout needs a transaction reference and amount. Publishing visit times, location, images or health details alongside the payment would create needless exposure. The proposal can be evaluated without using biometric data. A signed attendance record establishes provenance, while accuracy still depends on how the observation was made and how corrections are handled.

4. What would count as improvement?

Start with a retrospective or shadow comparison in which proposed decisions are logged while existing payments continue. Independent reviewers should examine a stratified sample of routine, incomplete and disputed claims, then record the uncertainty of their judgments. The denominator for rejection rates should include claims that providers abandoned after discovering a documentation burden. Compare outcomes by provider size and access to assisted channels.

A prospective pilot can compare the existing approval process with a version that allows partial release and time-limited review, using the same settlement rail. A second comparison can hold the approval process fixed and vary the payment rail where participants can opt in and access equivalent alternatives. Outcomes include the share of supported claims paid within the service standard, the share of unsupported payments, provider days awaiting usable funds, time to resolve appeals, and recipient burden. Separate time spent verifying care from time spent settling an approved claim.

The model predicts that partial release will reduce the days of delay for verified portions of multi-visit invoices. Its value should be greatest when disputes are concentrated in a minority of visits. An assisted review path should recover more valid claims in settings with intermittent documentation or limited device access, with review costs tracked alongside recovered payments. If these effects do not appear, or provider participation falls enough to reduce access to care, the approval design has failed the proposed test.

5. Relationship to the earlier papers

When Do Programmable Vouchers Improve Public Spending? examines the effects of checking eligibility and use restrictions at checkout. A purchase can satisfy those conditions while the service it pays for has yet to be delivered. This paper studies the later decision about performance and asks how to allocate approval rights and payment delays among the funder, beneficiary and provider.

Programmable Money in Three Layers distinguishes value, purpose and evidence across settlement systems. The care model uses that architecture and adds specific decision rights, a time-limited review path, and a rule for releasing the supported portion of a claim. These are design propositions for future evaluation. The paper reports no pilot results.

Keywords

Stablecoin settlement, care payments, agent–purpose binding, conditional escrow, electronic visit verification, assisted claims review.

References

Centers for Medicare & Medicaid Services. (n.d.). Electronic Visit Verification. https://www.medicaid.gov/medicaid/home-community-based-services/home-community-based-services-guidance-additional-resources/electronic-visit-verification

Holmström, B. (1979). Moral hazard and observability. The Bell Journal of Economics, 10(1), 74–91. https://doi.org/10.2307/3003320

Shim, J. (2026a). Target Money (Chapter 8). Raon Book.

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

Shim, J. (2026c). Programmable Money in Three Layers. Zenodo. https://doi.org/10.5281/zenodo.24580931