# What Price Settlement Speed?

A Gabrielli–von Kármán Frontier for Financial Settlement Systems, from Babylonian Storehouses to Proof-of-Stake

Rolf Bastiaanssen*

Bax Innovation, Barcelona

Working paper — November 2, 2026

**Abstract**

In 1950 Gabrielli and von Kármán asked "what price speed?" and answered with a single log–log chart: specific power plotted against velocity for every class of vehicle, bounded below by one limit line. This paper transposes that device to financial settlement. We define *settlement velocity* $v = 1/\tau$ as the inverse of the time to final settlement and *specific exchange drag* $\varepsilon$ as the resource cost of settling one unit of value, in basis points, decomposed into four additive frictions: direct fees, risk and adverse-selection reserves, energy and hardware, and the opportunity cost of locked collateral. Placing six pre-industrial exemplars (c. 1750 BCE–1880) and fifteen modern systems on one diagram yields a nonparametric *GvK Limit Line*, the lower convex hull in log–log space, and a set of institutional regime envelopes. Two analytical results organise the diagram. First, the elasticity of drag with respect to the policy rate equals the collateral share of drag, which explains why netting shields systems from rate shocks. Second, when faster settlement destroys netting, transaction-window and balance-sheet accounting of collateral costs can disagree in *sign*; we call this the settlement liquidity paradox and apply it to the 2024 U.S. move to T+1. Third, a hedonic decomposition of drag into core settlement cost and bundled services (credit float, dispute rights, novation, censorship resistance) shows that the Limit Line is an upper bound on the frontier of pure settlement cost, and splits each system's distance from the line exactly into a within-class distance and a cross-class *wrapper premium*, following the class-by-class logic of Ashby's material charts. We illustrate the results on a calibrated structural simulation panel of 1,466 system-months (2018–2026). The panel is synthetic and benchmarked to public magnitudes; we therefore report no statistical inference on it, and confine a parametric response surface of the calibration to an appendix. We set out what an empirical test would require, and draw implications for regulatory sandboxes evaluating instant, tokenised and T+0 settlement.

**Keywords:** settlement systems; payment economics; liquidity; netting; T+1; proof-of-stake; frontier analysis; regulatory sandboxes.

*Director, Bax Innovation, Barcelona, Spain. Corresponding author: r.bastiaanssen@baxcompany.com. ORCID: 0009-0000-6119-429X. The views expressed are the author's own and do not represent those of Bax Innovation or its clients. Replication code and the calibrated panel are provided in the accompanying replication package (see Data and code availability).

**JEL:** E42, G23, G28, N20, O33.

**Data provenance.** All quantitative results for modern systems in this paper are computed on a *calibrated structural simulation*: a fixed-seed synthetic monthly panel whose baseline magnitudes are benchmarked to public disclosures, not drawn from them record by record. Coordinates for historical exemplars are order-of-magnitude reconstructions with literature ranges. Nothing in this paper should be read as an empirical audit of any named system.

# 1 Introduction

Every settlement system answers the same engineering question: how quickly can an obligation be discharged with finality, and what does that speed cost? A Babylonian temple storehouse settled grain debts at harvest, roughly every six months. A Venetian merchant waited three months for a bill of exchange to be presented in Bruges. Fedwire settles in seconds, Solana in hundreds of milliseconds, and a Nasdaq matching engine in a millisecond. Yet the costs that accompany each of these speeds are rarely compared on a common scale. Payment economists measure liquidity costs of real-time gross settlement (Bech and Garratt, 2003; Martin and McAndrews, 2008); market-design economists measure the rents of speed (Budish et al., 2015; Aquilina et al., 2022); blockchain economists measure the energy or capital that secures a ledger (Budish, 2018; Huberman et al., 2021; Saleh, 2021); and economic historians measure exchange margins and specie-shipping costs (de Roover, 1953; Officer, 1996). Each literature uses its own units.

This paper proposes one unit and one chart. In 1950, Giuseppe Gabrielli and Theodore von Kármán plotted the specific power (power per unit weight per unit speed) of every class of vehicle against its maximum speed on log–log axes (Gabrielli and von Kármán, 1950). The resulting diagram showed, at a glance, which technologies were efficient for their speed, and that all of them lay above a single limit line. The device has since become a standard tool for comparing transport technologies and, in the form of Ashby's material-property charts, for comparing materials (Ashby, 2005). Its power is communicative as much as analytical: it puts heterogeneous technologies on one page and makes the trade-off visible.

We transpose the device to settlement. The horizontal axis is *settlement velocity*, $v = 1/\tau$, where $\tau$ is the time to legal or economic finality. The vertical axis is *specific exchange drag*, $\varepsilon$: the total resource cost of settling value over a month, divided by the value settled, in basis points. Drag is the sum of four frictions that the separate literatures treat in isolation: direct fees and operating costs, risk and adverse-selection reserves, energy and hardware, and the opportunity cost of collateral locked to support settlement. Section 3 formalises these definitions.

The paper makes four contributions. First, it supplies a *common accounting frame* that places a medieval bill of exchange, an interbank netting engine, a card network, a proof-of-work blockchain and a high-frequency matching engine in one coordinate system. Second, it derives two simple analytical results from that frame. The *rate-elasticity identity* (Proposition 1) shows that the elasticity of drag with respect to the policy rate equals the collateral share of drag; netting therefore acts as a shield against monetary tightening. The *settlement liquidity paradox* (Proposition 2) shows that when acceleration destroys netting, a transaction-window measure of collateral cost can fall while a balance-sheet measure rises. Third, it builds a nonparametric frontier, the *GvK Limit Line*, together with institutional regime envelopes and intra-class frontiers, and turns the unit-of-value objection into a result: a *hedonic upper-bound property* (Proposition 3) shows that bundled services can only lift observed drag above core settlement cost, so the Limit Line bounds the frontier of pure settlement cost from above, and each system's distance from it separates exactly into within-class inefficiency and a cross-class

wrapper premium. All three results are illustrated on a calibrated structural panel. Fourth, it translates the frame into questions that regulatory sandboxes can ask of instant, tokenised and T+0 settlement proposals.

We are explicit about what the paper is not. Its modern-system numbers come from a calibrated simulation, not from an audited dataset, and its historical coordinates are reconstructions that carry wide ranges. The contribution is a framework and a set of testable structural predictions. Section 7 sets out the data that an empirical test would require.

The remainder of the paper is organised as follows. Section 2 situates the frame in the literature and records where opinions diverge. Section 3 defines velocity, drag and the frontier, and states the two propositions. Section 4 presents the four-thousand-year taxonomy and its calibration. Section 5 presents the frontier diagram, the intra-class frontiers and the wrapper premium. Section 6 draws policy implications for sandboxes, and Section 7 discusses limitations and an empirical agenda. Section 8 concludes. Appendix B reports a parametric response surface of the calibration.

## 2 Related literature and the range of views

**Gross versus net settlement.** The central trade-off of wholesale payment design is between the liquidity cost of gross settlement and the credit and contagion risk of net settlement. Freixas and Parigi (1998) model the choice and show that gross systems trade liquidity for safety while net systems economise on liquidity at the cost of contagion exposure. Kahn et al. (2003) show that the comparison depends on how settlement risk is priced, and Bech and Garratt (2003) show that real-time gross settlement can induce strategic delay as banks wait for incoming payments. Liquidity-saving mechanisms (Martin and McAndrews, 2008) occupy the middle ground. Our framework places this trade-off on the vertical axis: netting lowers $\mathcal{C}_{\text{coll}}$ but can raise $\mathcal{C}_{\text{risk}}$. Views differ on how large the risk side is; Duffie and Zhu (2011) show that central clearing can even *increase* counterparty exposure when it fragments netting across asset classes. Our frame is agnostic on this question but requires that both sides be measured.

**How fast should settlement be?** Khapko and Zoican (2020) provide the closest theoretical precedent. They show that the optimal settlement lag balances counterparty risk against the liquidity cost of immediate delivery, and that instantaneous settlement is not generally optimal. The industry case for T+1 in the United States (Depository Trust & Clearing Corporation, 2021; U.S. Securities and Exchange Commission, 2023) emphasised reduced counterparty exposure and lower margin; critics emphasised pre-funding and operational costs. Proposition 2 shows that both sides can be correct, because they measure collateral cost on different accounting bases.

**The economics of speed in trading.** Budish et al. (2015) argue that continuous limit order books create a socially wasteful arms race for speed, and Aquilina et al. (2022) estimate that latency-arbitrage races account for a material share of trading costs on the London Stock

Exchange. We draw on this literature for the right-hand edge of the diagram, which we label the *Settlement Mach Barrier*: beyond roughly $10^3\,\mathrm{s}^{-1}$, additional velocity mainly redistributes rents rather than reducing drag.

**The cost of decentralised trust.** Budish (2018) argues that the security of proof-of-work ledgers requires recurring expenditure proportional to the value at stake. Huberman et al. (2021) analyse Bitcoin fees as the outcome of congestion pricing, Biais et al. (2019) study forks and coordination, and Saleh (2021) show that proof-of-stake can secure consensus without energy dissipation, at the cost of locking stake. Chiu and Koeppl (2019) model blockchain settlement for securities and quantify the trade-off between settlement speed and the cost of preventing double spending. Estimates of proof-of-work energy use vary widely across methods (de Vries, 2018). In our frame the debate becomes a question of composition: whether security is paid for in $\mathcal{C}_{\text{energy}}$ or in $\mathcal{C}_{\text{coll}}$.

**Payments as platforms.** For card networks, Rochet and Tirole (2003) show that interchange fees reflect two-sided platform pricing rather than resource costs. This matters for interpretation: the high drag of card networks in our diagram includes a transfer between users, merchants and issuers, not only a resource cost. We return to this caveat in Section 7.

**Economic history.** The historical exemplars rest on literatures with genuine disagreements. On the Champagne fairs, Milgrom et al. (1990) emphasise private-order enforcement through the law merchant, while Edwards and Ogilvie (2012) argue that public authority and the fair-holders' own institutions did the work; the disagreement bears on how much of the fairs' low drag reflected public subsidy. On bills of exchange, de Roover (1953) interprets exchange margins largely as disguised interest, whereas later work stresses negotiability and the transfer function of the bill (Munro, 2003); the interpretation decides how much of the margin belongs in $\varepsilon$. On the Bank of Amsterdam, the recent monetary history of Quinn and Roberds (2024) builds on Quinn and Roberds (2009) and revises the account in Smith (1776). On gold points, Officer (1996) and Spiller and Wood (1988) find arbitrage largely consistent with estimated shipping costs, while Morgenstern (1959) documented frequent apparent violations. For the ancient Near East, Van De Mieroop (2005) and Goetzmann (2016) describe credit and interest, but there is no direct measure of settlement cost. We use the most recent widely accepted treatment as the central value and report the range of views explicitly (Table 1).

# 3 Framework

## 3.1 Velocity and drag

Consider a settlement system $i$ observed in month $t$. Let $\tau_i$ denote the time from initiation to final, irrevocable settlement and define settlement velocity as

$$v_i = 1/\tau_i \qquad [\mathrm{s}^{-1}]. \tag{1}$$

Let $V_{it}$ be the value settled in the month. The system incurs four resource costs, all in currency units per month: direct fees and operating costs $\mathcal{C}_{\text{fees}}$; risk, fraud and adverse-selection reserves $\mathcal{C}_{\text{risk}}$; energy and hardware $\mathcal{C}_{\text{energy}}$; and the opportunity cost of collateral $\mathcal{C}_{\text{coll}}$.

**Definition 1** (Specific exchange drag)**.** *The specific exchange drag of system $i$ in month $t$ is*

$$\varepsilon_{it} \;=\; \frac{\mathcal{C}_{\text{fees}} + \mathcal{C}_{\text{risk}} + \mathcal{C}_{\text{energy}} + \mathcal{C}_{\text{coll}}}{V_{it}} \times 10^4 \quad [bps]. \tag{2}$$

Like specific power in the original diagram, $\varepsilon$ is an intensive quantity: it normalises by throughput, so systems of very different scale can be compared. Additivity assumes that the four frictions are separable resource costs. This is a modelling choice, discussed in Section 7.

### 3.2 Two accounting bases for collateral

Let $K_i$ be the collateral or liquidity that must be held to support settlement, and $r_t$ the opportunity-cost rate (we use SOFR). Two accounting bases are common, and they answer different questions.

**Definition 2** (Collateral accrual bases)**.** *The* transaction-window *basis charges collateral for the time it is in transit:*[1]

$$\mathcal{C}^{w}_{\text{coll}} = \frac{r_t}{360}\, K_i \, \max\!\left(\frac{\tau_i}{86\,400}, \frac{1}{86\,400}\right). \tag{3}$$

*The* balance-sheet carry *basis charges collateral for the whole month in which it must be held:*

$$\mathcal{C}^{c}_{\text{coll}} = \frac{r_t}{360}\, K_i \times 30. \tag{4}$$

The window basis measures the float cost of an individual transfer; the carry basis measures the standing cost of a clearing fund, a nostro balance or a staking position. For a system that holds a permanent buffer, the carry basis is the economically relevant one; for a system that pre-funds only transfer by transfer, the window basis is.

### 3.3 Two structural results

**Proposition 1** (Rate-elasticity identity)**.** *Holding $\mathcal{C}_{\text{fees}}$, $\mathcal{C}_{\text{risk}}$, $\mathcal{C}_{\text{energy}}$, $K_i$ and $\tau_i$ fixed, under either accrual basis*

$$\frac{\partial \ln \varepsilon_{it}}{\partial \ln r_t} \;=\; s^{\text{coll}}_{it} \;\equiv\; \frac{\mathcal{C}_{\text{coll}}}{\mathcal{C}_{\text{fees}} + \mathcal{C}_{\text{risk}} + \mathcal{C}_{\text{energy}} + \mathcal{C}_{\text{coll}}} \;\in\; [0, 1]. \tag{5}$$

*Proof.* $\mathcal{C}_{\text{coll}}$ is linear in $r_t$ under (3) and (4), so $\partial \mathcal{C}_{\text{coll}} / \partial \ln r_t = \mathcal{C}_{\text{coll}}$, and the remaining components do not depend on $r_t$. Dividing by total cost gives the result. □

[1] The one-second floor in equation (3) sets a minimum funding-accrual interval per transfer. Without it, the window cost of sub-second systems would tend to zero with $\tau$ and their drag would lose its collateral term altogether. The floor is conservative: money markets accrue interest on a daily basis, so for transfers that settle within the day, the window basis understates the funding cost actually incurred. This is one reason the carry basis of equation (4) is reported alongside it.

The identity has a direct implication. Netting reduces the collateral needed to settle a given gross value, which lowers $s^{\text{coll}}$ and hence the sensitivity of drag to monetary policy. We call this the *netting shield*. Conversely, gross real-time systems and proof-of-stake ledgers, whose security or finality depends on locked capital, are *capital-bound*: their drag inherits the rate cycle. The identity also bounds what any regression of $\ln \varepsilon$ on $\ln r$ can find: a coefficient outside $[0, 1]$ signals omitted variables or composition effects, not a structural elasticity.

**Proposition 2** (Settlement liquidity paradox)**.** *Suppose that shortening the settlement cycle reduces netting, so that the collateral requirement rises as* $K(\tau) = K_0 (\tau/\tau_0)^{-\eta}$ *with* $\eta > 0$*. Then, as* $\tau$ *falls, the transaction-window collateral cost* $\mathcal{C}^{w}_{\text{coll}} \propto \tau^{1-\eta}$ *falls if and only if* $\eta < 1$*, while the carry cost* $\mathcal{C}^{c}_{\text{coll}} \propto \tau^{-\eta}$ *always rises. For* $0 < \eta < 1$ *the two bases disagree in sign.*

*Proof.* Substitute $K(\tau)$ into (3) (for $\tau \geq 1\,\text{s}$) and (4) and differentiate with respect to $\ln \tau$: the exponents are $1 - \eta$ and $-\eta$ respectively. □

*A square-root benchmark for* $\eta$*.* The scaling law in Proposition 2 need not be ad hoc. Suppose that within a settlement cycle of length $\tau$, trades of size $a$ arrive at rate $\lambda$ with independent and equally likely directions. Gross value in the cycle is $G(\tau) = a\lambda\tau$, while the net position is a zero-drift random walk with expected absolute value $\mathbb{E}|N(\tau)| \approx a\sqrt{2\lambda\tau/\pi}$. If collateral is sized to a fixed multiple of the net position per unit of gross value settled, then, holding monthly value fixed, $K(\tau) \propto \mathbb{E}|N(\tau)|/G(\tau) \propto \tau^{-1/2}$, that is, $\eta = 1/2$. This benchmark lies inside the interval $0 < \eta < 1$ in which the two accounting bases disagree. Serially correlated order flow moves $\eta$ away from one half within the cycle: persistent flow nets less at every horizon and lowers $\eta$, while mean-reverting flow raises it. The benchmark concerns the netting ratio only; the level of a clearing fund also depends on margin horizons and stress add-ons.

*Worked example: T+1.* If halving the cycle from T+2 to T+1 raises the clearing-fund requirement by a factor of 1.35, then $\eta = \ln 1.35/\ln 2 = 0.433$, close to the square-root benchmark. The window-basis collateral cost is multiplied by $2^{-(1-\eta)} = 0.675$, a fall of about one third, while the carry cost rises by 35%. *Worked example: T+0.* If moving from one day to ten seconds multiplies required collateral tenfold, then $\eta = \ln 10/\ln 8640 = 0.254$. The window cost falls to 0.0012 of its former level while the carry cost rises tenfold. The policy debate over accelerated settlement is, in part, a debate between these two accounting bases.

### 3.4 The frontier

**Definition 3** (GvK Limit Line)**.** *Let* $\mathcal{S}_t$ *be the set of points* $(\log_{10} v_i, \log_{10} \varepsilon_{it})$ *for all systems observed at* $t$*, including the historical exemplars. The GvK Limit Line* $\varepsilon_F(v)$ *is the lower boundary of the convex hull of* $\mathcal{S}_t$*. The* frontier distance *of system* $i$ *is*

$$u_{it} = \ln \varepsilon_{it} - \ln \varepsilon_F(v_i) \;\geq\; 0. \tag{6}$$

The Limit Line is nonparametric and is recomputed whenever any input changes. It is the analogue of the dashed line in Gabrielli and von Kármán (1950): an empirical lower envelope,

not a law. Systems on the line ($u = 0$) are *best-in-class for their speed* within the observed set; it does not follow that the line cannot be crossed by a new technology.

**Definition 4** (Regime envelope)**.** *For each institutional regime $R$ (pre-industrial clearing; wholesale netting; gross RTGS and correspondent; retail payments; cryptographic ledgers; market microstructure), the regime envelope is the convex hull of its members' points in log–log space, padded by a disc of radius 0.3 decades.*

Envelopes play the role of the shaded vehicle classes in later versions of the transport diagram. They summarise where a regime operates; they are not fitted curves, and their slopes are descriptive.

### 3.5 The unit of value: a hedonic decomposition

Normalising by $V_{it}$ treats one dollar of settled value as the same good in every system. It is not. A Fedwire dollar is a transfer of central-bank reserves with irrevocable legal discharge. A card dollar also carries an unsecured credit float, a zero-liability fraud guarantee and chargeback rights; a cleared equity trade carries novation to a central counterparty and mutualised default protection; a public-blockchain transfer carries permissionless, censorship-resistant finality. We make this explicit rather than treat it as a caveat.

**Definition 5** (Hedonic decomposition of drag)**.** *Observed drag is core settlement cost plus the shadow cost of bundled non-settlement services:*

$$\varepsilon_{it} \;=\; \varepsilon_{it}^{\text{core}} \;+\; \sum_{k} \phi_k \, Z_{it}^{k}, \qquad \phi_k \geq 0,\; Z_{it}^{k} \geq 0,\; \varepsilon_{it}^{\text{core}} > 0, \tag{7}$$

*where $\varepsilon_{it}^{\text{core}}$ is the cost of discharging the obligation on the ledger, $Z_{it}^{k}$ is the quantity of bundled service $k$ per unit of value (for example $Z^{\text{credit}}$, credit float; $Z^{\text{dispute}}$, chargeback arbitration and zero-liability insurance; $Z^{\text{nov}}$, central-counterparty novation and default-fund mutualisation; $Z^{\text{cens}}$, censorship resistance) and $\phi_k$ its shadow price.*

The decomposition is an accounting identity, not an estimating equation: it requires no knowledge of $\phi_k$ or $Z^k$, only that bundled services are not free and not negative.

**Definition 6** (Intra-class frontier and wrapper premium)**.** *For an institutional class $R$ (wholesale netting; gross RTGS and instant; retail payments; cryptographic state machines; market microstructure; pre-industrial clearing), the intra-class frontier $\varepsilon_F^R(v)$ is the lower boundary of the convex hull of the points of class $R$ alone, on the velocity span of the class. For a member $i \in R$ define the within-class distance $u_{it}^R = \ln \varepsilon_{it} - \ln \varepsilon_F^R(v_i)$ and the wrapper premium $\pi_R(v_i) = \ln \varepsilon_F^R(v_i) - \ln \varepsilon_F(v_i)$.*

The construction follows Ashby (2005): on a material-property chart, ceramics, polymers and metals are not expected to share one property line; each class has its own envelope and lower boundary, comparisons within a class measure performance, and the gaps between classes measure what the class buys with its extra cost.

**Proposition 3** (Hedonic upper bound and wrapper decomposition)**.** *Under Definition 5:*

*(a)* $\varepsilon_{it}^{\text{core}} \leq \varepsilon_{it}$ *for every system and month.*

*(b) Let* $\varepsilon_F^{\text{core}}(v)$ *be the Limit Line constructed from the core points* $(\log v_i, \log \varepsilon_{it}^{\text{core}})$*. Then* $\varepsilon_F^{\text{core}}(v) \leq \varepsilon_F(v)$ *at every* $v$ *in the observed span: the observed Limit Line is an upper bound on the frontier of pure settlement cost.*

*(c) For every* $i \in R$*,*

$$u_{it} \;=\; u_{it}^R \;+\; \pi_R(v_i), \qquad u_{it}^R \geq 0, \quad \pi_R(v_i) \geq 0. \tag{8}$$

*(d) If the bundle is a class-common proportional markup,* $\ln \varepsilon_{it} = \ln \varepsilon_{it}^{\text{core}} + \omega_R$ *with* $\omega_R \geq 0$ *for all* $i \in R$*, then* $u_{it}^R$ *is identical on observed and on core drag, and* $\pi_R(v_i) \geq \omega_R$ *whenever the core frontier of class* $R$ *is no lower than the observed Limit Line at* $v_i$*.*

*Proof.* (a) follows from $\phi_k Z_{it}^k \geq 0$. For (b), the lower boundary of the convex hull of a finite point set is its greatest convex minorant, the pointwise supremum of the affine functions lying below every point. Lowering points weakly, with abscissae fixed, shrinks the set of affine minorants, so the supremum weakly falls; by (a) the core points lie weakly below the observed points. For (c), the identity holds by Definition 6; $u_{it}^R \geq 0$ because $\varepsilon_F^R$ is a minorant of the class's points, and $\pi_R \geq 0$ because every affine minorant of all points is also a minorant of the class's points, so $\varepsilon_F \leq \varepsilon_F^R$ on the class's span. For (d), a common vertical shift of the class's points shifts its greatest convex minorant by the same constant, so the shift cancels in $u^R$; and $\pi_R = \omega_R + [\ln \varepsilon_F^{R,\text{core}} - \ln \varepsilon_F]$, whose bracket is non-negative under the stated condition. □

Proposition 3 turns the unit-of-value objection into a measurement design. Part (b) means that bundled services can bias the Limit Line only upwards: the frontier of pure settlement cost lies on or below the line drawn here, so statements of the form "system $i$ is at least $u$ above best practice" remain conservative for core settlement. Part (c) assigns every system's distance two readings. The within-class distance $u^R$ is the efficiency benchmark, comparing like with like. The wrapper premium $\pi_R$ is the drag a class carries above the cross-class frontier at the same speed; it combines the price of the services the class bundles with any difference in core technology, and part (d) states when it bounds the former from above. Separating the two components requires measuring $Z^k$, which is part of the empirical agenda in Section 7.

# 4 A four-thousand-year taxonomy

## 4.1 Historical exemplars

Following the convention of the original diagram, which plotted representative vehicles rather than every model, we select one exemplar for each pre-electronic settlement mechanism (Table 1, Panel A): a commodity ledger (Babylonian temple storehouse), periodic multilateral

netting (the Champagne fairs), the bill of exchange (Venice), the public giro bank (the Amsterdam Wisselbank), interbank cheque clearing (the London Bankers' Clearing House) and specie arbitrage (transatlantic gold points). Settlement times follow the institutional cycle: one harvest, one fair, one maritime usance, two days of ledger reconciliation, one daily clearing and one steamship crossing.

Drag coordinates are the weakest part of the taxonomy and we treat them accordingly. For each exemplar we take as the central value the reading of the most recent widely accepted treatment and report a range that brackets the competing interpretations surveyed in Section 2. The ranges are our reconstructions; they are not figures reported by the cited authors. The Venetian value (200 bps; range 100–400) reflects exchange margins, brokerage and courier costs on a one-way bill, excluding any pure time-value component. The Wisselbank value (10 bps; range 1–25) combines small transfer fees with the cost of the receipt system that Quinn and Roberds (2024) describe. The gold-point value (50 bps; range 30–100) is a one-way shipping cost including freight, insurance, assay and interest in transit, consistent with the secular narrowing documented by Officer (1996). The Babylonian value is the most speculative and carries the widest range (350 bps; range 150–1,000). Because historical exemplars pre-date modern money markets, their drag is held static and carries no policy-rate term.

### 4.2 Modern systems

Panel B of Table 1 lists fifteen modern systems across wholesale netting, gross real-time and correspondent banking, retail card and instant payments, proof-of-work, proof-of-stake and layer-2 ledgers, and an exchange matching engine. For each system the calibration specifies settlement time, monthly value, locked collateral, fees, risk reserves and energy cost. Baseline magnitudes are benchmarked to the orders of magnitude found in public sources (central-bank payment statistics, operator disclosures, regulatory filings and energy-consumption estimates), and the monthly panel evolves them with a fixed-seed stochastic process, observed-style SOFR and industrial power-price paths, launch dates (FedNow, Pix, Solana), the Ethereum Merge (September 2022) and the U.S. move to T+1 (May 2024). The panel contains 1,466 system-months over 106 months for 15 systems; SOFR ranges from 0.03% to 5.41%. Appendix A documents the generating process.

## 5 The settlement frontier

### 5.1 The diagram

Figure 1 is the settlement analogue of the Gabrielli–von Kármán chart. Four features stand out.

*The Limit Line is U-shaped in log–log space.* From the slow left edge it falls through the historical clearing mechanisms and the wholesale netting engines to a minimum near CHIPS, rises gently through gross real-time systems and climbs steeply into the Mach Barrier, where only the matching engine operates. The descending branch reflects the historical substitution

**Table 1:** The settlement taxonomy: six historical exemplars and fifteen modern systems

| System | Mechanism / regime | $\tau$ | $\varepsilon$ (bps) | Range (sources) / dominant friction |
|---|---|---|---|---|
| *Panel A. Pre-industrial exemplars (static, literature-calibrated)* | | | | |
| Babylonian temple storehouse (c. 1750 BCE) | Commodity ledger, harvest settlement | 180 d | 350 | 150–1,000 (Van De Mieroop, 2005; Goetzmann, 2016) |
| Champagne fairs clearing (c. 1250) | Periodic multilateral netting | 60 d | 20 | 10–100 (Edwards and Ogilvie, 2012) |
| Venetian bills of exchange (c. 1500) | Bill of exchange (cambium) | 90 d | 200 | 100–400 (de Roover, 1953; Mueller, 1997; Munro, 2003) |
| Amsterdam Wisselbank (c. 1650) | Public giro bank, book entry | 2 d | 10 | 1–25 (Quinn and Roberds, 2009, 2024) |
| London Bankers' Clearing House (c. 1820) | Interbank cheque netting | 1 d | 8 | 2–20 (Cannon, 1910; Matthews, 1921) |
| Transatlantic gold points (c. 1880) | Specie shipment arbitrage | 14 d | 50 | 30–100 (Officer, 1996; Spiller and Wood, 1988) |
| *Panel B. Modern systems (2018–2026 calibrated panel, means; $\tau$-window basis)* | | | | |
| CHIPS | Wholesale netting | 1 h | 0.040 | fees (99%) |
| ACH | Wholesale netting | 2 d | 1.0 | risk (98%) |
| DTCC/NSCC equities | Wholesale netting | 1.7 d | 0.86 | risk (58%) |
| Fedwire Funds | Gross RTGS / corr. | 30 s | 0.10 | fees (100%) |
| FedNow | Gross RTGS / corr. | 5 s | 0.35 | fees (57%) |
| TIPS / SEPA Instant | Gross RTGS / corr. | 10 s | 0.20 | fees (75%) |
| SWIFT correspondent | Gross RTGS / corr. | 3 d | 60.2 | fees (83%) |
| Visa | Retail card | 2 d | 220 | fees (82%) |
| Mastercard | Retail card | 2 d | 223 | fees (83%) |
| Pix | Retail instant | 5 s | 3.5 | fees (71%) |
| Bitcoin L1 | Crypto PoW | 1 h | 245 | energy (95%) |
| Ethereum L1 | Crypto PoW→PoS | 12.8 min | 30.8 | fees (51%) |
| Solana L1 | Crypto PoS | 0.4 s | 2.2 | energy (86%) |
| Lightning | Crypto L2 | 1 s | 18.5 | energy (92%) |
| Nasdaq matching | Microstructure | 1 ms | 8.4 | risk (96%) |

*Notes:* Panel A: central drag values follow the most recent widely accepted treatment; ranges are the author's reconstructions bracketing competing interpretations in the cited sources and are not figures reported by those sources. Historical drag is static (no policy-rate term). Panel B: means over the calibrated simulation panel (2018–2026), drag on the transaction-window basis of equation (3); "dominant friction" is the largest of the four components in equation (2), with its share of total drag. Modern values are calibrated benchmarks, not audited measurements.

of ledgers and netting for physical movement of value; the ascending branch reflects the growing cost of liquidity, hardware and adverse selection at very high velocity.

*Netting defines the bottom of the chart.* Apart from the slowest exemplar, every vertex of the descending branch is a netting mechanism: the Champagne fairs, ACH and CHIPS. The London Clearing House and the Wisselbank sit just above the line. Netting is the oldest and most persistent technology for reducing drag.

*Regimes slope downward internally.* Within each envelope, faster members tend to have lower drag. This pattern is descriptive. It is consistent with technological maturation within a regime (newer vintages amortise fixed costs over more value), but the envelopes contain heterogeneous systems and the slope may reflect composition. It should not be read as evidence that speeding up a given system lowers its cost.

*Retail and cryptographic regimes sit far above the line.* Card networks lie about two orders

of magnitude above the line at their velocity, reflecting platform pricing (Rochet and Tirole, 2003) and fraud reserves. Bitcoin lies higher still, with energy accounting for roughly 94% of its drag. The Ethereum Merge moves a single system vertically, from 44.8 to 15.0 bps in panel means, at unchanged velocity.

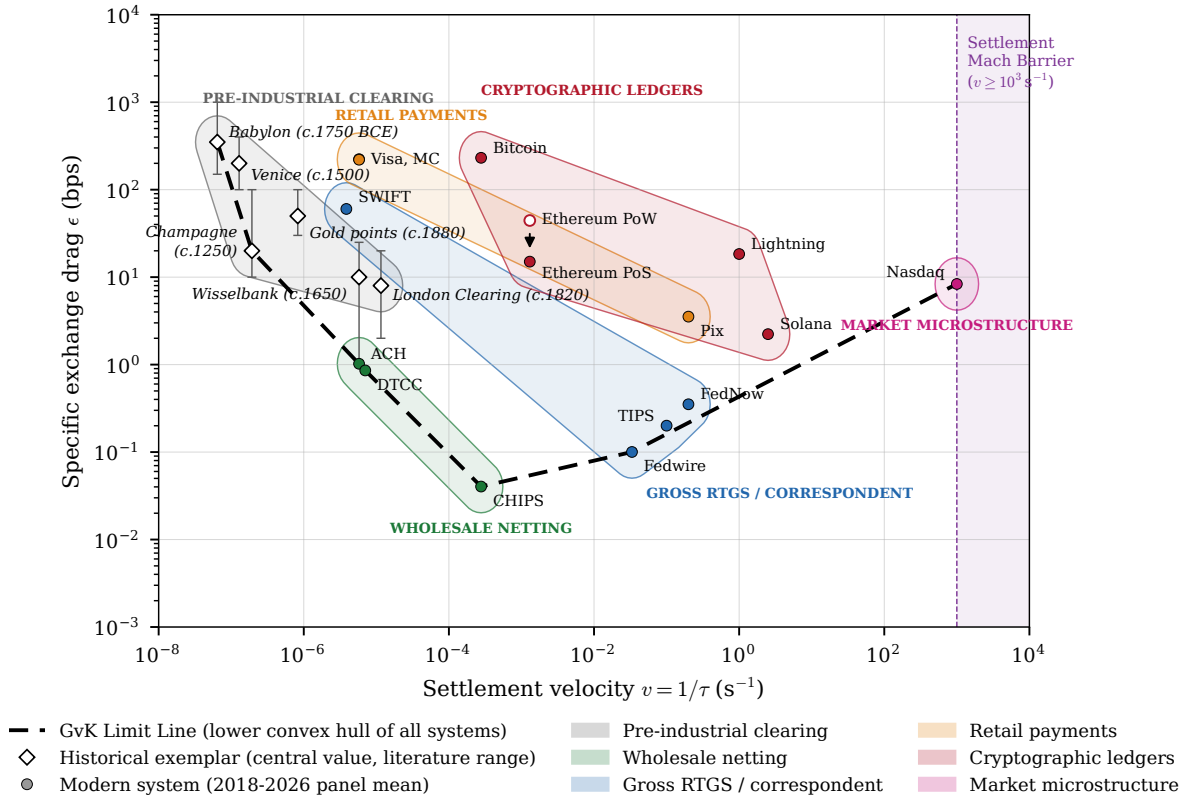



**Figure 1:** The Gabrielli–von Kármán settlement diagram. Horizontal axis: settlement velocity $v = 1/\tau$; vertical axis: specific exchange drag $\varepsilon$ (bps), both logarithmic. Diamonds: historical exemplars at their central values, with whiskers showing the literature range (Table 1). Circles: modern systems ("Visa, MC": Visa and Mastercard, which coincide), 2018–2026 means from the calibrated simulation panel (transaction-window basis). Shaded areas: institutional regime envelopes (padded convex hulls of central values). Dashed line: the GvK Limit Line, the lower convex hull of all plotted points. The arrow marks the Ethereum Merge (September 2022). The shaded band at $v \geq 10^3\,\mathrm{s}^{-1}$ marks the Settlement Mach Barrier.

## 5.2 Intra-class frontiers and the wrapper premium

Table 2 applies Proposition 3(c) to the points of Figure 1. Three patterns stand out. First, the wholesale netting class *is* the Limit Line over its span: CHIPS, ACH and DTCC carry no wrapper premium (the largest is 0.004), which is the netting floor stated in the class vocabulary. Second, the retail premium is large and is a premium for the wrapper, not for speed: card networks sit $\pi_R = 5.38$ above the cross-class frontier, a factor of about 216, while Pix, the retail system with the thinnest bundle (no credit float and limited dispute rights), carries a premium of 2.79, a factor of about 16. Third, the cryptographic class spans the widest range of premia, from 1.25 for Solana to 8.65 for Bitcoin; Ethereum's premium is unchanged by the Merge (5.63) because it is a property of the class frontier at Ethereum's velocity, while the

Merge removes Ethereum's within-class distance (from 1.08 to 0.00), moving it onto the class frontier.

**Table 2:** Intra-class frontiers: within-class distance and wrapper premium

| System | $\varepsilon$ (bps) | $u$ | $u^R$ | $\pi_R$ | $e^{\pi_R}$ |
|---|---|---|---|---|---|
| *Wholesale netting* (3 systems) | | | | | |
| ACH | 1.0 | 0.00 | 0.00 | 0.00 | 1 |
| DTCC/NSCC equities | 0.86 | 0.00 | 0.00 | 0.00 | 1 |
| CHIPS | 0.040 | 0.00 | 0.00 | 0.00 | 1 |
| *Gross RTGS / instant* (4 systems) | | | | | |
| SWIFT correspondent | 60.2 | 3.73 | 0.00 | 3.73 | 41 |
| Fedwire Funds | 0.10 | 0.00 | 0.00 | 0.00 | 1 |
| TIPS / SEPA Instant | 0.20 | 0.22 | 0.00 | 0.22 | 1 |
| FedNow | 0.35 | 0.49 | 0.00 | 0.49 | 2 |
| *Retail payments* (3 systems) | | | | | |
| Mastercard | 223 | 5.39 | 0.01 | 5.38 | 216 |
| Visa | 220 | 5.38 | 0.00 | 5.38 | 216 |
| Pix | 3.5 | 2.79 | 0.00 | 2.79 | 16 |
| *Cryptographic state machines* (5 systems) | | | | | |
| Bitcoin L1 | 231 | 8.65 | 0.00 | 8.65 | 5,732 |
| Ethereum (PoW, to Aug. 2022) | 44.4 | 6.71 | 1.08 | 5.63 | 278 |
| Ethereum (PoS, from Sep. 2022) | 15.0 | 5.63 | 0.00 | 5.63 | 278 |
| Lightning | 18.4 | 3.75 | 1.88 | 1.87 | 7 |
| Solana L1 | 2.2 | 1.25 | 0.00 | 1.25 | 3 |
| *Market microstructure* (1 system) | | | | | |
| Nasdaq matching | 8.4 | 0.00 | 0.00 | 0.00 | 1 |

*Notes:* Points of Figure 1: geometric means of monthly drag over 2018–2026 on the calibrated simulation panel (transaction-window basis), Ethereum split at the Merge; the six historical exemplars enter the Limit Line and form the pre-industrial class (not shown). $u$: log distance to the GvK Limit Line, equation (6); $u^R$: log distance to the intra-class frontier; $\pi_R$: wrapper premium, $\ln \varepsilon_F^R - \ln \varepsilon_F$, with $u = u^R + \pi_R$ (equation (8)); $e^{\pi_R}$: premium as a multiple of frontier drag. Instant payment systems follow their category in the panel. Calibrated benchmarks, not audited measurements.

The table also shows the limit of the exercise. With three to five systems per class, 13 of the 16 modern points are vertices of their own class frontier, so $u^R = 0$ for them by construction; only Mastercard, Ethereum before the Merge and Lightning have a measurable within-class distance. The wrapper premium is well defined on this sample; the within-class benchmark becomes informative only when each class is populated with several comparable systems, which is what the reporting protocol in Section 7 is designed to collect. No sampling-based statistics are attached to the table, because the panel is generated by a known process; Appendix B summarises that process with a parametric response surface.

### 5.3 Collateral transmission over the rate cycle

Figure 2 plots drag for four systems through the 2022–2023 tightening, on both accounting bases. On the window basis (panel a), the capital-bound systems barely react: Fedwire's thirty-second window makes its collateral cost negligible, and Ethereum's drag after the Merge is dominated by fees. On the carry basis (panel b), the same systems move with SOFR, as Proposition 1 predicts for systems with a large collateral share, while netted ACH and fee-dominated Visa remain flat. The choice of accounting basis, not the data, decides whether the

rate cycle appears to matter.

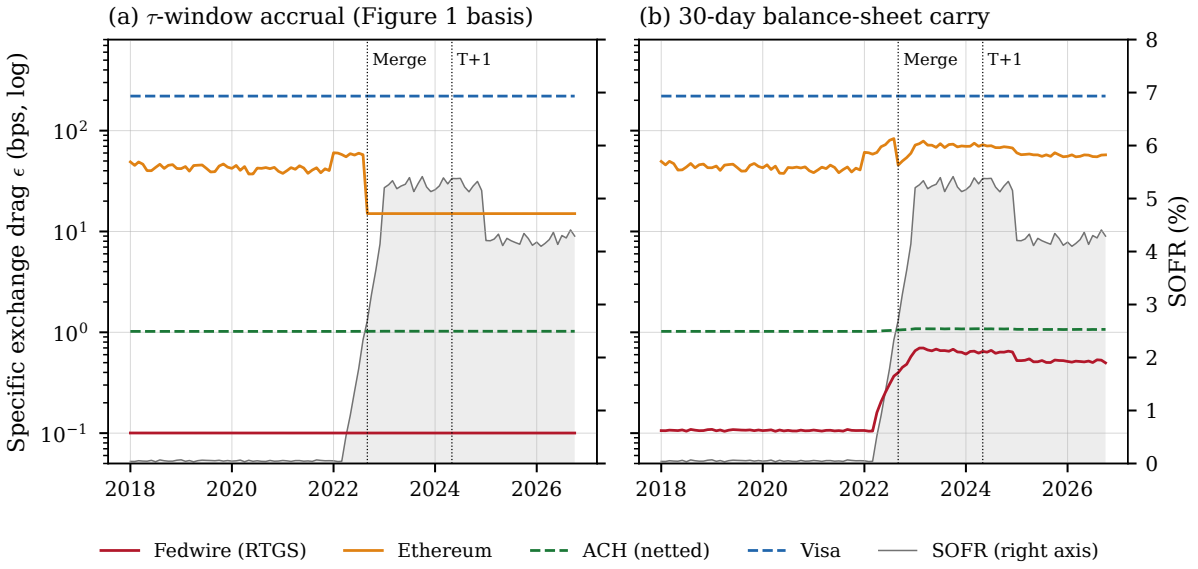



**Figure 2:** Collateral transmission, 2018–2026 (calibrated simulation panel). Specific exchange drag $\varepsilon$ (left axis, log) for Fedwire, Ethereum, ACH and Visa; SOFR shaded (right axis). Panel (a): transaction-window accrual, the basis of Figure 1 and Table 2. Panel (b): 30-day balance-sheet carry. Dotted lines mark the Ethereum Merge (September 2022) and the U.S. move to T+1 (May 2024).

### 5.4 Two episodes

**The Merge: substituting capital for energy.** In the calibration, Ethereum's energy cost falls by 99.94% at the Merge. Total drag falls on the window basis, from 44.8 to 15.0 bps in means. On the carry basis, however, drag at the 2023 rate peak reaches about 79 bps (against 15 bps on the window basis), because the locked stake, now the sole security resource, is charged its full monthly opportunity cost. Proof-of-stake did not eliminate the cost of security; it moved it from $\mathcal{C}_{\text{energy}}$ to $\mathcal{C}_{\text{coll}}$, and therefore from the energy market to the money market (Saleh, 2021).

**T+1: the liquidity paradox in the panel.** The calibration raises DTCC's clearing-fund collateral by a factor of 1.35 when the cycle shortens in May 2024. Between April and May 2024, DTCC drag changes by $-3.9\%$ on the window basis and by $+28.8\%$ on the carry basis: the two bases disagree in sign, exactly as Proposition 2 predicts for $0 < \eta < 1$. The carry change is smaller than the 35% implied by collateral alone because fees and risk reserves are unchanged.

## 6 Implications for regulatory sandboxes

Regulatory sandboxes and pilot regimes, such as the EU pilot regime for DLT market infrastructures (European Parliament and Council of the European Union, 2022), now routinely host proposals for instant, tokenised or atomic settlement. The framework suggests five questions that a sandbox authority can ask.

1. **On which basis is the collateral saving claimed?** Proposition 2 shows that acceleration can reduce float while increasing standing liquidity needs. Applicants should report collateral cost on both the transaction-window and the balance-sheet carry basis, together with the implied $\eta$, the elasticity of required liquidity with respect to the settlement cycle.

2. **What happens to netting?** The netting floor in Figure 1 is the most persistent regularity in four thousand years of settlement. Atomic delivery-versus-payment on a shared ledger removes the netting window by design. Sandboxes should require evidence of how liquidity-saving or netting functions are preserved, for example through liquidity-saving queues (Martin and McAndrews, 2008) or programmable netting on unified ledgers (Bank for International Settlements, 2023).

3. **Where is the cost of security booked?** The Merge episode shows that security costs migrate between components. Environmental disclosures for crypto-asset consensus mechanisms under MiCA (European Parliament and Council of the European Union, 2023) capture $\mathcal{C}_{\text{energy}}$; a complete assessment should also capture the capital locked in staking, $\mathcal{C}_{\text{coll}}$, and its sensitivity to the policy rate.

4. **How does drag respond to monetary tightening?** By Proposition 1, the rate elasticity of a design equals its collateral share. Capital-bound designs, including prefunded instant payment schemes now mandated in the euro area (European Parliament and Council of the European Union, 2024), should be stress-tested across rate regimes, not only at the rate prevailing during the pilot.

5. **Is the proposed speed beyond the Mach Barrier?** Beyond about $10^3\,\mathrm{s}^{-1}$, additional velocity mainly redistributes rents (Budish et al., 2015; Aquilina et al., 2022). Proposals that claim efficiency gains from sub-millisecond settlement should demonstrate a resource saving, not only a latency improvement.

Table 3 condenses the five questions into a due-diligence scorecard. Each row names the claim that most often fails in a proposal and the test, stated in the quantities of Section 3, that an authority can require before a pilot is authorised. The interactive explorer that accompanies this paper implements the same scorecard: an assessor enters the applicant's current and proposed latency, locked capital and cost lines, and the scorecard computes $\eta$, both collateral bases, the collateral share at alternative policy rates and the position relative to the Mach Barrier.

These questions are consistent with the liquidity and settlement principles of Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions (2012), which already require infrastructures to measure and manage liquidity risk. The framework adds a common unit, $\varepsilon$, that makes trade-offs comparable across technologies and over time.

**Table 3:** Sandbox due-diligence scorecard derived from the GvK framework.

| Dimension | Typical failure in a proposal | Required GvK test |
|---|---|---|
| Liquidity and float | Collateral saving claimed on transaction float alone. | Report $\mathcal{C}_{\text{coll}}$ on the window basis (3) and the 30-day carry basis (4); state the implied $\eta$ (Proposition 2). |
| Netting preservation | Atomic delivery-versus-payment removes the netting window, raising gross liquidity needs. | Document the liquidity-saving mechanism: queues, offsetting algorithms or programmable batch netting, and its effect on $K$. |
| Security migration | Proof-of-stake presented as near-zero cost because energy use is small. | Book the opportunity cost of staked or locked capital, $\mathcal{C}_{\text{coll}}$, next to $\mathcal{C}_{\text{energy}}$; report the collateral share $s^{\text{coll}}$. |
| Monetary robustness | Liquidity model tested only at the rate prevailing during the pilot, often near zero. | Recompute $\varepsilon$ across rate regimes (for example 0% and 5.5%); the rate elasticity equals $s^{\text{coll}}$ (Proposition 1). |
| Latency value | Hardware and co-location costs justified by sub-millisecond finality. | Beyond $10^3\,\text{s}^{-1}$, show that $\varepsilon$ falls (a resource saving), not only that $\tau$ falls (a redistribution of rents). |

# 7 Limitations and an empirical agenda

**Data.** The central limitation is provenance. Modern-system values come from a calibrated simulation, so any statistic computed on them recovers the structure of the calibration and cannot confirm it. Historical values are reconstructions with ranges spanning up to an order of magnitude. The diagram's broad geometry (the netting floor, the U-shaped Limit Line and the vertical gap of retail and proof-of-work systems) is robust to moving any exemplar within its range; the exact position of the line at the slow end is not.

**Inference.** We attach no standard errors, $t$ statistics or significance levels to any result. On a panel generated by the author's own calibrated process, sampling-based inference measures the number of simulated draws, not uncertainty about the world, and a test of $H_0\colon \beta = 0$ tests a parameter the researcher chose. The main results therefore rest on analytical statements (Propositions 1–3) and on nonparametric constructions (the Limit Line, the regime envelopes and the intra-class frontiers). The translog surface in Appendix B is a descriptive emulator of the calibration; it is not evidence for the propositions, and it should be re-estimated, with inference, only on measured data collected under a common reporting standard.

**Comparability and the unit of value.** Section 3 formalises the comparability problem rather than setting it aside. By Proposition 3, bundled services can only raise observed drag above core settlement cost, so the Limit Line is a conservative, upper-bound reference for pure settlement cost, and the distance of a system from it splits exactly into a within-class distance and a wrapper premium. What the framework does not supply is the split of the premium itself between the price of bundled services ($\sum_k \phi_k Z^k$) and differences in core technology. That split requires service quantities, for example chargeback volumes, credit-float balances, default-fund

sizes or validator decentralisation, and a pricing model for each, in the spirit of Khapko and Zoican (2020) for counterparty risk. Until such data exist, the premium is best read as an upper bound on what users pay for the wrapper when the class's core technology is no better than the cross-class frontier.

**Resource costs versus transfers.** Equation (2) counts every payment out of the transacting parties as friction, but not every payment is a resource cost. Much of card interchange funds cardholder rewards, interest-free periods and acquiring margins: it is a price structure on a two-sided platform (Rochet and Tirole, 2003), a transfer between merchants, issuers and cardholders. Proof-of-stake rewards are financed by token issuance and fees, and the base-fee burn introduced by EIP-1559 redistributes value from transactors to all token holders (Roughgarden, 2020); neither consumes physical resources in the way that proof-of-work hashing consumes electricity and hardware. The framework therefore supports two readings. *Private exchange drag*, the all-in $\varepsilon$ of Equation (2), is the cost faced by a merchant or transactor choosing between systems, and is the quantity plotted in Figure 1. *Social resource drag* retains only $\mathcal{C}_{\text{energy}}$, the operating component of $\mathcal{C}_{\text{fees}}$ and the deadweight part of $\mathcal{C}_{\text{risk}}$, and excludes transfers; it is the relevant quantity for welfare and environmental assessment. The protocol in the replication package records the cost basis of each line so that the two readings can be separated in measured data.

**Additivity and tail risk.** Equation (2) treats the four frictions as additive and separable, and measures risk by expected-loss reserves. Both assumptions understate the interaction between collateral and risk. Netting lowers $\mathcal{C}_{\text{coll}}$ precisely by replacing bilateral exposures with exposures to a central counterparty, which concentrates credit, liquidity and operational risk in a single node (Freixas and Parigi, 1998; Duffie and Zhu, 2011). Expected-loss reserves, in turn, underweight low-frequency, high-impact events such as the default of Lehman Brothers in 2008 or the dash for cash in U.S. Treasury markets in March 2020 (Duffie, 2020; Financial Stability Board, 2020). A low $\varepsilon$ can therefore reflect visible daily float traded for an unpriced catastrophic tail. A lower $\varepsilon$ is not, by itself, evidence of a more efficient or more resilient system; a macro-prudential assessment would replace $\mathcal{C}_{\text{risk}}$ by a tail-risk charge, for example the cost of the default fund and committed liquidity facilities sized to a stressed scenario, and treat the frictions jointly.

**An empirical agenda.** A full empirical test requires, for each system and with several comparable systems per class so that intra-class frontiers are identified: settled value from central-bank payment statistics or operator reports; fee revenue from annual reports or on-chain data; collateral and liquidity usage from intraday-liquidity data, clearing-fund disclosures or staking data; loss and fraud reserves from regulatory filings; and energy use from operator or independent estimates. Quarterly frequency is realistic for most series; proprietary data for trading venues will be the binding constraint. With such data, Proposition 1 can be tested within systems over the 2022–2023 tightening, and Proposition 2 on the 2024 North American move to T+1. The replication package proposes a common reporting format for this purpose (the GvK Open Protocol, specification `gvk-spec-v1`): operators report monthly latency, settled

value, locked capital and three cost lines per anonymised system, with a cost-basis flag that separates observed, estimated and allocated figures, and a publication gate that releases non-public data only in aggregates of at least three contributors. Populating the modern panel with such data is the main priority for the next version of this paper; it will show whether a measured cross-regime frontier corroborates the calibrated one.

## 8 Conclusion

Gabrielli and von Kármán showed that a single, well-chosen pair of axes could summarise a century of transport engineering. This paper argues that settlement velocity and specific exchange drag do the same for four thousand years of settlement technology. The resulting diagram shows a persistent netting floor, a U-shaped limit line and a Mach barrier at very high speed. Three results organise it: the rate elasticity of drag equals its collateral share; acceleration that destroys netting can lower float costs while raising balance-sheet costs; and because bundled services can only add to drag, the limit line bounds the frontier of pure settlement cost from above, while each system's distance from it splits exactly into within-class inefficiency and a cross-class wrapper premium. The quantitative illustrations rest on a calibrated simulation and literature-based reconstructions; the next step is to measure the same quantities from public data. The question "what price speed?" applies to settlement as well as to vehicles, and the answer depends on how the cost of collateral is counted.

## Data and code availability

The calibrated panel, its fixed-seed generator, the estimation and figure scripts, the LaTeX source and the interactive explorer are distributed as a replication package that accompanies this working paper and is available from the author on request; a public archive with a persistent identifier will be cited in the next version. A single command (`bash scripts/reproduce_all.sh`) regenerates the panel, Tables 1, 2 and B.1, every number quoted in the text, both figures and this document. Appendix A lists the steps, and each equation in Section 3 is cross-referenced to the function that implements it.

## Use of generative AI

The author used generative-AI tools to assist with software development, data-pipeline scripting and language editing. All modelling choices, calibrations, interpretations and conclusions are the author's own; the author reviewed and verified all content and takes full responsibility for the paper.

## References


Aigner, D., Lovell, C. A. K., and Schmidt, P. (1977). Formulation and estimation of stochastic frontier production function models. *Journal of Econometrics*, 6(1):21–37.

Aquilina, M., Budish, E., and O'Neill, P. (2022). Quantifying the high-frequency trading "arms race". *Quarterly Journal of Economics*, 137(1):493–564.

Ashby, M. F. (2005). *Materials Selection in Mechanical Design*. Butterworth-Heinemann, Oxford, 3rd edition.

Bank for International Settlements (2023). Blueprint for the future monetary system: Improving the old, enabling the new. In *Annual Economic Report 2023*, chapter III. Bank for International Settlements, Basel.

Bech, M. L. and Garratt, R. (2003). The intraday liquidity management game. *Journal of Economic Theory*, 109(2):198–219.

Biais, B., Bisière, C., Bouvard, M., and Casamatta, C. (2019). The blockchain folk theorem. *Review of Financial Studies*, 32(5):1662–1715.

Budish, E. (2018). The economic limits of Bitcoin and the blockchain. NBER Working Paper 24717, National Bureau of Economic Research.

Budish, E., Cramton, P., and Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4):1547–1621.

Cannon, J. G. (1910). Clearing houses. Senate Document 491, National Monetary Commission, U.S. Senate, Washington, DC.

Chiu, J. and Koeppl, T. V. (2019). Blockchain-based settlement for asset trading. *Review of Financial Studies*, 32(5):1716–1753.

Christensen, L. R., Jorgenson, D. W., and Lau, L. J. (1973). Transcendental logarithmic production frontiers. *Review of Economics and Statistics*, 55(1):28–45.

Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions (2012). Principles for financial market infrastructures. Technical report, Bank for International Settlements, Basel.

de Roover, R. (1953). *L'évolution de la lettre de change, XIVe–XVIIIe siècles*. Armand Colin, Paris.

de Vries, A. (2018). Bitcoin's growing energy problem. *Joule*, 2(5):801–805.

Depository Trust & Clearing Corporation (2021). Advancing together: Leading the industry to accelerated settlement. White paper, DTCC, New York.

Duffie, D. (2020). Still the world's safe haven? Redesigning the U.S. Treasury market after the COVID-19 crisis. Hutchins Center Working Paper 62, Hutchins Center on Fiscal and Monetary Policy, Brookings Institution.

Duffie, D. and Zhu, H. (2011). Does a central clearing counterparty reduce counterparty risk? *Review of Asset Pricing Studies*, 1(1):74–95.

Edwards, J. and Ogilvie, S. (2012). What lessons for economic development can we draw from the Champagne fairs? *Explorations in Economic History*, 49(2):131–148.

European Parliament and Council of the European Union (2022). Regulation (EU) 2022/858 on a pilot regime for market infrastructures based on distributed ledger technology. Official Journal of the European Union, L 151.

European Parliament and Council of the European Union (2023). Regulation (EU) 2023/1114 on markets in crypto-assets (MiCA). Official Journal of the European Union, L 150.

European Parliament and Council of the European Union (2024). Regulation (EU) 2024/886 as regards instant credit transfers in euro. Official Journal of the European Union, L series.

Financial Stability Board (2020). Holistic review of the March market turmoil. Technical report, Financial Stability Board, Basel.

Freixas, X. and Parigi, B. (1998). Contagion and efficiency in gross and net interbank payment systems. *Journal of Financial Intermediation*, 7(1):3–31.

Gabrielli, G. and von Kármán, T. (1950). What price speed? Specific power required for propulsion of vehicles. *Mechanical Engineering*, 72(10):775–781.

Goetzmann, W. N. (2016). *Money Changes Everything: How Finance Made Civilization Possible*. Princeton University Press, Princeton, NJ.

Greene, W. H. (1980). Maximum likelihood estimation of econometric frontier functions. *Journal of Econometrics*, 13(1):27–56.

Huberman, G., Leshno, J. D., and Moallemi, C. (2021). Monopoly without a monopolist: An economic analysis of the Bitcoin payment system. *Review of Economic Studies*, 88(6):3011–3040.

Kahn, C. M., McAndrews, J., and Roberds, W. (2003). Settlement risk under gross and net settlement. *Journal of Money, Credit and Banking*, 35(4):591–608.

Khapko, M. and Zoican, M. (2020). How fast should trades settle? *Management Science*, 66(10):4573–4593.

Martin, A. and McAndrews, J. (2008). Liquidity-saving mechanisms. *Journal of Monetary Economics*, 55(3):554–567.

Matthews, P. W. (1921). *The Bankers' Clearing House: What It Is and What It Does*. Sir Isaac Pitman & Sons, London.

Milgrom, P. R., North, D. C., and Weingast, B. R. (1990). The role of institutions in the revival of trade: The law merchant, private judges, and the Champagne fairs. *Economics & Politics*, 2(1):1–23.

Morgenstern, O. (1959). *International Financial Transactions and Business Cycles*. Princeton University Press for the National Bureau of Economic Research, Princeton, NJ.

Mueller, R. C. (1997). *The Venetian Money Market: Banks, Panics, and the Public Debt, 1200–1500*. Johns Hopkins University Press, Baltimore.

Munro, J. H. (2003). The medieval origins of the financial revolution: Usury, rentes, and negotiability. *International History Review*, 25(3):505–562.

Officer, L. H. (1996). *Between the Dollar–Sterling Gold Points: Exchange Rates, Parity, and Market Behavior*. Cambridge University Press, Cambridge.

Quinn, S. and Roberds, W. (2009). An economic explanation of the early Bank of Amsterdam, debasement, bills of exchange, and the emergence of the first central bank. In Atack, J. and Neal, L., editors, *The Origins and Development of Financial Markets and Institutions*, pages 32–70. Cambridge University Press, Cambridge.

Quinn, S. and Roberds, W. (2024). *How a Ledger Became a Central Bank: A Monetary History of the Bank of Amsterdam*. Cambridge University Press, Cambridge.

Rochet, J.-C. and Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029.

Roughgarden, T. (2020). Transaction fee mechanism design for the Ethereum blockchain: An economic analysis of EIP-1559. arXiv:2012.00854.

Saleh, F. (2021). Blockchain without waste: Proof-of-stake. *Review of Financial Studies*, 34(3):1156–1190.

Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*. W. Strahan and T. Cadell, London.

Spiller, P. T. and Wood, R. O. (1988). Arbitrage during the dollar–sterling gold standard, 1899–1908: An econometric approach. *Journal of Political Economy*, 96(4):882–892.

U.S. Securities and Exchange Commission (2023). Shortening the securities transaction settlement cycle. Final Rule, Release No. 34-96930, SEC, Washington, DC.

Van De Mieroop, M. (2005). The invention of interest: Sumerian loans. In Goetzmann, W. N. and Rouwenhorst, K. G., editors, *The Origins of Value: The Financial Innovations that Created Modern Capital Markets*, pages 17–30. Oxford University Press, New York.

# A Calibration and replication

**Panel generation.** The monthly panel (January 2018 to October 2026, 106 months) is produced by a single script with a fixed random seed. For each system the script specifies settlement time, baseline monthly value, locked collateral, fee and risk rates in basis points, and baseline energy cost. Each month it draws SOFR from a regime path (near-zero to March 2022, a linear tightening through 2022, a plateau near 5.3% in 2023–2024 and a lower plateau thereafter) and an industrial power price; it then applies a 40% trend in volume over the sample with 3% monthly noise, scales proof-of-work energy cost with the power price (and, for Bitcoin, with a rising hash-rate trend), and applies three structural events: system launches (FedNow from July 2023, Pix from November 2020, Solana from 2020), the Ethereum Merge (September 2022: energy cost falls to a residual level and the system is reclassified as proof-of-stake) and the U.S. T+1 transition (May 2024: settlement time halves and clearing-fund collateral rises by a factor of 1.35). Drag is computed with equation (2) on the transaction-window basis.

**Tables and figures.** Tables 1, 2 and B.1 and all numbers in the text are produced by `build_manuscript_tables.py`, which re-derives velocity from settlement time, computes the Limit Line, the intra-class frontiers and the decomposition of equation (8), fits the response surface of Appendix B by least squares (point values only), and writes the tables and number macros used by this document. Figures 1 and 2 are produced by `plot_manuscript_figures.py`, which recomputes the regime envelopes and the Limit Line from the same panel and the static historical exemplars.

**One-command reproduction.** `scripts/reproduce_all.sh` runs, in order: the fixed-seed panel generator (and checks that the regenerated panel is byte-identical to the archived one), the table and number builder, the figure script, the unit tests of the interactive engine, and `latexmk` on this document. Python dependencies and the last verified versions are listed in `requirements.txt`; the optional cross-check `refine_interaction_model.py` refits surface M2 with `statsmodels` and must reproduce the coefficients of Table B.1.

**Equation-to-code map.** Equations (1)–(4) are implemented in `src/lib/gvk/engine.ts`; Propositions 1 and 2 are verified numerically in `src/test/gvk-propositions.test.ts`; the Limit Line and equation (6) in `src/utils/calculateFrontier.ts`; the regime envelopes in `src/lib/gvk/regimes.ts`; the intra-class frontiers and equation (8) in `src/lib/gvk/classFrontier.ts`, with Proposition 3 verified numerically in `src/test/gvk-class-frontier.` and equation (B.1) in `src/lib/gvk/regression.ts` (dashboard) and `scripts/gvk_common.py` (estimation).

**Interactive explorer.** A companion dashboard renders the same frontier interactively, lets the user replay the panel month by month, vary SOFR and power prices in a snapshot mode, switch the collateral accrual basis, and simulate a T+0 transition. The dashboard and the scripts share the same historical coordinates and frontier construction.

# B A parametric response surface of the calibration

This appendix summarises how the calibrated simulation behaves across its design grid of settlement speed, netting and policy rate. It is a *simulation metamodel*, an emulator of a known data-generating process in the sense of the response-surface literature, not an econometric estimate. We therefore report point values only. The surface is a translog (Christensen et al., 1973) in velocity with regime and rate terms, shifted down by its minimum residual as in corrected ordinary least squares (COLS; Aigner et al., 1977; Greene, 1980) so that every simulated observation lies on or above it:

$$\ln \varepsilon_{it} = \beta_0 + \beta_1 \ln v_i + \beta_2 (\ln v_i)^2 + \delta N_i + \gamma_{\text{base}} \ln r_t + \gamma_{\text{cap}} (\ln r_t \times K_i) + e_{it}, \tag{B.1}$$

where $N_i$ indicates a netted wholesale regime and $K_i$ a capital-bound regime (gross wholesale or proof-of-stake).

**Table B.1:** Translog response surface of the calibrated simulation (point values)

| | (M1) Baseline translog | (M2) Rate × capital | (M3) Adds $K_i$ level |
|---|---|---|---|
| Intercept $\beta_0$ | 0.3820 | 0.9941 | 1.9029 |
| $\ln v$ $(\beta_1)$ | 0.0133 | −0.0609 | −0.0788 |
| $(\ln v)^2$ $(\beta_2)$ | 0.0346 | 0.0254 | 0.0231 |
| Netted regime $N_i$ $(\delta)$ | −5.5025 | −6.4132 | −6.5829 |
| $\ln r_t$ $(\gamma_{\text{base}})$ | −0.0128 | −0.1243 | −0.0005 |
| $\ln r_t \times K_i$ $(\gamma_{\text{cap}})$ | | 0.4886 | 0.1304 |
| Capital-bound regime $K_i$ $(\kappa)$ | | | −2.2467 |
| Capital-bound rate slope $\gamma_{\text{base}} + \gamma_{\text{cap}}$ | | 0.3643 | 0.1299 |
| $R^2$ of the surface | 0.6908 | 0.8477 | 0.8687 |
| Min. residual (COLS shift) | −3.1400 | −2.7326 | −2.1125 |
| COLS intercept $\hat{\beta}_0^{\text{COLS}}$ | −2.7579 | −1.7385 | −0.2096 |
| Simulated system-months / systems | 1,466 / 15 | 1,466 / 15 | 1,466 / 15 |

*Notes:* Least-squares fit of equation (B.1) to $\ln \varepsilon_{it}$ (transaction-window basis) over all simulated system-months. $r_t$: SOFR, floored at 0.01%. COLS shifts the intercept by the minimum residual. No standard errors or significance levels are reported: the panel is generated by a known process, so sampling-based inference is not meaningful. Coefficients describe the calibration, not causal effects.

Read as an emulator, the surface reproduces three design features of the calibration. The quadratic term is positive ($\beta_2 = 0.0254$ in M2), so the surface is U-shaped in log velocity with a minimum near $\tau^* = \exp(\beta_1 / 2\beta_2) \approx 0.30\,\text{s}$, a location set by a few very fast systems. The netting term ($\delta = -6.4132$) places netted systems about 610 times below gross systems of equal velocity, which is the calibration's netting assumption returned by construction. The capital-bound rate slope is 0.364 in M2, inside the $[0, 1]$ bound of Proposition 1; adding a level term for capital-bound regimes (M3) reduces it to 0.130, because on this design grid the capital-bound systems differ from the others in level as well as in rate exposure, and a single cross-sectional surface cannot apportion the two. That is a property of the design, not a finding about transmission; Proposition 1 is an identity verified exactly in the engine's tests, and its within-system content is shown in Figure 2.

The interactive explorer uses surface M2 as a second, parametric reference next to the Limit Line. At SOFR of 5%, M2 places the frontier for a non-netted, non-capital-bound system at Visa's velocity at 21.4 bps, against Visa's mean drag of 220.1 bps, a log distance of 2.33. This parametric distance is a property of the emulator; the nonparametric decomposition of Table 2 is the reference used in the main text.