

The Ontology of Agentic Economics: Money, Claims, and Obligation

Paper I of *Foundations of Agentic Economics*

Matthew Long
The YonedaAI Collaboration
YonedaAI Research Collective
Chicago, IL
`matthew@yonedaai.com`

August 2026

Abstract

Monetary theory has spent two centuries arguing about what money *is*, and the argument has been hard to settle because the competing answers explain different things. This paper proposes an ontology in which the disagreement partly dissolves. The primitive economic object is not money but the obligation: a directed, dated, conditional promise running from one party to another. An economy is then a time-dependent directed graph $\mathcal{E}_t = (V, E_t)$ whose vertices are humans, firms, governments, machines, and AI agents, and whose edges $e_{ij} = (q, u, t, c, r)$ carry a quantity, a unit of account, a maturity, a set of conditions, and a risk. Money appears late in this construction rather than early, as the standardized and highly transferable claim used to discharge whatever survives clearing. We state that idea as the *Residual Settlement Principle*, $M_t = \mathcal{R}(\mathcal{E}_t)$, and give it analytic content by proving four results about the clearing operator, which we model as a partition of obligations into netting groups. Gross compression is settlement-neutral. Coarsening the partition weakly reduces the settlement requirement, so multilateral netting dominates bilateral netting, which dominates gross settlement. No arrangement drives the requirement to zero unless the graph is balanced at every vertex, which is non-generic. And shortening the window over which obligations may be offset is a refinement, so it weakly increases the requirement, with instantaneous settlement requiring gross funding. Together these imply that the quantity of settlement money an economy needs is not a fact about its production but a joint fact about its obligations and its clearing institutions. Existing infrastructure demonstrates the effect at scale: CLS settles trillions of dollars daily against funding of well under one percent of that value. Autonomous agents matter here for two reasons, neither of which is speed: they lower the reconciliation cost that bounds how wide a netting perimeter can be, and cheap machine verification enlarges the set of contracts that can exist at all. We argue this against the strongest objections we know, including the charge that the formalism merely restates balance-sheet identities, and we state the conditions under which the proposal would be wrong. The framework is offered as a research program rather than a finished theory.

Keywords: monetary theory, credit theory of money, obligation networks, netting and clearing, settlement, autonomous agents, hierarchy of money

1 Introduction

Ask what money is and the usual reply is a list. Currency, bank deposits, sometimes short-dated government paper, occasionally money market fund shares. The list is a reasonable answer to a practical question, and it is what monetary statistics are built from, but it is not an answer to the theoretical one. It tells us which objects have been behaving like money in a particular jurisdiction at a particular time. It does not tell us what property those objects share, why the list changes, or what would be on it in an economy organized differently from ours.

The theoretical question becomes more pressing when the economy starts to reorganize. Software systems can now originate contracts, evaluate counterparties, hold positions, and instruct payments with limited human intervention, and researchers have begun to study what happens when they manage balance sheet positions in settlement systems, so far in experimental rather than live settings (Aldasoro and Desai, 2025). Suppose that trend continues. Suppose obligations can be created, priced, transferred, offset, and discharged continuously, by parties that never sleep and whose books reconcile in milliseconds. What, then, is money?

The answer this paper defends is that the question was misposed from the start, because money was never the primitive. The primitive is the obligation. Money is a derived object: the standardized, highly transferable claim that an economy uses to discharge whatever obligations its clearing arrangements could not cancel. On this view the quantity of money an economy requires is a residual, and a residual is exactly the kind of quantity that can change dramatically without anything real changing at all.

Axiom 1 (Money as claim). Money is a transferable claim, not a primitive economic object.

Stated that baldly the claim sounds either obvious or reckless, depending on the reader's training. To economists in the credit-theory tradition it will sound like a restatement of Innes (1913). To those trained on search-theoretic monetary models it will sound like an evasion, since the interesting question there is precisely why a valueless token circulates at all. Part of the work of this paper is to show that both reactions are partly right, and that the tradition has been arguing past itself because the rival accounts answer different questions. Menger's account explains how a medium of exchange could emerge among agents without institutions. Knapp's explains why a particular unit becomes the one everybody settles in. Innes's explains what most economic activity actually consists of once institutions exist. These are not competing answers to one question; they are answers to three.

What unifies them, we argue, is the hierarchical structure that Mehrling (2011, 2012) describes, in which one tier's money is the next tier's credit and the tiers are reconciled only at settlement. Hierarchy is the real-world anchor of the formal claim this paper makes. If moneyness is a graded and tier-relative property rather than a binary one, then the boundary

between money and credit is a matter of institutional design, and the amount of money an economy needs depends on where that boundary is drawn.

1.1 The main claim

We model an economy as a time-dependent directed graph

$$\mathcal{E}_t = (V, E_t), \quad (1)$$

where V contains humans, firms, governments, machines, and AI agents, and where an edge

$$e_{ij} = (q, u, t, c, r) \quad (2)$$

is an obligation running from i to j , characterized by quantity q , unit of account u , maturity t , conditions c , and risk r . The obligation is the primitive; the graph is the economy's promise structure at an instant.

Within this setting the paper's central proposal is what we call the Residual Settlement Principle:

$$\boxed{M_t = \mathcal{R}(\mathcal{E}_t)} \quad (3)$$

where $\mathcal{R}$ returns the residual obligation remaining after permissible clearing and netting, and M_t is the stock of standardized transferable settlement claims required to discharge that residual. Money, on this reading, is unresolved economic obligation. It has no independent economic existence outside the accounting relationships it represents.

The word doing the most work in (3) is *permissible*. Clearing is not a mathematical operation performed on a graph; it is a legal operation performed on contracts, and what may be offset against what is determined by netting agreements, insolvency law, membership rules, and the boundaries of clearing houses. Change the legal partition and the residual changes, with the real economy untouched. Section 6 makes this precise and proves the direction of the effect.

1.2 What the paper establishes, and what it does not

We prove four propositions about the clearing operator. The first says that compressing gross positions while holding net positions fixed does not change what has to be settled, which distinguishes portfolio compression from netting and explains why the two are measured in incomparable units. The second says that enlarging netting groups weakly reduces the settlement requirement, with the corollary that multilateral netting dominates bilateral netting, which dominates gross settlement. The third gives the floor: settlement cannot be eliminated unless every node's net position is zero, and the set of obligation graphs with that property has measure zero. The nettable proportion is therefore bounded strictly below one for generic graphs, which matters for how Paper IV of this series should interpret its empirical netting ratios.

The fourth is a constraint rather than an encouragement, and it cost us an earlier version

of this paper’s central claim about autonomy. Shortening the window over which obligations may be offset is a refinement of the partition, so by the second proposition it weakly increases the settlement requirement, rising to gross funding in the limit of instantaneous settlement. Speed and netting pull against each other, which is the classical tradeoff between real-time gross settlement and deferred net settlement (Freixas and Parigi, 1998; Kahn and Roberds, 2009). Any claim that autonomous agents reduce settlement demand by operating continuously is therefore false as stated, and Section 7.3 sets out the narrower claim that survives.

These are modest results. They are propositions about a definition, and a hostile reader could describe them as arithmetic. That is a fair description of their difficulty and an unfair description of their content, for a reason we take up at length in Section 8: the propositions locate the economically interesting variable, and it turns out not to be the obligation graph. It is the partition.

What the paper does not establish is that any of this is empirically correct as a theory of observed monetary aggregates. The Residual Settlement Principle is a claim about required settlement media. Observed money stocks exceed that requirement, for reasons we address in Section 6.8, and no part of the argument here shows that the gap behaves as the framework would predict. Nor do we test the agentic extension, because the phenomenon it describes barely exists yet. This is a proposal about how to set up the problem, offered in the hope that the setup is fertile. It is not settled science, and several of its components are conjectural in ways we flag as they arise.

1.3 Plan

Section 2 sets out the three traditions and argues that the hierarchy of money reconciles them. Section 3 constructs the obligation graph and states what it can and cannot represent. Section 4 develops the chain from production to settlement. Section 5 treats money as a graded property. Section 6 contains the formal results and the statement of the Residual Settlement Principle. Section 7 asks what autonomy changes. Section 8 takes on the objections we consider strongest. Section 9 states what would falsify the proposal, and Section 10 places the paper in the five-paper series of which it is the first.

Papers II through V of *Foundations of Agentic Economics* build on this one. Paper II takes the obligation as given and asks what production looks like when the producing entity is itself an autonomous agent. Paper III asks what happens to taxation when income stops being a well-defined accounting event. Paper IV specifies a protocol implementing identity, contracting, clearing, and settlement. Paper V constitutionalizes the resulting system. The dependency runs forward: everything downstream inherits the ontology constructed here.

2 Three traditions and the space between them

2.1 Money as an equilibrium of exchange

The account with the strongest formal pedigree begins with Menger (1892). Barter requires a double coincidence of wants. Some goods are more saleable than others, and an agent who

cannot find a direct match will accept a saleable good intending to trade it on. Saleability is self-reinforcing, so the process converges on one or a few commodities that everybody accepts. Money emerges from decentralized behavior with no legislation and no prior institution.

The modern descendants of this argument are considerable achievements. Kiyotaki and Wright (1989) embed the story in a search environment and derive which good becomes the medium of exchange as an equilibrium outcome rather than assuming it, and Lagos and Wright (2005) give a tractable framework in which monetary questions can be asked alongside ordinary general-equilibrium ones. Ostroy and Starr (1990) survey what the transactions role of money requires of an exchange process, and Hicks (1935) had already argued that the theory would only become serious once money was treated as an asset chosen alongside others. The tradition takes seriously a problem that credit theorists tend to skip: in a world of anonymous strangers with no enforcement, a promise is worthless, so something other than a promise has to change hands. That is the force of Hahn (1965) on the difficulty of getting money to have value at all, and of Clower (1967) on the cash-in-advance constraint.

Critics often claim that the historical record refutes this tradition. It does not, because the tradition is a rational reconstruction and its authors mostly say so. Menger was explaining how money *could* arise without design, and the answer he gave is coherent. The trouble is that the reconstruction has often been read as anthropology, and as anthropology it is in poor shape.

2.2 Money as circulating credit

Graeber (2011) assembled the case from the ethnographic and archaeological record, and the case is that the barter economy of the textbooks has never been observed. What the record shows instead is dense webs of mutual obligation, tracked informally or on tablets, settled periodically or not at all, long predating coinage. Credit came first. Coins came much later, in contexts that had a lot to do with paying soldiers.

The economics of this position had been laid out much earlier by Innes (1913, 1914), who argued that a sale is not the exchange of a good for a token but the exchange of a good for a credit, and that the buyer's obligation is discharged when the seller accepts it back in payment for something else. On Innes's account credit and credit alone is money, and the coin is simply a token of the state's debt. Merchants historically accumulated mutual claims and met at fairs to cancel them against each other, with only the residue changing hands. That practice, which Innes describes at length, is the direct historical ancestor of both the argument of this paper and the clearing houses of Section 6.

The credit tradition explains something the exchange tradition does not, which is what the overwhelming bulk of economic activity in a developed economy actually consists of. Trade credit, payables, receivables, deposits, loans, and derivative exposures dwarf currency by orders of magnitude. If money is the interesting object, almost nothing that happens in a modern economy is monetary.

2.3 Money as a creature of law

Knapp (1924) advanced a third claim, which is neither of the above: whatever the origin of

money, its *validity* as a means of payment is a legal fact. The state names the unit, declares what discharges a debt to it, and thereby determines what everybody else will accept, since everybody else owes taxes. Money is a creature of law.

Goodhart (1998) sharpened the contrast into two theories, which he called M and C, and argued that the chartalist one predicts the observed relationship between states and currencies considerably better. Currency areas coincide with sovereign boundaries far more tightly than transaction-cost logic would imply, and the exceptions tend to be cases where sovereignty is contested or delegated. The modern monetary theory literature (Wray, 2015; Kelton, 2020) develops the fiscal consequences: state spending creates claims and taxation extinguishes them, so the sovereign’s balance sheet is the terminal node of the system rather than one participant among many.

Chartalism has a weakness the credit theorists exploit, which is that it explains denomination and final settlement better than it explains the vast interior of the obligation network, most of which the state has no direct involvement in. And the credit theory has the symmetric weakness: it explains the interior beautifully and leaves unexplained why everything is denominated in one particular unit and why the apex claim is the one it is.

2.4 Hierarchy as the reconciliation

Mehrling (2011, 2012) supplies the structure that lets both be right. Money is hierarchical. At each level, what counts as money to the participants below is credit to the institution above. A depositor treats a bank deposit as money; the bank treats it as a liability discharged in central bank reserves. A dealer treats repo as near-money; the clearing bank does not. Central bank money sits at the apex domestically, and gold or reserve currency historically sat above that internationally.

Two consequences follow, and both are load-bearing for this paper. The first is that “is x money?” has no context-free answer, since the answer depends on which tier is asking. The second is that the system expands and contracts along the hierarchy: in calm periods the boundary between money and credit blurs and lower-tier claims trade at par with higher-tier ones, and in crises the hierarchy reasserts itself and everything below the apex is discovered to be credit after all. Gorton and Metrick (2012) document one such reassertion in detail. Pozsar (2014) maps the tiering of the shadow banking system in these terms.

Chartalism, in this light, is a theory about the apex. Credit theory is a theory about the interior. Search theory is a theory about what happens when there is no hierarchy at all, which is a real limiting case and not an idle one, since it describes the situation of any two parties with no shared institutions. The traditions are compatible once one stops treating them as rival answers to a single question.

2.5 Economies without a settlement asset

There is a fourth strand, smaller than the others and closer to this paper’s conclusion than to its premises, which asks what an economy would look like if the settlement asset shrank to nothing.

Wicksell (1936) posed the question first, in the form of a pure credit economy in which all payment runs through bank ledgers and no reserve medium circulates. His purpose was to show that the price level in such an economy is governed by the relation between the loan rate and the return on capital rather than by any quantity of money, since there is no quantity to speak of. Black (1970) and Fama (1980) revived the construction in a modern setting, treating banking as an accounting system in which portfolio management and payment services are separable and in which nothing forces a determinate nominal quantity to exist. Woodford (2003) works out monetary policy in the cashless limit and finds the interest rate instrument intact.

This literature is the natural home of the present paper, and its existence tempers any claim to novelty in the destination. What is new here is not the observation that an economy can function with vanishing settlement media. It is the account of what determines how little is enough, which the pure-credit tradition mostly leaves as a limiting assumption rather than deriving from the structure of obligations and the reach of clearing.

2.6 The informational hinge

One further result belongs to none of these traditions and matters more to the present argument than any of them. Kocherlakota (1998) proved that in a broad class of environments, money is equivalent to memory. Where agents cannot commit, and where no shared record of past behavior exists, a circulating object achieves exactly the allocations that a complete public record of everyone's history would achieve. Money substitutes for the record. Kocherlakota and Wallace (1998) extend the argument to partial record-keeping.

Read one way this is a defense of money. Read the other way it is a statement of the conditions under which money is unnecessary: if the record exists, is shared, and can be enforced against, the object is redundant. That framing puts the agentic question in exactly the right form. The interesting feature of autonomous agents is not that they are clever. It is that they operate on a substrate where record-keeping and verification cost almost nothing, which is precisely the parameter Kocherlakota's result turns on.

3 The obligation graph

3.1 Definition

Definition 1 (Obligation graph). An *obligation graph* at time t is a directed multigraph $\mathcal{E}_t = (V, E_t)$ in which V is a finite set of economic parties, including humans, firms, governments, machines, and AI agents, and each edge $e_{ij} \in E_t$ is an ordered relation from i to j carrying the tuple

$$e_{ij} = (q, u, t, c, r), \quad (4)$$

where $q \in \mathbb{R}_{>0}$ is the quantity owed, u the unit of account in which q is denominated, t the maturity at which performance is due, c the conditions under which performance is owed, and r a characterization of the risk that performance fails.

The tuple deserves component-by-component comment, since each field is carrying real weight and two of them are usually suppressed in the accounting treatments this construction resembles.

The quantity q and unit u are the least interesting fields and the most presupposed. Writing q as a scalar in unit u builds in the assumption that the obligation is denominated, which is to say that the pricing problem has already been solved somewhere outside the model. We return to this in Section 8, because it is a genuine limitation rather than a technicality.

The maturity t is what makes an obligation credit rather than a spot exchange. Note the notational collision with the time index of $\mathcal{E}_t$, which we retain because the five-tuple in (4) is the canonical form used throughout this series. Where confusion is possible we write $t(e)$ for the maturity of a particular edge and reserve the bare subscript for the graph's time index.

The conditions c are the field that most repays attention. An obligation is rarely unconditional. Performance may be owed only if goods arrive, only if a service meets specification, only if a reference rate exceeds a threshold, only if a counterparty has not defaulted elsewhere. Conditionality is what makes contracts economically useful and what makes them expensive, because someone must determine whether the condition holds. Townsend (1979) formalized the cost of that determination and showed how it shapes which contracts are worth writing. The set of feasible obligations is therefore bounded by verification technology, and that boundary moves.

The risk r is not a property of the contract but an estimate about it, held by someone, and different parties will hold different estimates for the same edge. Carrying r in the tuple is a modeling convenience that hides a genuine subjectivity. It is what makes the same nominal claim worth different amounts to different holders, and it is why claims trade at all.

3.2 The accounting identity

Every edge is an asset to its head and a liability to its tail. Writing $A_j^{(i)}$ for the claim j holds on i and $L_i^{(j)}$ for the corresponding liability,

$$A_j^{(i)} = L_i^{(j)}, \quad \sum_i A_i - \sum_i L_i = 0. \quad (5)$$

The second equality holds under conditions worth stating rather than assuming. It requires that the system be closed, so that no edge has an endpoint outside V ; that all obligations be denominated in a common unit, or else be converted by some valuation map; and that assets and liabilities be recorded at the same value on both sides, which in practice they are not, since the holder's estimate r enters the carrying value.

The multi-unit case is a real gap. If edges are denominated in different units, there is no canonical way to add them, and (5) holds only after applying an exchange rate vector that is itself a market outcome and can change. The identity is then an identity in each unit separately and an approximation in aggregate. We flag this rather than resolve it. It is the same difficulty that makes cross-currency netting hard in practice, which is why CLS exists at all.

3.3 The life of an edge

Edges are not static. The graph is time-dependent because obligations are created, transformed, and destroyed, and distinguishing the transformations is where the ontology does its work.

An edge is *originated* when a party incurs an obligation, typically in exchange for present performance. It may be *transferred*, meaning that the head of the edge changes while the tail does not: the creditor assigns the claim to somebody else. This is what a payment is. Paying with a bank deposit does not extinguish anything; it reassigns a claim on the bank from the payer to the payee, leaving the bank's liability intact and merely re-pointed. An edge may be *novated*, meaning that the tail changes too, which is the operation a central counterparty performs when it interposes itself. It may *mature*, at which point performance falls due if the conditions c are satisfied. It may be *netted* against opposing edges, which is the subject of Section 6. It may be *settled*, meaning genuinely extinguished, which removes it from the graph. And it may be *written down* or defaulted, which removes it without performance.

Payment and settlement are different operations and conflating them is the single most common source of confusion in this area. Payment moves a claim. Settlement destroys one. Almost all of what we call paying is the former, which is why the total stock of obligations in a developed economy grows more or less continuously without anybody finding this remarkable.

3.4 What the graph does not contain

The obligation graph is a ledger of promises. It is not a model of value, and it should not be mistaken for one.

It contains no preferences, no technology, no goods, and no prices except insofar as prices have already been embedded in q . It says nothing about whether the promised performance is worth having. A graph of obligations to deliver worthless output is formally identical to a graph of obligations to deliver valuable output, and only the risk field r registers the difference, indirectly and through somebody's judgment. Nothing in the construction implies that a large obligation graph indicates a productive economy; the 2007 vintage of the graph was very large.

It also omits obligations that are real but not economic in this sense, and the omission is not innocent. Duties of care, political allegiances, and the unpriced expectations that hold communities together are obligations in every sense except the one modeled here. Graeber's point in placing debt in a moral rather than financial frame is that the financial frame is a historically specific narrowing. We adopt that narrowing because it buys formal tractability, and we note that it costs something.

Finally, the graph as constructed is a snapshot of promises rather than of capacity. Paper II of this series argues that capacity is the more fundamental quantity and that obligations are underwritten by it. The two objects are related but distinct, and this paper treats capacity as exogenous.

4 From production to settlement

The ontology can be stated as a chain in which each arrow is a genuine transformation rather than a relabeling:

$$\text{Production} \rightarrow \text{Claim} \rightarrow \text{Credit} \rightarrow \text{Money} \rightarrow \text{Settlement}. \quad (6)$$

Production to claim. Something is made or done. Whoever did it holds a claim on the result, either directly, by owning the output, or contractually, by holding a promise from whoever commissioned it. The claim is the economic shadow of the production event. This is the sense in which value is productive capacity: what a claim is ultimately a claim *on* is somebody’s ability to produce.

Claim to credit. A claim whose performance is deferred is credit. The deferral is what introduces both time and risk, and therefore what introduces the possibility of pricing. A spot exchange creates no credit because the claim is discharged at the moment it arises. Almost no exchange in a developed economy is spot in this sense, since even a card payment leaves a chain of interbank claims outstanding for a day or more.

Credit to money. Credit becomes money by acquiring three properties: standardization, so that units are fungible and no due diligence is required per unit; transferability, so that the claim can be reassigned without the debtor’s consent; and low conditionality, so that the conditions c collapse toward the empty set and acceptance requires no judgment about contingencies. A claim with all three is money for whoever will take it. This is the content of the redefinition

$$\boxed{M = \text{transferable settlement claim}} \quad (7)$$

replacing the notion of money as a scarce object.

Money to settlement. Settlement is the cancellation of obligation. Money is the instrument through which cancellation occurs, and it works because accepting the standardized claim is treated as equivalent to performance. At the apex of the hierarchy, settlement in central bank money is final by legal construction rather than by economic argument, which is Knapp’s point restated in modern infrastructure terms (Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions, 2012; Norman et al., 2011).

The chain reads in one direction only, which is the substantive claim. Conventional intuition runs money $\rightarrow$ transaction $\rightarrow$ production: you need money to buy things, and buying things induces production. The ordering proposed here is

$$\text{credible future production} \rightarrow \text{credit} \rightarrow \text{transaction} \rightarrow \text{settlement}, \quad (8)$$

in which the ability to produce supports an obligation, the obligation supports a transaction, and settlement cleans up afterward. McLeay et al. (2014) document that this is how bank

lending actually works, against the loanable-funds intuition: the loan creates the deposit rather than the deposit funding the loan. What Paper II adds is that the same inversion applies to any agent whose future production can be underwritten, not only to licensed banks.

5 Moneyiness as a graded property

If money is a kind of claim rather than a kind of object, then the question of whether something is money admits of degree. We write the moneyiness of a claim x as

$$\mathcal{M}(x) = f(\text{acceptability, liquidity, credit quality, transferability, settlement finality, stability}). \quad (9)$$

A bank deposit scores highly on all six. A personal IOU scores low on most of them: it is acceptable to almost nobody, illiquid, of uncertain credit quality, transferable only with effort, final only when the debtor performs, and stable only in nominal terms. Treasury bills, repo, commercial paper, money market fund shares, and trade credit occupy intermediate positions, and they do not lie on a single line. There is no fundamental boundary between money and credit in this picture, only different grades of claim.

Two honest qualifications about (9) are needed before it is put to work.

The first is that writing f as a scalar-valued function overstates what is known. The six arguments are not commensurable, and there is no defensible exchange rate between, say, settlement finality and stability. What the construction genuinely delivers is a partial order: x is more money-like than y if it weakly dominates on every component. Many pairs are incomparable, and the incomparability is economically real. A Treasury bill dominates a deposit on credit quality and is dominated on acceptability, which is why both exist and why the repo market exists to convert between them. We use the scalar form as a notational convenience and do not rely on it for any result.

The second is that $\mathcal{M}$ is state-dependent, and violently so. The six components are not properties of the instrument but assessments of it, and assessments move together. In calm conditions the moneyiness ordering compresses, lower-tier claims trade at par with higher-tier ones, and market participants economize on the genuinely money-like assets. In stress the ordering decompresses discontinuously. Gorton and Metrick (2012) trace the mechanism for securitized banking in 2007–2008: haircuts on repo collateral rose sharply, which is precisely the observation that instruments previously treated as cash-equivalent stopped being so. Holmström and Tirole (1998) give the theoretical account of why private claims cannot in general supply all the liquidity an economy wants in bad states.

This matters for the argument of the next section, because the clearing operator $\mathcal{R}$ depends on which claims participants will accept in discharge. If moneyiness collapses in stress, so does the reach of clearing, and the residual grows at exactly the moment it is hardest to fund. Any claim that netting reduces the need for settlement assets has to be read as a claim about normal conditions.

Remark 1. The hierarchy of Mehrling (2012) and the grading (9) are different objects that are easy to conflate. The hierarchy is a vertical ordering of *issuers*, in which each tier settles in the

liabilities of the tier above. Moneyiness is a graded property of *instruments*, which correlates with tier but is not determined by it. A highly rated corporate's commercial paper and a weak bank's deposit sit at different tiers and may have similar moneyiness.

6 Clearing, netting, and the residual

6.1 The elementary case

Take three parties and three obligations, all in the same unit and maturing on the same date:

$$A \rightarrow B : 100, \quad B \rightarrow C : 80, \quad C \rightarrow A : 70. \quad (10)$$

Gross obligations total 250. But nothing requires these to be discharged one at a time. Computing each party's net position gives

$$\nu_A = -30, \quad \nu_B = +20, \quad \nu_C = +10, \quad (11)$$

which sum to zero, and the entire structure can be discharged by A transferring 30 to be divided between B and C . Settlement capacity of 30 suffices for 250 of obligations.

This is the whole idea, and it is very old. Innes describes medieval merchants doing it at fairs; the London bankers' clearing house institutionalized it in the eighteenth century (Norman et al., 2011). What has changed over time is not the concept but the reach: how many parties can be included, how frequently the operation runs, and which legal claims may be included in it.

6.2 Netting sets and the settlement requirement

The reach is the object to formalize. What may be netted against what is a legal question, determined by netting agreements, clearing house membership, and the enforceability of close-out netting in the relevant insolvency regime (International Swaps and Derivatives Association, 2010).

The natural first attempt is to partition the participants, placing parties who may offset against one another in the same block. That attempt fails, and the reason is instructive. Bilateral netting is conducted pair by pair across the whole network, so a party with mutual obligations to two others nets separately with each. No partition of V represents this: putting all three in one block imposes multilateral netting, and separating them forces at least one pair to settle gross. Netting perimeters overlap, and partitions of participants cannot overlap.

We therefore partition the obligations rather than the parties. This costs nothing and represents every arrangement of interest, including bilateral netting on arbitrary graphs and the fragmented multi-venue arrangements of Section 6.6.

Definition 2 (Netting arrangement). A *netting arrangement* on $\mathcal{E}_t$ is a partition Π of the edge set E_t into *netting groups*. Obligations lying in the same group $g \in \Pi$ may be legally offset against one another; obligations in different groups may not.

Definition 3 (Net position and settlement requirement). For a netting group $g \in \Pi$ and a party $i \in V$, write $\text{head}(e)$ and $\text{tail}(e)$ for the creditor and debtor of edge e . The net position of i within g is

$$\nu_i(g) = \sum_{\substack{e \in g: \\ \text{head}(e)=i}} q(e) - \sum_{\substack{e \in g: \\ \text{tail}(e)=i}} q(e), \quad (12)$$

and, writing $(x)^- = \max(0, -x)$ for the negative part, the *settlement requirement* under Π is

$$S(\mathcal{E}, \Pi) = \sum_{g \in \Pi} \sum_{i \in V} (\nu_i(g))^- . \quad (13)$$

Summing over edges rather than over ordered pairs keeps (12) well posed on a multigraph, where a party may hold several distinct obligations against the same counterparty. Within any group the net positions sum to zero, so (13) counts each group's one-sided funding need once. Three arrangements recur: Π_{gross} , in which every edge is its own group and nothing offsets; Π_{bi} , in which edges are grouped by the unordered pair of parties they connect; and $\Pi_{\text{multi}} = \{E_t\}$, the single multilateral group. Note that $S(\mathcal{E}, \Pi_{\text{gross}}) = \sum_e q(e) = X$, since an isolated edge contributes its full amount through its debtor.

6.3 Units and dates

Definition 3 adds quantities, which requires them to be commensurable. The edge tuple (4) permits arbitrary units and maturities, and dollars due tomorrow do not offset against euros due in a decade. Left unaddressed this would leave $\mathcal{R}$ undefined on the very object the paper takes as primitive.

We resolve it by slicing. For a unit u and a settlement window τ , let $\mathcal{E}_t^{u,\tau}$ denote the subgraph of obligations denominated in u and falling due within τ , and let $\Pi|_{u,\tau}$ be the induced arrangement. The settlement requirement is then a vector indexed by unit and window,

$$S(\mathcal{E}_t, \Pi) = \left(S(\mathcal{E}_t^{u,\tau}, \Pi|_{u,\tau}) \right)_{u,\tau}, \quad (14)$$

and all results below are stated within a slice. This is not a technical evasion but a description of how settlement actually works: participants in a multi-currency system fund a position per currency per value date, which is exactly why payment-versus-payment mechanisms had to be built rather than assumed. Aggregating (14) into a scalar requires an exchange rate vector and a discount curve, both of which are market outcomes lying outside the framework. We return to the consequences in Section 8.

6.4 Three results

The first result separates two operations that are routinely reported in the same sentence and measured in incompatible units.

Proposition 1 (Gross compression is settlement-neutral). *Let $\mathcal{E}'$ be obtained from $\mathcal{E}$ by any*

transformation that preserves $\nu_i(g)$ for every party i and every group $g \in \Pi$. Then

$$S(\mathcal{E}', \Pi) = S(\mathcal{E}, \Pi), \quad (15)$$

however much the gross outstanding $\sum_{e \in E} q(e)$ has changed.

Proof. Immediate from (13), which depends on $\mathcal{E}$ only through the within-group net positions that the hypothesis fixes. $\square$

The content is in the reading. Portfolio compression, in which redundant offsetting trades are torn up and replaced with fewer trades having the same net risk, reduces gross notional outstanding by very large amounts. Compression services have removed a cumulative notional well in excess of a quadrillion dollars from interest rate swap markets since 2008. Proposition 1 says that this achievement, real as it is for operational risk and for capital charges computed on notional, does not by itself reduce what must be funded at settlement. Compression ratios and netting ratios are not comparable quantities, and adding or averaging them is a category error.

The second result is the one that makes the arrangement, rather than the graph, the economically interesting variable.

Proposition 2 (Coarsening weakly reduces the settlement requirement). *Let Π' coarsen Π , meaning that every group of Π is contained in some group of Π' . Then*

$$S(\mathcal{E}, \Pi') \leq S(\mathcal{E}, \Pi). \quad (16)$$

Proof. Fix $g' \in \Pi'$ and write $g' = \bigsqcup_k g_k$ with $g_k \in \Pi$. Net positions are sums over edges, so they are additive over a disjoint decomposition of the edge set: $\nu_i(g') = \sum_k \nu_i(g_k)$ for every party i . The map $x \mapsto (x)^- = \max(0, -x)$ is subadditive, since $\max(0, -(x+y)) \leq \max(0, -x) + \max(0, -y)$, whence

$$(\nu_i(g'))^- \leq \sum_k (\nu_i(g_k))^- . \quad (17)$$

Summing over $i \in V$ and over $g' \in \Pi'$, and using that the groups g_k enumerate Π exactly once, gives $S(\mathcal{E}, \Pi') \leq S(\mathcal{E}, \Pi)$. $\square$

Corollary 3 (Ordering of settlement regimes). *With Π_{gross} , Π_{bi} , and Π_{multi} as defined above,*

$$S(\mathcal{E}, \Pi_{\text{multi}}) \leq S(\mathcal{E}, \Pi_{\text{bi}}) \leq S(\mathcal{E}, \Pi_{\text{gross}}) = \sum_{e \in E} q(e) = X, \quad (18)$$

where X denotes gross exchange.

Proof. Π_{multi} coarsens Π_{bi} , which coarsens Π_{gross} , in the sense required by Proposition 2; the final equality was noted after Definition 3. $\square$

Grouping edges rather than parties is what makes Corollary 3 true in general. Under a partition of participants, bilateral netting on a network where one party faces several

counterparties is not representable at all, and the ordering would hold only for graphs whose mutually offsetting pairs happen to be disjoint.

Corollary 3 justifies defining the nettable proportion used in Paper IV of this series as a property of the arrangement rather than of the economy. Setting

$$n(\Pi) = 1 - \frac{S(\mathcal{E}, \Pi)}{X} \quad \text{so that} \quad S = X(1 - n), \quad (19)$$

the relation $S = X(1 - n)$ becomes a definition of n rather than an empirical law, and n inherits a monotonicity in the coarseness of Π . Reported netting ratios are therefore statements about particular legal and operational perimeters, and averaging them across venues is not meaningful.

The third result gives the floor.

Proposition 4 (Irreducibility of the residual). *Under the coarsest arrangement $\Pi_{\text{multi}} = \{E\}$, writing $\nu_i = \nu_i(E)$,*

$$S(\mathcal{E}, \Pi_{\text{multi}}) = \sum_{i \in V} (\nu_i)^- = \frac{1}{2} \sum_{i \in V} |\nu_i|, \quad (20)$$

which vanishes if and only if $\nu_i = 0$ for every $i \in V$. Moreover, if the graph is connected, the set of edge-weight vectors $q \in \mathbb{R}^{|E|}$ satisfying $\nu_i = 0$ for all i is a linear subspace of codimension $|V| - 1$, and therefore has Lebesgue measure zero in $\mathbb{R}^{|E|}$ whenever $|E| > |V| - 1$.

Proof. The first display follows from $\sum_i \nu_i = 0$, which makes the positive and negative parts equal in magnitude. For the genericity claim, let $D \in \mathbb{R}^{|V| \times |E|}$ be the incidence matrix of the directed graph, with $D_{ie} = +1$ if $\text{head}(e) = i$, $D_{ie} = -1$ if $\text{tail}(e) = i$, and 0 otherwise. Then (12) reads $\nu = Dq$, so the balanced weight vectors are exactly $\ker D$. The incidence matrix of a connected directed graph on $|V|$ vertices has rank $|V| - 1$, the single deficiency corresponding to the left null vector $\mathbf{1}^\top D = 0$, which is the identity $\sum_i \nu_i = 0$. Hence $\ker D$ has codimension $|V| - 1$ in $\mathbb{R}^{|E|}$, and a proper linear subspace has measure zero. $\square$

Proposition 4 is the formal statement of a fact practitioners know: you cannot net your way to zero. Perfectly balanced obligation graphs exist but are exceptional, and any real graph will leave somebody short and somebody long. The nettable proportion n is bounded strictly below one for generic graphs, which is consistent with the empirical record, where even the most efficient venues report ratios approaching but not reaching unity.

We should be careful about how much this genericity argument is asked to carry. It says that balanced graphs are exceptional among all weight vectors, under a uniform notion of “exceptional” that no economy realizes. Real payment flows are concentrated, seasonal, and strongly reciprocal, which is why clearing houses achieve the ratios they do. Proposition 4 establishes that the floor is positive; it says nothing about whether the floor is near or far, and the empirical evidence is that for well-designed perimeters it is very near.

6.5 The temporal direction runs the other way

A netting group cannot contain obligations that do not yet exist. Every arrangement therefore has a temporal extent, the window over which obligations accumulate before the group is closed

and its residual funded. Making the window shorter refines the arrangement, and Proposition 2 then runs against the intuition that faster is better.

Proposition 5 (Temporal refinement weakly increases the settlement requirement). *Let a settlement window τ be subdivided into consecutive sub-windows $\tau = \bigsqcup_k \tau_k$, and let Π^τ and $\Pi^{\{\tau_k\}}$ be the arrangements that net within the whole window and within each sub-window respectively, holding the obligation graph fixed. Then*

$$S(\mathcal{E}, \Pi^\tau) \leq S(\mathcal{E}, \Pi^{\{\tau_k\}}). \quad (21)$$

If in addition the maturities $t(e)$ are pairwise distinct, then as the mesh of the subdivision goes to zero, $S \rightarrow X$.

Proof. Subdividing the window partitions each netting group by the sub-window in which its edges fall, which is a refinement in the sense of Proposition 2; apply that proposition. For the limit, pairwise distinct maturities are separated by some positive minimum gap, so any subdivision of smaller mesh isolates each edge in its own group, giving Π_{gross} and hence $S = X$ by Definition 3. $\square$

The distinctness condition is not decorative. Obligations that mature at the same instant are never separated by any subdivision, so they continue to offset however fine the partition becomes, and the limit is then the settlement requirement of the arrangement that groups simultaneous obligations rather than X . Synchronized settlement is a real design choice for exactly this reason: batching obligations to a common instant preserves offsetting that would otherwise be destroyed by sequencing. Instantaneous settlement of a stream of unsynchronized obligations, by contrast, forfeits netting entirely.

This is the central tension of payment system design and it has been understood for a long time. Real-time gross settlement extinguishes obligations as they arise, which removes intraday credit exposure between participants and requires funding close to gross. Deferred net settlement waits, which economizes on funding and leaves participants exposed to one another in the interval. Freixas and Parigi (1998) analyze the tradeoff directly and find it genuinely two-sided; Bech and Garratt (2003) model the resulting timing game, in which participants delay payments to economize on intraday liquidity and can collectively produce gridlock; Kahn and Roberds (2009) survey the field.

Proposition 5 therefore places a real constraint on any claim that faster clearing is cheaper clearing, and it constrains the argument of Section 7 in particular. Speed and netting pull in opposite directions. Any account of what autonomous agents do to settlement demand has to say which of the two effects it is invoking, and an account that invokes both at once without distinguishing them is incoherent. Section 7.3 takes up the resolution, which is that continuous operation and gross settlement are separable, and that the infrastructure separating them already exists.

6.6 Netting is not monotone in the number of clearing houses

Proposition 2 is often misread as saying that more clearing is better. It says something narrower: coarsening a partition helps. Adding a clearing house that carves a new perimeter out of an existing one is a *refinement*, and refinements can hurt.

Duffie and Zhu (2011) established this for central counterparties and derivatives. Introducing a CCP for one asset class removes that class from the bilateral netting sets in which it was previously offset against other classes, and for plausible parameterizations the fragmentation raises average exposure to counterparty default rather than lowering it. Clearing distinct classes in separate CCPs always increases exposures relative to clearing the combined set in one. The mechanism is exactly the one in Proposition 2, run backwards.

Appendix B works a five-party example in which the same obligations require 65 of settlement under a single netting perimeter and 295 under two perimeters, a factor of 4.5, with no change whatsoever to the underlying economic relationships. The number is an artifact of the example, but the direction is not.

We take this to be the strongest available evidence for the paper’s central contention. The settlement requirement is not a property of the economy. It is a property of the economy together with its clearing institutions, and the institutional term can dominate.

6.7 Relation to the clearing-vector literature

The operator $\mathcal{R}$ has an obvious ancestor in Eisenberg and Noe (2001), who studied a network of interbank obligations and proved by a fixed-point argument that a clearing payment vector always exists and is generically unique. Rogers and Veraart (2013) extend the construction to default costs, and Acemoglu et al. (2015) study how network structure governs contagion.

The relationship is one of lineage rather than of identity, and the difference should be stated plainly. The Eisenberg–Noe clearing vector solves an allocation problem under default: given limited resources and pro-rata priority, how much does each node actually pay? The operator $\mathcal{R}$ in this paper solves a prior and easier problem: assuming performance, what must be transferred? The two coincide when nobody defaults and diverge otherwise. A complete treatment would compose them, applying $\mathcal{R}$ to determine settlement obligations and then an Eisenberg–Noe fixed point to resolve the residual under shortfall. We do not attempt that composition here, and we regard it as the most tractable open problem the framework generates.

6.8 The Residual Settlement Principle

Principle 1 (Residual Settlement). Let Π_t be the netting arrangement admissible at time t and let $\mathcal{R}$ map an obligation graph to the residual obligation surviving clearing under Π_t . Then

$$\boxed{M_t = \mathcal{R}(\mathcal{E}_t)} \tag{22}$$

where M_t denotes the standardized transferable settlement claims required to discharge the residual obligation graph at time t . Slice by slice, the required quantity of the settlement asset

of unit u for window τ is

$$|M_t^{u,\tau}| = S(\mathcal{E}_t^{u,\tau}, \Pi_t|_{u,\tau}). \quad (23)$$

Remark 2 (The residual and the asset that discharges it). Two distinct objects stand on the two sides of (22), and the compact notation invites conflating them. The residual $\mathcal{R}(\mathcal{E}_t)$ is a structure of surviving liabilities: party i still owes party j . The settlement claims M_t are assets of a quite different kind, typically liabilities of a bank or a sovereign standing high in the hierarchy, with moneyiness close to the maximum of (9). A residual obligation of a defaulting counterparty has very low moneyiness and could not discharge anything.

What (22) asserts is a determination, not an identity of instrument: the residual fixes how much of the settlement asset is required, slice by slice, as (23) states. The sentence “money is unresolved economic obligation” should be read as saying that the demand for the settlement asset is the shadow of what clearing failed to cancel, not that the two are the same thing. Keeping them separate is what makes the hierarchy of Mehrling (2012) indispensable to the framework rather than decorative.

The equality in (22) is one of referent rather than of type: $\mathcal{R}(\mathcal{E}_t)$ is a residual obligation structure, and M_t is what is required to extinguish it. The reduction to magnitudes is the settlement requirement of Definition 3, taken slice by slice as in (14). Read as a sentence: money is unresolved economic obligation.

Before drawing consequences we should be exact about what kind of statement this is, because a reader who takes it for an empirical law will find it circular and be right to. Principle 1 has two components and they have different standing. The first is definitional. Having defined the settlement requirement as what survives clearing, identifying $|M_t|$ with $S(\mathcal{E}_t, \Pi_t)$ adds no information, and no observation could disturb it. The second is a substantive conjecture: that the demand for apex settlement assets tracks $S(\mathcal{E}_t, \Pi_t)$, and in particular responds to changes in Π_t that leave the real economy untouched. That conjecture can fail, and Section 9 says how.

It can also fail because S is not the only determinant of settlement asset demand, and the framework as it stands omits the others. Participants hold reserves against uncertainty about the timing of their own inflows, not only against known net positions. They respond to the opportunity cost of holding reserves against alternatives. They face regulatory requirements and central bank operating frameworks that set floors unrelated to netting. None of these appear in (13), and a demand function built from the principle alone would be misspecified. What the principle claims is that S is a determinant and a manipulable one, not that it is the only one.

With that stated, three consequences follow, and one objection has to be met immediately.

The consequences first. If (22) holds, then the quantity of money an economy requires is jointly determined by its promise structure and its clearing institutions, and the second term is manipulable. Two economies with identical production, identical preferences, and identical obligations can require settlement balances differing by an order of magnitude, purely because one has a wider netting perimeter. Second, money is endogenous in a stronger sense than the usual one: not merely created by bank lending (McLeay et al., 2014; Tobin, 1963) but determined in quantity by an institutional variable that no monetary authority directly

controls. Third, one dimension along which economic systems differ can be written as

$$\text{Economic efficiency} \approx \frac{\text{real productive exchange}}{\text{settlement capital required}}, \quad (24)$$

a ratio that improves when clearing improves without any change in production.

The label on the left of (24) is broader than the quantity on the right, and we use it only because it is the form in which the idea is stated elsewhere in this series. What the ratio measures is settlement-liquidity efficiency: how little settlement capital an economy ties up to support a given volume of real exchange. It is silent about allocative efficiency, and it is silent about the concentration of credit risk that a low denominator usually implies. Section 8 argues that both silences matter enough that (24) should not be treated as an objective function.

Now the objection. Observed monetary aggregates are far larger than any plausible settlement requirement, because money is held for reasons other than imminent settlement, chiefly as a store of value and as a precautionary buffer. If M_t is defined as required settlement media, then (22) is not a theory of $M1$ or $M2$ and should not be tested against them naively.

The framework’s answer is that the store-of-value holdings are not outside the graph. A deposit held as savings is an edge: a claim on a bank, sitting in E_t like any other. A Treasury bill held as a buffer is an edge. What these holdings represent is not settlement media but positions in the obligation network, and the ontology absorbs them without modification. What sits outside the graph is only the apex settlement asset, and it is that asset which (22) is about. This is why the hierarchy of Mehrling (2012) matters structurally and not merely rhetorically: the residual principle applies at the apex, and everything below the apex is graph.

We regard this answer as adequate but not complete. It rescues the principle from an obvious falsification, and it does so by narrowing the claim. The narrowed claim is testable, which Section 9 takes up.

6.9 Evidence from existing infrastructure

The effect is not hypothetical. Wholesale settlement infrastructure already achieves compressions of the kind Proposition 2 describes, and the magnitudes are large.

CLS, which settles foreign exchange on a payment-versus-payment basis, reports that multilateral netting reduces participants’ gross payment obligations by approximately 96%, and that with liquidity-saving mechanisms the funding actually required falls by roughly 99%. On its record day of 20 June 2024 it settled \$19.1 trillion against funding of about \$72 billion, an implied $n \approx 0.996$ (CLS Group, 2025). CHIPS, the large-value dollar system, reports liquidity efficiency on the order of 26:1, meaning that a dollar of intraday funding supports around twenty-six dollars of settled value, an implied $n \approx 0.96$ (The Clearing House, 2025). In derivatives, the BIS reports that gross credit exposure after legally enforceable close-out netting stood 86.4% below gross market value at end-June 2025 (Bank for International Settlements, 2025), consistent with the ISDA finding that netting reduces OTC credit exposure by more than 85% (International Swaps and Derivatives Association, 2010).

Three cautions apply to these figures. They are venue-specific and measure different things, so they must not be averaged or treated as estimates of a common parameter; the CLS and

CHIPS numbers concern payment funding while the BIS number concerns credit exposure. They are reported by the operators of the infrastructure in question, and while we have no reason to doubt them, they have not been independently audited for the purposes of this paper. And they describe normal conditions, which is the qualification Section 5 attached to any statement about the reach of clearing.

With those cautions, the qualitative point survives. Foreign exchange turnover did not fall when CLS was introduced; the funding required to settle it did. That is an instance of X holding while S collapsed, which is what (22) says should be possible and what a theory treating money as a primitive has no natural way to express.

7 What autonomy changes

Nothing said so far requires artificial intelligence. The obligation graph describes a medieval fair as well as it describes CLS, and the three propositions are indifferent to who computes the netting. The question this section addresses is what, if anything, changes when the parties in V include entities that act without human intervention.

7.1 Five layers

It helps to separate the functions that a clearing economy performs, because autonomy affects them unequally.

The *capability layer* records what each party can produce. The *obligation layer* is the graph $\mathcal{E}_t$, holding the promises made in exchange for production. The *trust layer* holds the estimates r that obligations will be honored. The *clearing layer* continuously offsets reciprocal obligations. The *settlement layer* discharges what clearing could not, using assets at the top of the hierarchy.

Autonomy leaves the first and last layers largely alone. What an economy can produce is a question about technology and resources, and what constitutes final settlement is a question about law. The middle three are where machine agency bites, and the largest effect is on clearing.

7.2 Three constraints on the netting perimeter

Proposition 2 says that widening the netting perimeter reduces the settlement requirement. If widening is beneficial, the question is why perimeters are not already wide. Three constraints bind, and they are of different kinds.

The first is legal. Close-out netting is enforceable only where insolvency law permits it, and cross-border enforceability is uneven (International Swaps and Derivatives Association, 2010). This constraint is not technological, and no amount of computation relaxes it. It is the reason that the partition Π_t is genuinely exogenous to the parties and why Paper IV of this series treats jurisdictional interoperability as a protocol-level concern rather than an engineering detail.

The second is operational: reconciliation. Netting requires agreement about what is owed, and establishing agreement across institutional boundaries has historically been expensive and slow. This constraint is technological and it has been eroding for decades.

The third is temporal. Netting runs in cycles, and cycles exist because reconciliation and funding have historically been batch operations aligned to human working hours and to the operating windows of settlement systems. Between cycles, exposures accumulate.

Autonomous agents relax the second constraint and leave the first untouched. Machine reconciliation reduces the cost of establishing what is owed, which widens the perimeter in extent, and Proposition 2 then delivers a reduction in S . The legal constraint remains, and it is the binding one in the long run.

The third constraint is where care is required, because the obvious claim about it is false. It is tempting to say that agents also shorten the cycle toward zero and that this widens the perimeter in frequency. Proposition 5 shows the opposite: shortening the window is a refinement, so it weakly *increases* the settlement requirement, and in the limit of instantaneous settlement the requirement rises to gross. Faster is not cheaper. An earlier formulation of the argument in this paper asserted that continuous operation reduces S , and that assertion does not survive the proposition.

7.3 Continuous operation without gross settlement

The resolution is that continuous operation and gross settlement are separable, and conflating them is what produces the false claim.

What Proposition 5 penalizes is shortening the interval over which obligations may be *offset*. It says nothing against shortening the interval at which the offsetting algorithm *runs*. A system can hold a wide netting group open while continuously searching it for offsetting cycles, releasing obligations as soon as they can be cancelled against others and leaving only what cannot. Payments are then discharged throughout the day, but the funding required is closer to the net than to the gross, because cancellation rather than payment is doing most of the work.

This is not hypothetical. It is what liquidity-saving mechanisms in large-value payment systems do, and it is the reason the CHIPS figures in Section 6 look as they do. CHIPS settles payments with finality within seconds, which is continuous operation, and yet supports something on the order of twenty-six dollars of settled value per dollar of intraday funding (The Clearing House, 2025), which is a netting ratio unavailable under true gross settlement. The queue holds obligations briefly, the algorithm searches for offsets, and the residual is funded from a small prefunded balance. Continuous and net are not opposites.

The claim about agents that survives is therefore narrower than the enthusiast's version and narrower than the one this paper made a paragraph ago. Autonomous agents reduce reconciliation cost, which widens perimeters in extent, and they can run offsetting searches continuously rather than in batches, which captures netting benefits without imposing a batch delay on discharge. Neither effect requires shortening the offsetting window, and the framework predicts that a system which did shorten it would need more settlement assets, not fewer. Bech and Garratt (2003) model the timing incentives that make this design problem hard, and Freixas and Parigi (1998) the risk that the retained exposure creates.

7.4 Conditions, verification, and the set of feasible contracts

The deeper change is not to clearing but to the conditions field c , and it is the part of the argument we consider most novel.

An obligation conditional on a state of the world is only worth writing if the state can be determined at a cost small relative to the value at stake. Townsend (1979) showed how verification cost shapes the optimal contract, and Hart and Moore (1988) showed how unverifiability forces contracts to remain incomplete and pushes the resulting disputes into renegotiation and governance. The population of contracts that exist in an economy is therefore a function of what can be cheaply checked.

Continuous machine verification moves that boundary. Conditions that can be evaluated against sensor data, delivery confirmations, code execution, or attested credentials (World Wide Web Consortium, 2022; Szabo, 1997) become writable at scales where the administrative cost would previously have exceeded the contract value. The consequence is that E_t acquires edges that could not previously exist, in kinds and in numbers.

The mechanism admits a compact statement, which is worth making because the verbal version invites the suspicion that nothing has been said. Let a candidate obligation e carry condition $c(e)$, let $\sigma(e) > 0$ be the surplus the parties expect from writing it, and let $\kappa(c)$ be the cost of verifying c to the standard required for enforcement. The contract is worth writing only if it does not consume its own surplus:

$$E_t(\kappa) = \{ e : \kappa(c(e)) \leq \sigma(e) \}. \quad (25)$$

Proposition 6 (Verification cost and the feasible edge set). *If $\kappa' \leq \kappa$ pointwise, then $E_t(\kappa) \subseteq E_t(\kappa')$. The feasible obligation set is monotone decreasing in verification cost, and strictly so whenever some candidate obligation satisfies $\kappa'(c(e)) \leq \sigma(e) < \kappa(c(e))$.*

Proof. Immediate from (25): the defining inequality is weakened pointwise. $\square$

The proposition is elementary, and we claim nothing for its difficulty. Its role is to locate the mechanism precisely enough to be argued with. What falls out of it is that the contracts gained when κ drops are exactly those of low surplus and high conditionality, which are the contracts an economy of many small machine-mediated transactions would consist of. A serious treatment would endogenize σ and κ in the manner of Townsend (1979) and ask what the optimal verification standard is rather than taking it as given, and would have to confront the point of Hart and Moore (1988) that some conditions are not merely expensive to verify but not verifiable at all by any third party, so that no reduction in κ reaches them. We do not attempt that here, and (25) should be read as a statement of what would need modeling rather than as a model.

This is the crux of the reply to the objection that the graph formalism merely restates balance-sheet accounting, and the two mechanisms need separating. Cheaper verification does not describe existing contracts more precisely. It changes which contracts are feasible. A framework organized around obligations with an explicit conditionality field can express that

change; a framework organized around money cannot, because the change does not show up as a change in any monetary quantity until it has already happened.

7.5 Obligations against expected production

The second consequence of autonomy is that credit can be extended against production that has not occurred. An agent reliably generating output should not have to accumulate a cash balance before acquiring inputs, since its expected production is itself underwriting. The credit limit becomes

$$\text{Credit Limit}_i = f(\text{expected future production, variance, reputation, dependencies}), \quad (26)$$

and acceptance of a newly originated obligation becomes probabilistic rather than automatic.

This is endogenous money creation at machine speed, and it is not new in kind. Banks have done exactly this for centuries, which is the burden of McLeay et al. (2014) and, at greater length, of Minsky (1986): private balance sheet expansion creates purchasing power, and the constraint is acceptance rather than prior saving. Anyone may issue an obligation; the difficulty is getting somebody to take it. What is new is who may do it. Stiglitz and Weiss (1981) showed why credit is rationed rather than cleared by price when lenders cannot observe borrower quality; if agent behavior is substantially more observable than human behavior, the informational premise of that result weakens for a class of borrowers.

One assumption carries the weight here and should be named rather than smuggled. Underwriting against future production requires that future production be *pledgeable*: that a claim on output not yet produced can be established, monitored, and enforced. Pledgeability is not free and it is not general. It needs the output to be verifiable in the sense of Section 7, and it needs recourse that survives the obligor's failure, which for a going concern means the ability to seize or redirect the productive capacity rather than merely to obtain a judgment against an entity whose assets are its own future activity. Where those conditions fail, capacity is not collateral, and lending falls back on assets that can be repossessed.

The distinction matters because an influential part of the collateral literature runs the other way, deriving borrowing constraints precisely from the fact that human capital and future output *cannot* be pledged, which is what forces borrowers onto durable assets like land. Nothing in this paper establishes that agentic production escapes that problem.

The empirical record is friendlier to the present argument than that tradition would suggest. Lian and Ma (2021) find that for United States non-financial firms the dominant form of borrowing constraint is cash-flow-based, with obligations sized against earnings and going-concern value, while lending against liquidation value of durable assets is the smaller category. Underwriting against expected production is therefore the normal case in the largest corporate credit market rather than an exotic one, which weakens the objection that capacity-based credit is a theoretical curiosity. It does not dissolve the objection, because what makes cash-flow lending work is an enforcement apparatus built around going concerns that agentic issuers may or may not fit.

What this paper claims is accordingly narrow: verification and enforcement technology

determine where the pledgeability boundary sits, and the boundary is not fixed. Paper II of this series takes up the question directly, develops the balance sheet $B_i(t) = (A_i, L_i, C_i, R_i, P_i)$ and the economic capacity function $C_i(t)$, and carries the burden of defending pledgeability for agentic producers.

The ontological point that belongs here is narrower still: (26) is an edge-admission rule. It governs which obligations may enter E_t , and it is therefore part of the specification of the graph rather than a separate theory bolted onto it. Whether the rule is generous or stingy is a question about pledgeability, and this paper takes no position on the answer.

7.6 Scale, and a figure we decline to give

It is tempting at this point to quantify: to say how many obligations per second an agentic economy would generate, or by what factor the graph would grow. We decline, because we have no defensible number. Published work on agent economies is at the stage of framework and simulation rather than measurement (Hadfield and Koh, 2025; Tomasev et al., 2025; Yang et al., 2025), and the one piece of infrastructure-level empirical work we are aware of concerns a single agent managing liquidity buffers in a payment system (Aldasoro and Desai, 2025). Extrapolating from that to economy-wide transaction volumes would be invention.

What can be said without invention is directional and follows from the propositions, with the signs taken carefully. If autonomy widens offsetting perimeters, S falls relative to X by Proposition 2. If it also enlarges the feasible contract set by Proposition 6, X itself rises, which lowers the ratio S/X again provided the new obligations are no less nettable than the old. Both effects push the level of settlement assets required away from any fixed relation to the volume of real activity. But if autonomy is used to shorten offsetting windows rather than to widen them, Proposition 5 runs the other way and S rises. The net direction is a design question, not a technological inevitability, and whether the effect is large enough to matter is an empirical question that cannot yet be answered.

7.7 Where the state remains

None of this dissolves the sovereign. Two functions survive intact and one is strengthened.

The state deserves a name in the model, being a participant rather than an environment. Write $G \in V$ for the public vertex: the state is a party in the obligation graph like any other, issuing claims when it spends and extinguishing them when it taxes. What distinguishes G is not that it sits outside $\mathcal{E}_t$ but that its liabilities sit at the top of the hierarchy. Paper III of this series develops the financing of G and uses the same symbol.

Settlement finality at the apex remains a legal construction. Central bank money is final because the law says so, and no amount of netting produces an asset with that property from below (Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions, 2012; Goodhart, 1998). Taxation remains a demand for a specific liability: whatever the state accepts in payment of taxes acquires a floor of acceptability that no private claim can replicate, which is Knapp's point and the operational core of the modern monetary theory literature (Knapp, 1924; Wray, 2015).

In graph terms, taxation makes every other vertex a debtor to G in a unit only G can create, which is why the apex claim is the one it is.

The function that is strengthened is monetary policy, which we take up in Section 8. Paper III of this series asks what happens to the tax base when income stops being a well-defined accounting event, and Paper V asks what constitutional constraints the resulting arrangement requires.

8 Objections

The proposal deserves serious opposition, and this section states the objections we consider strongest in their strongest form. Several of them we concede.

8.1 The formalism merely restates balance-sheet accounting

This is the objection we expect to be raised first and the one most worth answering carefully. Every asset is somebody's liability; the sectoral balances sum to zero; netting is what clearing houses have always done. Drawing a graph and calling the edges primitives adds notation, not content.

The objection has real force against the first half of the paper. Sections 3 and 4 are, in their formal content, a restatement of double-entry accounting with two extra fields. We do not claim otherwise. The answer has three parts, and only the third is decisive.

First, an identity constrains without determining. That assets equal liabilities tells you nothing about how much settlement is required, because the settlement requirement depends on the partition and the identity holds for every partition. Proposition 2 is not derivable from the identity, and neither is the Duffie–Zhu fragmentation result. The economically interesting variable is invisible to the accounting framing precisely because accounting is partition-independent by construction.

Second, the framing determines what is treated as endogenous. In a money-primitive framework, the money stock is the object of interest and clearing arrangements are plumbing. In an obligation-primitive framework, the clearing arrangement is a policy variable of the first order, and the money stock is downstream of it. That is a substantive difference in research program even where the accounting is identical, and it generates different questions. Nobody asks what the optimal netting perimeter is if the perimeter is plumbing.

Third, and most importantly: continuous machine-verifiable conditions change what is feasible, not only how it is described. A contract conditional on a state that costs more to verify than the contract is worth does not exist. When verification cost falls below that threshold the contract comes into existence, which is Proposition 6, and the population of E_t changes. No restatement of an accounting identity has that consequence, because identities hold over whatever population happens to exist and are silent about which populations are possible. The obligation framework has a field for conditionality; the balance sheet does not.

The third answer has a limit that should be stated with it. Verification cost is not the only thing keeping contracts from existing. Hart and Moore (1988) locate contractual incompleteness in states that no third party can verify at any price, and for those states the cost of verification

is not high but undefined, so no reduction in κ reaches them. The mechanism therefore expands the feasible set within the class of contracts whose conditions are in principle checkable, and leaves the genuinely incomplete ones where they were. How large that class is relative to the whole is not something we can establish.

We regard the third answer as sound within that limit and the first two as suggestive. A reader who accepts only the third is entitled to conclude that the contribution of this paper is narrower than its framing suggests, and we would not argue strenuously.

8.2 The origin story favors commodity money

If money emerged from barter through the saleability mechanism (Menger, 1892), then money is prior to credit, and the ontology proposed here inverts the actual order.

The reply is that origin and ontology are different questions, and that even a fully successful Mengerian account would not establish what money *is* now. Historical priority does not confer analytic priority: fire was used long before oxidation was understood, and the order of discovery constrained nothing about the order of explanation. The Mengerian argument is at its strongest as an account of what happens among parties with no shared institutions, which is a real situation and one the credit theory handles poorly.

There is a further point that cuts against the objection on its own ground. The anthropological record, as far as it goes, does not support the barter stage as a historical phase (Graeber, 2011). The credit relationships appear first in the documentary evidence, and the coinage arrives late and in state-military contexts. This is contested terrain and Graeber's synthesis has been criticized on particulars, so we do not rest weight on it. But the objection cannot claim historical support that the record does not supply.

8.3 Netting relocates risk rather than removing it

This objection is correct and we concede it substantially.

Multilateral netting reduces the funding required at settlement. It does not reduce the underlying credit exposure so much as concentrate it, replacing many small bilateral exposures with fewer and larger exposures to a central node. Freixas and Parigi (1998) analyze the tradeoff between gross and net systems directly and find it genuinely two-sided: net systems economize on liquidity and are more exposed to contagion. Duffie and Zhu (2011) show that the netting benefit itself is not monotone in the number of clearing venues. Acemoglu et al. (2015) show that denser networks are stabilizing for small shocks and destabilizing for large ones, which means the risk consequences of widening a perimeter depend on the shock distribution. The regulatory apparatus around financial market infrastructure exists because of exactly this concern (Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions, 2012).

Nothing in Principle 1 prices this. $S(\mathcal{E}, \Pi)$ measures the settlement requirement and is silent about the risk cost of achieving a low one. A complete framework would carry a second functional measuring concentration and would exhibit the tradeoff explicitly, and we do not have one. This is a real gap rather than an oversight, and it qualifies the efficiency ratio (24):

an economy that minimizes settlement capital is not thereby well designed.

The failure mode has a name and a date. Herstatt Bank was closed in June 1974 during the settlement day, after receiving Deutschmarks and before paying dollars, and the resulting principal losses to counterparties are the reason payment-versus-payment mechanisms such as CLS were eventually built (Norman et al., 2011). The infrastructure that produces the impressive netting ratios of Section 6 was constructed in response to a risk that netting alone did not address.

8.4 The unit of account is presupposed

The edge tuple (4) contains a unit u , and nothing in the framework explains where u comes from. Obligations are denominated before the model starts. If money is supposed to be derived rather than primitive, having the unit of account handed to the construction from outside is awkward.

We concede this and regard it as the most serious internal weakness of the proposal. The graph cannot generate its own numéraire.

Two distinct problems hide behind the objection and they have different severity. The first is technical: with several units in play, the settlement requirement of Definition 3 adds incommensurable quantities. Section 6.3 handles this by slicing the graph by unit and settlement window and stating the results within a slice, so that M_t is a vector of per-unit, per-date requirements rather than a scalar. That is faithful to practice, since participants in a multi-currency system fund per currency per value date, and it leaves the propositions intact. What it costs is the ability to speak of *the* money stock without first choosing an exchange rate vector, and that cost is real: comparisons across slices are not available inside the framework.

The second problem is the deep one and slicing does not touch it. Nothing explains why any particular u is used, or why the number of units in circulation is small. The chartalist tradition has an answer, and it is one this paper is otherwise happy to accept: the sovereign names the unit (Knapp, 1924; Goodhart, 1998). Taking that answer means the ontology is not self-contained, since it rests on a political fact, but a theory of money resting on a political fact is in respectable company.

One observation softens the concession without removing it. The unit-of-account function and the settlement function are analytically distinct and can be carried by different objects, as they were under commodity standards and as they are today in economies that price in dollars and settle locally. The framework needs a unit for denomination and an asset for settlement, and it is only the first that it cannot supply. Article VI of the constitutional framework in Paper V takes a normative position here, holding that no particular representation of value is intrinsically the primitive. That position does not resolve the theoretical gap.

8.5 Netting economizes on buffers, which may be destabilizing

A framework that treats settlement capital as a cost to be minimized invites a Minskyan objection. Minsky (1986) argued that stability is destabilizing: periods of calm lead participants to economize on the margins of safety, and the resulting balance sheet structures are fragile in

ways that only appear when conditions change.

The efficiency ratio (24) is exactly an instruction to economize on margins of safety. Read as a normative target, it is an argument for holding as little settlement capital as clearing arrangements allow, which is the behavior Minsky identified as the mechanism of financial fragility. Combined with the state-dependence of moneyiness discussed in Section 5, the concern sharpens: the reach of clearing contracts in stress at the same time as the residual grows, so an economy optimized for the normal-state value of S will be short of settlement assets in precisely the states where it needs them.

We accept this and offer (24) as a descriptive ratio rather than an objective function. Its role is to make visible a dimension along which economies differ and which money-primitive accounting does not display. Turning it into a target without a companion measure of fragility would be a mistake, and we say so rather than qualifying it into vagueness.

8.6 If money is residual, is monetary policy impotent?

If the money stock is determined by the promise structure and the clearing partition, it might seem that a central bank controls nothing of consequence.

The implication runs the other way. What a central bank supplies is the apex settlement asset, and Principle 1 says that this asset is what the entire graph ultimately resolves into. Setting the terms on which the residual can be discharged is therefore a lever on every edge in E_t , not merely on a monetary aggregate. This is closer to the modern practice of monetary policy, which operates on the price of settlement balances rather than their quantity, than the money-primitive framing is. Woodford (2003) develops the theory of policy in a limiting case where the quantity of settlement media goes to zero and the interest rate remains the operative instrument, which is a formal precedent for the situation this framework describes as the destination of improved clearing.

There is a real consequence, though, and it is not reassuring. If the netting partition is a first-order determinant of settlement demand, and if that partition is set by private arrangements, clearing house membership rules, and cross-border insolvency law, then a material part of monetary conditions is determined by parties who are not monetary authorities and who are optimizing something else. That is a governance problem rather than an impotence problem, and it is one of the reasons Paper IV of this series treats the clearing perimeter as a matter for protocol design rather than private negotiation.

8.7 Autonomous agents are speculative

The final objection is that the agentic premise is science fiction and the paper would be sounder without it.

Partly conceded. The strongest version of the agentic economy, in which autonomous entities hold balance sheets and contract at scale without human involvement, does not exist. The literature is largely framework and simulation (Hadfield and Koh, 2025; Tomasev et al., 2025; Yang et al., 2025), and Section 7 declines to quantify for that reason.

What is not speculative is the direction of the two mechanisms. Reconciliation cost has

been falling for decades and settlement cycles have been shortening, both without any agentic premise at all. Verification cost has been falling wherever conditions can be attached to machine-readable evidence. Aldasoro and Desai (2025) study a language model agent managing liquidity buffers and settlement in a real-time gross settlement environment, which is an experimental result rather than a deployment and should be read as one, but it is a concrete instance of an agent taking balance sheet decisions of the kind at issue. The framework’s dependence on autonomy is weaker than its framing suggests: the propositions hold regardless, and autonomy is one route by which the constraints they identify get relaxed.

9 What would falsify this

Conceptual papers rarely state their own defeat conditions, which makes them easy to write and hard to use. We state four, in decreasing order of decisiveness.

Invariance of settlement balances to clearing infrastructure. Principle 1 implies that changing the netting perimeter changes the demand for settlement assets while leaving real activity unchanged. If the introduction of major netting infrastructure had left participants’ settlement balances and intraday funding requirements unchanged, the principle would be wrong in its central claim. The CLS and CHIPS figures cited in Section 6 run the right way, but they are operator statistics rather than a controlled comparison, and a serious test would exploit the staggered introduction of netting arrangements across venues and jurisdictions. We have not conducted that test and consider it the most valuable empirical work the framework calls for.

Failure of coarsening to reduce settlement. Proposition 2 is a theorem under its assumptions, so a genuine empirical failure would indicate that the assumptions are the wrong idealization rather than that the arithmetic is wrong. The most likely culprit would be the single-unit, single-maturity restriction. If widening perimeters in practice systematically failed to reduce funding needs because of maturity mismatch or currency mismatch, the formal apparatus would need rebuilding on a multi-unit, multi-date basis, and the qualitative conclusions might not survive.

Binary rather than graded moneyness. Section 5 asserts a continuum. If instruments priced as though they were either money or not money, with no intermediate grades and no state-dependent movement between them, the grading would be descriptively idle. Haircut schedules and the cross-sectional pricing of near-money assets are the natural place to look, and the crisis evidence (Gorton and Metrick, 2012) currently favors the graded reading.

Absence of new contracts as verification cost falls. Section 7 claims that cheaper verification enlarges the feasible contract set, and Proposition 6 states the mechanism. This is the least tested claim in the paper. If verification costs fall substantially in a domain and the population of conditional contracts does not change in kind, only in volume, then the

mechanism we identify as the decisive answer to the accounting objection does not operate, and the objection stands. The prediction is specific enough to check: what should appear first are contracts of low surplus and high conditionality, which were previously priced out by the fixed cost of determining whether their conditions held.

Settlement demand falling as cycles shorten. Proposition 5 predicts the opposite: compressing settlement cycles, holding the offsetting perimeter fixed, should raise funding requirements. Securities markets have moved from longer cycles to $T+1$ in several jurisdictions, and the framework predicts that such moves raise intraday funding needs unless accompanied by wider or more continuous offsetting. If shortened cycles were systematically accompanied by lower funding requirements with no change in offsetting arrangements, Proposition 5 would be the wrong idealization of how netting windows work.

10 Position in the series

This paper supplies the ontology on which four others depend.

Paper II, *The Agentic Production Economy*, takes the obligation as primitive and asks what production looks like when the producing entity is an autonomous agent. It replaces the firm-centric production function with $\mathcal{P} = H + A + K + E + I$ over human capability, autonomous intelligence, physical capital, energy, and information, and develops the agent balance sheet and economic capacity function that underwrite the credit limit (26). Where this paper treats capacity as exogenous, Paper II makes it the object of study. The theory of the firm it engages is the Coasean one (Coase, 1937), on the argument that agentic negotiation drives the cost of using the price mechanism toward zero.

Paper III, *Post-Income Taxation*, asks what happens to public finance when income ceases to be a well-defined accounting event. If most economic activity consists of creating and netting obligations, and if settlement is a residual, then the moment at which income is realized becomes an artifact of clearing arrangements rather than a fact about production. Paper III argues that the taxable event should move to final consumption and to scarce stocks, which connects to the Diamond–Mirrlees production efficiency result (Diamond and Mirrlees, 1971), and it tests the resulting revenue arithmetic against real fiscal data.

Paper IV, *The Agentic Economic Protocol*, implements the layers of Section 7 as a specification: verifiable identity, machine-readable contracts, continuous clearing, and settlement. Its central quantitative object is the relation $S = X(1 - n)$, which Corollary 3 here shows to be a definition of n in terms of the netting partition rather than an empirical regularity. The empirical values of n that Paper IV reports should accordingly be read as measurements of particular perimeters.

Paper V, *The Global Agentic Economic Constitution*, constitutionalizes the system. Its Article VI, monetary neutrality, is the normative counterpart of Axiom 1: if no representation of value is intrinsically primitive, then no particular monetary implementation should be constitutionally entrenched. The theoretical gap identified in Section 8 concerning the unit of account is not resolved by that article, and we flag the residue for future work.

11 Conclusion

The question this paper began with was what money is when obligations can be created, exchanged, netted, and settled continuously by parties that are not human. The answer offered is that money is what is left over.

The obligation is the primitive. An economy is a graph of directed, dated, conditional promises, and money is the standardized transferable claim used to discharge whatever survives clearing. Written compactly, $M_t = \mathcal{R}(\mathcal{E}_t)$. The propositions of Section 6 give that statement analytic teeth: compressing gross positions changes nothing about settlement, widening the netting perimeter reduces it, and no perimeter eliminates it, because balanced obligation graphs are exceptional. The consequence is that the settlement money an economy requires is a joint fact about its promises and its clearing institutions, and the institutional term is large enough to dominate. Existing infrastructure already demonstrates the magnitude.

Autonomous agents enter the argument in a specific and limited way. They relax the reconciliation and cycle constraints on the netting perimeter, leaving the legal constraint untouched, and they enlarge the set of contracts that can exist by making conditions cheap to verify. The second of these is the more consequential and the less tested.

Much is unresolved. The framework has no account of where the unit of account comes from and borrows one from the chartalists. It measures settlement efficiency without measuring the concentration risk that efficient clearing creates, which means its central ratio should not be used as an objective. Its treatment of default is a placeholder where the clearing-vector literature has a solution. And the agentic premise on which its framing rests is, as of this writing, mostly prospective.

What the framework offers in exchange is a change of subject. If money is primitive, the questions are about its quantity, its velocity, and its control. If obligation is primitive, the questions are about who may promise what to whom, which promises may be cancelled against which, and who decides. Those seem to us the questions an economy of autonomous agents will actually pose.

A Notation

The letter E carries three jobs across this series, and readers moving between papers should be warned. Here it is the edge set E_t of the obligation graph. In Paper II it is also energy, as the fourth argument of $\mathcal{P} = H + A + K + E + I$, and economic power E_i . The overload is inherited from the source material rather than introduced, and context separates the three cleanly enough that renaming would cost more in cross-paper consistency than it would buy in local clarity.

B A five-party worked example

The example makes Propositions 2 and 4 concrete and exhibits the fragmentation effect of Duffie and Zhu (2011) in miniature. All obligations are denominated in the same unit and fall

Symbol	Meaning
$\mathcal{E}_t = (V, E_t)$	Obligation graph at time t
V	Parties: humans, firms, governments, machines, AI agents
$e_{ij} = (q, u, t, c, r)$	Obligation from i to j : quantity, unit, maturity, conditions, risk
$q(e)$	Quantity owed under edge e
$t(e)$	Maturity of edge e , where the bare subscript would be ambiguous
$A_j^{(i)}, L_i^{(j)}$	Claim held by j on i ; corresponding liability of i
$\mathcal{M}(x)$	Moneyness of claim x
Π	Netting arrangement: a partition of the edge set E_t
$g \in \Pi$	A netting group
$\nu_i(g)$	Net position of party i within group g
$(x)^- = \max(0, -x)$	Negative part
$S(\mathcal{E}, \Pi)$	Settlement requirement under arrangement Π
$\mathcal{E}_t^{u, \tau}$	Slice of the graph in unit u due within window τ
$\kappa(c), \sigma(e)$	Verification cost of condition c ; surplus from obligation e
$X = \sum_e q(e)$	Gross exchange
$n(\Pi) = 1 - S/X$	Nettable proportion under Π
$\mathcal{R}$	Residual clearing operator
$M_t = \mathcal{R}(\mathcal{E}_t)$	Residual Settlement Principle
$G \in V$	The public vertex: the state as a party in the graph
$\mathcal{P} = H + A + K + E + I$	Production function (Paper II)
$B_i(t) = (A_i, L_i, C_i, R_i, P_i)$	Agent balance sheet (Paper II)

Table 1: Notation used in this paper and its relation to the series.

due on the same date.

B.1 The obligations

Edge	Amount	Edge	Amount
e_{12}	100	e_{45}	60
e_{23}	80	e_{52}	30
e_{31}	70	e_{35}	45
e_{14}	40	e_{53}	15
e_{41}	25	e_{24}	50
Gross exchange X			515

Table 2: Obligations among five parties.

B.2 Gross settlement

Under Π_{gross} , every party is its own block and nothing offsets. Every edge settles at its full amount, so

$$S(\mathcal{E}, \Pi_{\text{gross}}) = X = 515. \quad (27)$$

B.3 Bilateral netting

Under Π_{bi} , edges are grouped by the unordered pair of parties they connect, so opposing obligations between the same two parties offset. Only two pairs have obligations in both directions: (1, 4), where 40 against 25 nets to 15 owed by 1, and (3, 5), where 45 against 15 nets to 30 owed by 3. Summing the eight resulting one-way amounts,

$$S(\mathcal{E}, \Pi_{\text{bi}}) = 100 + 80 + 70 + 15 + 60 + 30 + 30 + 50 = 435. \quad (28)$$

Bilateral netting removes 80 of the 515, a reduction of 15.5%. Most obligation pairs in the example are one-directional, which is typical, and bilateral netting has little to work with.

B.4 Multilateral netting

Under $\Pi_{\text{multi}} = \{V\}$ all obligations offset against one another. Party by party:

$$\nu_1 = (70 + 25) - (100 + 40) = 95 - 140 = -45, \quad (29)$$

$$\nu_2 = (100 + 30) - (80 + 50) = 130 - 130 = 0, \quad (30)$$

$$\nu_3 = (80 + 15) - (70 + 45) = 95 - 115 = -20, \quad (31)$$

$$\nu_4 = (40 + 50) - (25 + 60) = 90 - 85 = +5, \quad (32)$$

$$\nu_5 = (60 + 45) - (30 + 15) = 105 - 45 = +60. \quad (33)$$

These sum to zero, as Proposition 4 requires. Parties 1 and 3 are short; parties 4 and 5 are long; party 2 is flat despite being involved in 260 of gross obligations. The settlement requirement is the sum of the short positions:

$$S(\mathcal{E}, \Pi_{\text{multi}}) = 45 + 20 = 65, \quad (34)$$

so $n = 1 - 65/515 = 0.874$. Multilateral netting removes 87.4% of gross obligations against bilateral netting's 15.5%, on identical underlying relationships.

B.5 Fragmentation into two perimeters

Now suppose parties $\{1, 2, 3\}$ clear in one venue and $\{4, 5\}$ in another, with cross-venue obligations settling gross. The corresponding arrangement groups the edges internal to each venue together and leaves each cross-venue edge in its own group, which refines Π_{multi} , so Proposition 2 predicts a weakly larger requirement.

Within $\{1, 2, 3\}$ the internal edges are $e_{12} = 100$, $e_{23} = 80$, and $e_{31} = 70$, giving $\nu_1 = -30$, $\nu_2 = +20$, $\nu_3 = +10$, so this venue requires 30. Within $\{4, 5\}$ the only internal edge is $e_{45} = 60$, requiring 60. The cross-venue edges $e_{14} = 40$, $e_{41} = 25$, $e_{52} = 30$, $e_{35} = 45$, $e_{53} = 15$, and $e_{24} = 50$ settle gross, totalling 205. Hence

$$S(\mathcal{E}, \Pi_{\text{frag}}) = 30 + 60 + 205 = 295. \quad (35)$$

Splitting one clearing perimeter into two raises the settlement requirement from 65 to

295, a factor of 4.5, with no change to any obligation. Permitting bilateral netting across the venue boundary recovers some of the loss: the cross-venue edges are then grouped by counterparty pair, giving 15 for (1, 4), 30 for (5, 2), 30 for (3, 5), and 50 for (2, 4), so the cross-venue component falls to 125 and the total to 215. That is still more than three times the single-perimeter figure. This last arrangement is one that a partition of the participants could not express, which is why Definition 2 partitions obligations instead.

Arrangement	S	$n = 1 - S/X$
Gross settlement	515	0.000
Bilateral netting throughout	435	0.155
Two perimeters, gross across the boundary	295	0.427
Two perimeters, bilateral across the boundary	215	0.583
Single multilateral perimeter	65	0.874

Table 3: Settlement requirement under five arrangements, identical obligations.

Table 3 is the paper’s argument in one place. The obligations are the same in every row. The economy is the same in every row. The settlement requirement varies by a factor of eight. Whatever determines that variation is not a fact about production, and a theory that treats money as primitive has nowhere to put it.

References

- Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *The American Economic Review*, 105(2):564–608, 2015.
- Iñaki Aldasoro and Ajit Desai. AI agents for cash management in payment systems. BIS Working Papers 1310, Bank for International Settlements, 2025.
- Bank for International Settlements. Otc derivatives statistics at end-june 2025. Technical report, Bank for International Settlements, Basel, 2025. BIS Statistical Release, November 2025.
- Morten L. Bech and Rod Garratt. The intraday liquidity management game. *Journal of Economic Theory*, 109(2):198–219, 2003.
- Fischer Black. Banking and interest rates in a world without money: The effects of uncontrolled banking. *Journal of Bank Research*, 1:8–20, 1970.
- Robert W. Clower. A reconsideration of the microfoundations of monetary theory. *Western Economic Journal*, 6(1):1–8, 1967.
- CLS Group. Annual report 2024. Technical report, CLS Group Holdings AG, Zurich, 2025.
- R. H. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.

- Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions. Principles for financial market infrastructures. Technical report, Bank for International Settlements, Basel, 2012.
- Peter A. Diamond and James A. Mirrlees. Optimal taxation and public production i: Production efficiency. *The American Economic Review*, 61(1):8–27, 1971.
- Darrell Duffie and Haoxiang Zhu. Does a central clearing counterparty reduce counterparty risk? *The Review of Asset Pricing Studies*, 1(1):74–95, 2011.
- Larry Eisenberg and Thomas H. Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001.
- Eugene F. Fama. Banking in the theory of finance. *Journal of Monetary Economics*, 6(1):39–57, 1980.
- Xavier Freixas and Bruno Parigi. Contagion and efficiency in gross and net interbank payment systems. *Journal of Financial Intermediation*, 7(1):3–31, 1998.
- Charles A. E. Goodhart. The two concepts of money: Implications for the analysis of optimal currency areas. *European Journal of Political Economy*, 14(3):407–432, 1998.
- Gary Gorton and Andrew Metrick. Securitized banking and the run on repo. *Journal of Financial Economics*, 104(3):425–451, 2012.
- David Graeber. *Debt: The First 5,000 Years*. Melville House, Brooklyn, NY, 2011.
- Gillian K. Hadfield and Andrew Koh. An economy of AI agents. In *The Economics of Transformative AI*. University of Chicago Press for the National Bureau of Economic Research, Chicago, IL, 2025. Also available as arXiv:2509.01063.
- Frank H. Hahn. On some problems of proving the existence of an equilibrium in a monetary economy. In Frank H. Hahn and Frank P. R. Brechling, editors, *The Theory of Interest Rates*. Macmillan, London, 1965.
- Oliver Hart and John Moore. Incomplete contracts and renegotiation. *Econometrica*, 56(4):755–785, 1988.
- John R. Hicks. A suggestion for simplifying the theory of money. *Economica*, 2(5):1–19, 1935.
- Bengt Holmström and Jean Tirole. Private and public supply of liquidity. *Journal of Political Economy*, 106(1):1–40, 1998.
- A. Mitchell Innes. What is money? *Banking Law Journal*, 30(5):377–408, May 1913.
- A. Mitchell Innes. The credit theory of money. *Banking Law Journal*, 31:151–168, 1914.
- International Swaps and Derivatives Association. The importance of close-out netting. ISDA Research Notes 1, ISDA, 2010.

- Charles M. Kahn and William Roberds. Why pay? an introduction to payments economics. *Journal of Financial Intermediation*, 18(1):1–23, 2009.
- Stephanie Kelton. *The Deficit Myth: Modern Monetary Theory and the Birth of the People’s Economy*. PublicAffairs, New York, 2020.
- Nobuhiro Kiyotaki and Randall Wright. On money as a medium of exchange. *Journal of Political Economy*, 97(4):927–954, 1989.
- Georg Friedrich Knapp. *The State Theory of Money*. Macmillan, London, 1924. Abridged English translation by H. M. Lucas and J. Bonar of *Staatliche Theorie des Geldes*, Duncker & Humblot, Leipzig, 1905.
- Narayana R. Kocherlakota. Money is memory. *Journal of Economic Theory*, 81(2):232–251, 1998.
- Narayana R. Kocherlakota and Neil Wallace. Incomplete record-keeping and optimal payment arrangements. *Journal of Economic Theory*, 81(2):272–289, 1998.
- Ricardo Lagos and Randall Wright. A unified framework for monetary theory and policy analysis. *Journal of Political Economy*, 113(3):463–484, 2005.
- Chen Lian and Yueran Ma. Anatomy of corporate borrowing constraints. *The Quarterly Journal of Economics*, 136(1):229–291, 2021. doi: 10.1093/qje/qjaa030. Bibliographic record verified against the Crossref REST API, which reports published-online 2020-09-24 and published-print 2020-12-22; cited here as 2021 following the volume 136, issue 1 convention.
- Michael McLeay, Amar Radia, and Ryland Thomas. Money creation in the modern economy. *Bank of England Quarterly Bulletin*, 54(1):14–27, 2014.
- Perry Mehrling. *The New Lombard Street: How the Fed Became the Dealer of Last Resort*. Princeton University Press, Princeton, NJ, 2011.
- Perry Mehrling. The inherent hierarchy of money. Prepared for *Duncan Foley Festschrift* volume and conference, April 20–21, 2012, 2012.
- Carl Menger. On the origin of money. *The Economic Journal*, 2(6):239–255, 1892. Translated by Caroline A. Foley.
- Hyman P. Minsky. *Stabilizing an Unstable Economy*. Yale University Press, New Haven, CT, 1986.
- Ben Norman, Rachel Shaw, and George Speight. The history of interbank settlement arrangements: Exploring central banks’ role in the payment system. Working Paper 412, Bank of England, 2011.
- Joseph M. Ostroy and Ross M. Starr. The transactions role of money. In Benjamin M. Friedman and Frank H. Hahn, editors, *Handbook of Monetary Economics*, volume 1, pages 3–62. North-Holland, Amsterdam, 1990.

- Zoltan Pozsar. Shadow banking: The money view. OFR Working Paper 14-04, Office of Financial Research, U.S. Department of the Treasury, 2014.
- L. C. G. Rogers and Luitgard A. M. Veraart. Failure and rescue in an interbank network. *Management Science*, 59(4):882–898, 2013.
- Joseph E. Stiglitz and Andrew Weiss. Credit rationing in markets with imperfect information. *The American Economic Review*, 71(3):393–410, 1981.
- Nick Szabo. Formalizing and securing relationships on public networks. *First Monday*, 2(9), 1997.
- The Clearing House. CHIPS: Real-time payments for large-value transactions. The Clearing House Payments Company L.L.C., 2025. Volume and liquidity-efficiency statistics as published at theclearinghouse.org.
- James Tobin. Commercial banks as creators of ‘money’. In Deane Carson, editor, *Banking and Monetary Studies*, pages 408–419. Richard D. Irwin, Homewood, IL, 1963.
- Nenad Tomasev, Matija Franklin, Joel Z. Leibo, Julian Jacobs, William A. Cunningham, Iason Gabriel, and Simon Osindero. Virtual agent economies. arXiv preprint arXiv:2509.10147, 2025.
- Robert M. Townsend. Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory*, 21(2):265–293, 1979.
- Knut Wicksell. *Interest and Prices: A Study of the Causes Regulating the Value of Money*. Macmillan, London, 1936. Translated by R. F. Kahn from *Geldzins und Güterpreise*, Gustav Fischer, Jena, 1898.
- Michael Woodford. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press, Princeton, NJ, 2003.
- World Wide Web Consortium. Decentralized identifiers (DIDs) v1.0: Core architecture, data model, and representations. W3C Recommendation, 19 July 2022, 2022.
- L. Randall Wray. *Modern Money Theory: A Primer on Macroeconomics for Sovereign Monetary Systems*. Palgrave Macmillan, Basingstoke, 2nd edition, 2015.
- Yingxuan Yang, Ying Wen, Jun Wang, and Weinan Zhang. Agent exchange: Shaping the future of AI agent economics. arXiv preprint arXiv:2507.03904, 2025.