The Tax of Distrust

Measuring the hidden costs of conditional payments

Jace Shim · November 2026 · https://doi.org/10.5281/zenodo.24581402

Abstract

A payment system can report lower fraud and smaller audit budgets while leaving merchants and recipients with more paperwork, more rejections, and longer waits. This working paper develops the tax of distrust from Chapters 1 and 2 of Target Money into a comparative accounting method. It tracks verification, unresolved losses, settlement charges, financing delays, and exclusion over the same population and time period. The measure distinguishes the costs borne by individual participants from the resource costs of the system as a whole. An illustrative subsidy program shows how shifting checks to the point of authorization can improve an agency's accounts while worsening the combined outcome. Stablecoin settlement is assessed as a distinct change to the payment rail. The paper provides a measurement design and reports no empirical estimates.

1. The cost that disappears from the agency's books

Suppose a local authority pays suppliers for eligible equipment. Its staff used to audit a sample of invoices after payment. Under a new system, every invoice is checked before money moves. The audit budget falls and fewer payments are flagged after settlement. Meanwhile, small suppliers spend more time correcting records, and a handful give up on valid claims. These costs can be missing from the same dashboard that records the savings.

Williamson (1979) treats governance as a response to the costs of contracting and adapting under uncertainty. Target Money applies that insight to digital payments, where convenience at the moment of transfer can coexist with considerable work before and after it. This paper asks a narrower question: what needs to be counted before calling a new control cheaper?

2. What is the tax of distrust?

The tax of distrust is the measurable burden of verifying, enforcing and settling an exchange when participants cannot take eligibility, performance or payment for granted. It is measured against a specified comparison and over a fixed population and period. It includes resources spent on routine checks, losses and delays left unresolved by those checks, and the burden imposed on legitimate participants.

The book identifies audit costs, fraud and leakage, intermediation charges and the cost of delayed settlement. This paper follows those categories and records two additions explicitly: the work of resolving exceptions and the value lost when an eligible exchange never takes place. A system that reduces fraud by refusing more claims can then be assessed on both outcomes.

3. Two accounts for the same program

Build an incidence account for each actor: the payer, recipient or supplier, payment provider, and any third-party reviewer. For each actor record staff and user time, unrecovered losses, net fees paid or received, the cost of delayed access to funds, and the location of any unserved eligible transaction. Keep cash outlays and estimates separate. A payment reversed after review should contribute its residual loss, recovery effort and delay only once.

Build a resource account for the program as a whole. Count labor, infrastructure, mediation and actual payment-processing resources once. Fees paid by one included actor to another are transfers in this account, even though they affect participation and bargaining. Record gross misdirection and recoveries as separate public-budget measures. Include any social loss from defeated program purpose, excluding the principal already counted as a transfer. If the cost of a provider's processing is unknown, report a range or leave the resource entry unestimated. A fee offers a useful private cost measure and only a rough proxy for the resources consumed.

Use the sum of daily unavailable balances to measure delay exposure. If an approved claim for 10,000 currency units becomes usable ten days later, the exposure is 100,000 currency-days. At a 12% annual marginal funding rate, the ten-day financing cost is about 33 units on a 365-day basis. Use the incremental balance and rate for the party that actually loses access, and avoid adding the same financing loss to another category. Report delays in days as well, since an estimated funding cost can miss consequences such as missed payroll or temporary supplier exit.

4. A break-even comparison

Consider two versions of the subsidy program for 1,000 eligible claims. All numbers here are hypothetical. Checking before release saves 6,000 units in public audit resources and prevents 4,000 units of estimated harm from misdirected funding. Operating the new checks and handling exceptions uses 4,500 units. Suppliers incur 2,000 units in additional documentation time and 1,000 units in incremental delay costs. The remaining benefit is 2,500 units before accounting for eligible claims that the process excludes.

If twenty additional legitimate claims are abandoned, the break-even value of each lost delivery is 125 units. At 150 units per lost delivery, the revised design reduces program value by 500 units. The comparison should report the break-even value and the number of abandoned claims, especially where a defensible social value for each delivery is unavailable. If the new design changes who applies at all, the denominator should remain the eligible population or be estimated from independent records.

5. Comparing payment rails and control design

A stablecoin can carry an authorized payment. It may shorten the transfer leg while adding conversion, custody, onboarding and reconciliation steps. Measure time until funds are usable, and compare an identical authorization policy across payment rails. Compare prepayment checks with postpayment review on the same rail to estimate what changes when the control moves. Use district or program rollout comparisons where feasible, and inspect case mix, supplier selection and pre-existing trends.

Collect timestamps, itemized operating costs, sampled independent reviews, appeals and follow-up with noncompleters. Apply the same audit standard in both systems. Otherwise, a lower count of detected fraud could reflect changes in detection. Set a common

observation window, follow unresolved claims through a stated cutoff and test sensitivity to recovery rates, wages, financing rates and the value of missed deliveries. Report uncertainty and cost incidence alongside the aggregate.

6. Relationship to earlier papers

When Do Programmable Vouchers Improve Public Spending? evaluates checkout enforcement for eligible households. This paper follows costs across institutions and includes supplier financing, fees and recovery after a claim is paid. Who Should Authorize Payment for Care? specifies decision rights when care records conflict, and When Should Freelancers Get Paid? measures exposure under milestone escrow. Those are contract designs that can be evaluated with the accounting method proposed here. Programmable Money in Three Layers describes the roles of value, purpose and evidence. This paper measures the burdens that arise when organizations implement those roles.

Keywords

Tax of distrust, transaction costs, programmable payments, stablecoin settlement, payment delays, audit costs, cost incidence.

References

Shim, J. (2026). Target Money, Chapters 1–2. Raon Book.

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

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

Shim, J. (2026). Who Should Authorize Payment for Care? Zenodo. https://doi.org/10.5281/zenodo.24581092

Shim, J. (2026). When Should Freelancers Get Paid? Zenodo. https://doi.org/10.5281/zenodo.24581270

Williamson, O. E. (1979). Transaction-cost economics. The Journal of Law and Economics, 22(2), 233–261. https://doi.org/10.1086/466942