

The Solvency of Appearances

Trust, Concealed Liabilities, Debt Spirals, and the Mathematics of Selective Repayment

Flyxion
Independent Researcher

October 2026

Abstract

A completed repayment is evidence that a payment was completed. This essay asks how far that evidence can be extended to the solvency of the person or institution that made it, and answers that under explicit and realistic assumptions it can be extended remarkably little. The central thesis is that an individual payment can be genuine while the financial structure producing it is insolvent, and that when obligations are spread across creditors who cannot observe one another, selective repayment can manufacture evidence of trustworthiness while aggregate solvency deteriorates. The mechanism does not require fraud. Honest intentions, authentic documents, and ordinary accounting suffice. The essay develops the argument through a sequence of formal results: a non-identifiability theorem for aggregate liabilities and assets given any finite repayment record; a selection result showing that testimony about repayment is biased upward whenever satisfied creditors speak more readily than unsatisfied ones; a refinancing-neutrality identity and a present-value criterion separating sustainable from unsustainable debt; a collective verification failure in which every participant rationally declines to check an appeal that the group would profit from checking; a conservation law for losses in the sale of doubtful claims; a floor on senior-tranche loss probability under a systemic factor; a spectral-radius criterion for amplification of losses through leveraged balance sheets; the coexistence of a calm and a run equilibrium in a solvent but illiquid bank; and a general projection theorem for financial networks. The essay closes by arguing that the danger lies less in borrowing than in the substitution of local evidence for knowledge of aggregate obligations.

Contents

1	Introduction	4
2	Four Quantities and a Vocabulary	5
3	The Invisible Ledger	7

3.1	What a payment record does establish	8
3.2	Time as evidence	9
3.3	What would restore identification	10
4	Selective Repayment and the Production of Reputation	10
4.1	Allocation of scarce cash	10
4.2	Conservation of recovery and the manufacture of satisfied creditors	11
4.3	Reputational survivorship	11
5	Debt Spirals Without Fraud	12
5.1	The accounting identity	12
5.2	The constant-parameter solution	13
5.3	The no-Ponzi condition	14
5.4	Minsky’s taxonomy and the legal notion of a Ponzi scheme	14
6	Authentic Evidence and Fabricated Credibility	15
6.1	Five kinds of evidence	15
6.2	Reused evidence and provenance	16
6.3	Why the five kinds of evidence do not entail one another	17
7	Distributed Informal Credit Origination	17
7.1	The arithmetic of micro-solicitation	17
7.2	Collective verification failure	18
7.3	Messaging platforms, payment rails, and the observability gap	19
8	The Sale of Doubtful Claims	20
8.1	Four valuations of one claim	20
8.2	Conservation of loss	20
8.3	When trade occurs	21
9	Mortgages, Securitization, and the Crisis of 2008	21
9.1	Successive representations of one household cash flow	22
9.2	The senior tranche and the systemic factor	22
9.3	The funding layer	23

10 Leverage, Forced Sales, and the Institutional Spiral	23
10.1 Leverage targeting and forced sales	24
10.2 Price impact and the feedback operator	24
11 Runs and the Collective-Action Problem	26
11.1 A minimal model	26
11.2 Opacity cuts both ways	27
12 A General Theory of Financial Opacity	28
12.1 Networks and projections	28
12.2 Clearing and the dilution of the visible creditor	29
12.3 Projection and closure	29
13 Trust Without Omniscience	30

1 Introduction

Credit is a promise about the future, and a promise about the future can only be evaluated through evidence available in the present. Among the kinds of evidence that lenders find most persuasive, a history of timely repayment ranks near the top. A borrower who returned a small sum on schedule is trusted with a larger one, a household whose cards are always current is offered higher limits, and an institution whose short-term obligations have always been rolled over without incident is assumed to be able to roll them over again. The inference is natural, and in many circumstances it is sound. The purpose of this essay is to identify precisely the circumstances in which it is not, and to measure how much is lost when it fails.

The difficulty can be stated in one sentence. A payment is a local event, whereas solvency is a global property. A creditor observes the payment made to that creditor, together with whatever statements the borrower chooses to provide, but solvency depends on the sum of all obligations the borrower has incurred and on the resources and income available to meet them. When the creditors cannot observe one another, the local record and the global state are connected only by assumptions, and the assumptions are exactly what a stressed or strategic borrower is in a position to violate. The consequence is that a sequence of genuine, on-time, fully documented payments can coexist with a financial structure that is already unable to meet its obligations.

The thesis can be put in the following form. *An individual payment can be genuine while the financial structure that produces it is insolvent. When obligations are distributed across creditors who cannot observe one another, selective repayment can manufacture evidence of trustworthiness even as aggregate solvency deteriorates.* The thesis is deliberately weaker than the claim that such borrowers are fraudsters. It is also more general. Fraud is one route to the outcome, but the same pattern is produced by a household that services one card with another, by a small business that rolls short-term loans, by an individual who solicits modest sums from a very large number of acquaintances, and by a leveraged financial institution whose funding is renewed every night. The intentions in these cases differ enormously. The informational structure is the same.

The essay distinguishes four quantities that are routinely conflated in everyday and institutional talk about credit. *Willingness* to repay is a property of the borrower's disposition. *Ability* to repay is a property of the borrower's resources and income. *Observed repayment* is a property of the record available to a particular creditor. *Sustainability* of the obligation network is a property of the whole collection of claims and the resources standing behind them. The first two are not directly observable by a creditor, the third is observable but local, and the fourth is neither. Most of the errors examined below consist of inferring one of these quantities from another.

A word on scope is necessary, because a treatment that runs from interpersonal borrowing to securitization risks suggesting that personal deception, household insolvency, bank runs, and structured finance are the same phenomenon. They are not. Personal fraud typically depends on concealment, and concealment is what keeps the arrangement alive. A bank run,

by contrast, can destroy an arrangement that disclosure would have preserved, because the damage arises from collective suspicion rather than from concealment. Household insolvency may involve neither deception nor suspicion. Securitization involves legitimate transfers of risk that, under particular modeling errors, hide correlated exposure. What the cases share is a formal feature: a global structure is observed by each participant only through a restricted projection, and the properties that matter are not recoverable from the projections. What they do not share is the dynamics that follow. The essay keeps these two statements apart throughout.

The plan is as follows. Section 2 fixes vocabulary. Section 3 introduces the invisible ledger and proves that local repayment histories do not identify global solvency, then quantifies what they do establish. Section 4 treats selective repayment and the production of reputation. Section 5 treats debt spirals that involve no deception at all, and relates them to Minsky's taxonomy. Section 6 separates the authenticity of evidence from the truth of the claim it supports. Section 7 examines distributed informal credit origination, in which messaging platforms and payment rails allow a single actor to aggregate small transfers from many people who cannot see one another. Section 8 turns to the sale of doubtful claims and proves a conservation law for losses. Section 9 follows a mortgage through its successive representations and isolates the role of correlation. Section 10 develops a model of forced sales and amplification. Section 11 treats runs as a collective-action problem. Section 12 states a general theory of financial opacity in terms of projections of a liability network. Section 13 concludes.

2 Four Quantities and a Vocabulary

Consider a borrower b and a discrete sequence of dates $t = 0, 1, \dots, T$. At each date the borrower holds liabilities to a set of creditors and makes payments to some of them. Let $r_t \geq 0$ denote the per-period interest rate on outstanding debt, let $P_t \geq 0$ denote payments made out of resources other than new borrowing, and let $B_t \geq 0$ denote new borrowing used specifically to fund payments to existing creditors. The gross payment made to creditors at date t is then

$$X_t = P_t + B_t.$$

The decomposition matters because a creditor who receives a payment cannot in general tell which of the two components funded it.

Definition 2.1 (The four quantities). For the borrower b at date t :

Willingness $\omega_t \in [0, 1]$ is the probability that the borrower would make a contractually due payment if the resources to do so were available.

Ability $\alpha_t \geq 0$ is the maximum payment the borrower could make at date t out of non-borrowed resources, namely cash, income, and liquidation of assets at their realizable value.

Observed repayment for creditor i is the record $x_i(0), \dots, x_i(t)$ of payments actually received

by i .

Sustainability is the property that the entire obligation network can be serviced indefinitely, formalized in Section 5 as the no-Ponzi condition.

The four are logically independent in the sense that every combination of qualitative values is realizable. A borrower may be fully willing and entirely unable, in which case sustained observed repayment can occur only through refinancing. A borrower may be able but unwilling, in which case observed repayment may be absent despite solvency. A borrower may be both willing and able to meet the obligation to one creditor while the network as a whole is unsustainable because of liabilities that creditor cannot see. Table 1 records who can observe what.

Quantity	Determined by	Observed by creditor i ?	Inferable from a payment record?
Willingness	disposition of b	no, only through behavior	weakly
Ability	resources and income of b	no	weakly, if funding source known
Observed repayment	payments to i	yes	not applicable
Sustainability	whole liability network	no	not without further restrictions

Table 1: The four quantities and their observability. The final column is the subject of Section 3.

Two accounting identities are used repeatedly. The first is the definition of aggregate liability at date t ,

$$D(t) = \sum_{j \in J} d_j(t),$$

where J is the full set of creditors, visible and hidden, and $d_j(t)$ is the outstanding claim of j . The second is the definition of net worth,

$$W(t) = A(t) - D(t),$$

where $A(t)$ is the realizable value of the borrower's assets and the present value of resources available for debt service is included as appropriate. Solvency at date t means $W(t) \geq 0$ in the present-value sense made precise in Section 5.

3 The Invisible Ledger

Let $I = \{1, \dots, n\} \subseteq J$ denote the set of creditors who are visible to one another only through the borrower's own account. Creditor $i \in I$ observes the information set

$$\mathcal{I}_i(t) = \{d_i, p_i(0), \dots, p_i(t), s_i(t)\},$$

where $p_i(\cdot)$ are the payments received and $s_i(t)$ collects statements, photographs, promises, balances, and other material supplied by the borrower. The aggregate liability $D(t) = \sum_{j \in J} d_j(t)$ is generally not a function of $\mathcal{I}_i(t)$:

$$\mathcal{I}_i(t) \not\Rightarrow D(t).$$

Knowing that someone repaid a loan of one hundred dollars does not establish whether they owe another hundred dollars, ten thousand dollars, or nothing.

A *structure* for the borrower is a specification $\sigma = (M, g, c)$ of the initial liquid assets $M \geq 0$, the hidden draws $g(t) \geq 0$ on credit from lenders outside I , and the hidden outflows $c(t) \geq 0$ such as consumption, losses, or payments to parties outside I . Cash evolves according to

$$m(t) = m(t-1) + g(t) - c(t) - \sum_{i \in I} p_i(t), \quad m(0) = M,$$

and a structure is *feasible* when $m(t) \geq 0$ for every t . Hidden borrowing accrues interest at the rate $r \geq 0$, so the hidden liability at date T is

$$G = \sum_{t=1}^T g(t)(1+r)^{T-t}.$$

A structure is *consistent* with the record $R_i = (d_i, p_i(\cdot), s_i)$ if it is feasible and generates exactly the payments in R_i . Writing $d(T)$ for the visible debt outstanding at T according to the record, the aggregate liability is $D = d(T) + G$ and the assets are $A = m(T)$.

Theorem 3.1 (Non-identifiability of aggregate solvency). *Let R_i be any record with $d(T) \geq 0$ and total payments $X = \sum_t p_i(t)$. For every pair (D, A) with $D \geq d(T)$ and $A \geq 0$ there exists a structure σ consistent with R_i whose aggregate liability at T is D and whose assets at T are A . In particular, both a solvent structure ($A \geq D$) and an insolvent structure ($A < D$ with $D > 0$) are consistent with every record.*

Proof. Fix (D, A) with $D \geq d(T)$ and $A \geq 0$. Set $M = X$, so that the visible payments can be funded entirely out of initial liquid assets, and make no draws before the final date. If $r = 0$ or draws occur only at date T , the hidden liability is simply the amount drawn at T . Let $g(T) = D - d(T) \geq 0$, so that $G = D - d(T)$ and the aggregate liability is D . Cash just before the final date is $M - X = 0$, and at the final date cash equals $g(T) - c(T)$. If $g(T) \geq A$, choose $c(T) = g(T) - A$, giving terminal cash A . If $g(T) < A$, increase the initial

assets to $M = X + A - g(T)$ and set $c(T) = 0$, again giving terminal cash A . In both cases the structure is feasible, reproduces the payments in R_i , and has aggregate liability D and assets A . Taking $A = D$ gives a solvent structure, and taking $A = 0$ with $D > 0$ gives an insolvent one. $\square$

The theorem concerns hidden liabilities incurred at any time, including after the record ends. A second and distinct mechanism concerns the funding of the payments in the record itself.

Proposition 3.2 (Funding-source non-identifiability). *Let R_i be any record with total payments X and suppose the borrower may draw hidden credit of at least X over the horizon. Then for any initial liquid assets $M \geq X$, the record R_i is consistent both with a structure in which every payment is funded from own resources ($g \equiv 0$) and with a structure in which every payment is funded from hidden borrowing ($g(t) = p_i(t)$ for all t , with own resources untouched). In the first, own resources committed to the payments equal X and the hidden liability is zero. In the second, own resources committed equal zero and the hidden liability at T is $G_X = \sum_t p_i(t)(1+r)^{T-t} \geq X$. The difference in net worth between the two structures is $G_X - X$, the interest accrued on the hidden loans.*

Proof. In the second structure the cash recursion gives $m(t) = m(t-1) + g(t) - p_i(t) = m(t-1)$, so cash stays at M and feasibility holds. In the first, $m(t) = M - \sum_{s \leq t} p_i(s) \geq 0$ because $M \geq X$. The record generated is R_i in both cases by construction. Terminal cash is $M - X$ in the first and M in the second, so net worth is $M - X - d(T)$ and $M - d(T) - G_X$ respectively. $\square$

The proposition shows that a payment record cannot reveal the funding source, although refinancing leaves net worth nearly unchanged apart from interest. The reason the distinction nevertheless matters is dynamic: a structure funded by hidden credit is exposed to the interruption of that credit, whereas a structure funded from own resources is not.

These constructions are not pathological. They are the ordinary behavior of a borrower with access to a credit line, a second lender, or a balance-transfer offer. The creditor's information set is identical in all of them.

3.1 What a payment record does establish

The negative result is sharpened by asking what the record does establish when the borrower's access to hidden credit is bounded.

Proposition 3.3 (The evidential value of repayment). *Suppose total hidden borrowing over the horizon is known to be at most $\bar{B}$. Then a record of total payments X implies that the borrower committed own resources of at least $(X - \bar{B})^+$ to those payments, where $y^+ = \max(y, 0)$. This bound is attained, so no stronger inference about own resources is available from the record alone.*

Proof. Payments are funded from own resources or from borrowing, so own-funded payments are at least $X - \bar{B}$ and at least zero. The bound is attained by the structure that borrows the maximum $\bar{B}$ to fund payments and uses own resources for the remainder. $\square$

The quantity $(X - \bar{B})^+$ has a plain interpretation. It is the amount of ability demonstrated by the repayment record. It vanishes whenever $\bar{B} \geq X$, which is the situation of a borrower with a credit line larger than the payments made so far, and it tends to zero as the bound on hidden credit is relaxed. Observed repayment is evidence of ability only to the extent that hidden credit is excluded.

3.2 Time as evidence

A repayment record is not without value, and the value can be stated exactly under a simple model of the insolvent borrower's access to credit. Let $S \in \{0, 1\}$ indicate solvency, with prior $\mathbb{P}(S = 1) = \pi \in (0, 1)$. Suppose that a solvent borrower makes every payment, and that an insolvent borrower has access to hidden credit for a random number of periods τ , after which the first payment is missed. Let $G(t) = \mathbb{P}(\tau > t)$ denote the probability that the insolvent borrower is still paying after t periods.

Proposition 3.4 (Posterior after a payment record). *After t consecutive on-time payments, the posterior probability of solvency is*

$$\mathbb{P}(S = 1 \mid \text{paid through } t) = \frac{\pi}{\pi + (1 - \pi)G(t)},$$

and the odds of solvency are multiplied by $1/G(t)$ relative to the prior. If the insolvent borrower can refinance indefinitely, so that $G \equiv 1$, the posterior equals the prior at every t and the mutual information between the payment record and solvency is zero.

Proof. By Bayes' rule, $\mathbb{P}(S = 1 \mid \text{paid}) = \mathbb{P}(\text{paid} \mid S = 1)\pi / [\mathbb{P}(\text{paid} \mid S = 1)\pi + \mathbb{P}(\text{paid} \mid S = 0)(1 - \pi)]$, with $\mathbb{P}(\text{paid} \mid S = 1) = 1$ and $\mathbb{P}(\text{paid} \mid S = 0) = G(t)$. When $G \equiv 1$ the likelihood of the record is the same under both values of S , so S and the record are independent and their mutual information vanishes. $\square$

If τ is geometric with per-period hazard h , then $G(t) = (1 - h)^t$ and the odds multiplier is $(1 - h)^{-t}$. The record becomes informative only on a time scale comparable to $1/h$. A borrower whose hidden credit lasts for many periods can therefore be indistinguishable from a solvent one for many periods, and a creditor who observes only a handful of payments learns almost nothing. Time is evidence, but it is evidence priced in the units of the borrower's access to credit, which the creditor does not observe.

3.3 What would restore identification

The impossibility can be reversed only by restrictions on borrowing, income, or disclosure. Three are worth naming. A registry of claims that every lender is required to consult and update identifies D directly, provided that the borrower cannot borrow from unregistered lenders. A cap on total borrowing bounds G and thereby yields Proposition 3.3. Verification of the borrower's assets by a trusted third party bounds A . The first two address the liability side and the third addresses the asset side, and it is essential to see that they are separate gaps: even a perfect registry of liabilities does not identify solvency without information about assets, and even perfect information about assets does not identify solvency without information about liabilities. Both are hidden from the individual creditor in the ordinary course.

4 Selective Repayment and the Production of Reputation

A simple arithmetic example fixes ideas. A borrower obtains one hundred dollars from each of one hundred people, collecting

$$C = 100 \times 100 = \$10,000.$$

Suppose that one creditor becomes suspicious and the borrower repays that creditor \$110, a ten percent premium. The cost to the borrower is \$110 out of \$10,000, or 1.1 percent of the funds collected. The creditor, now holding evidence that the borrower repays with interest, has reason to withdraw suspicion and may recommend the borrower to others. The remaining ninety-nine creditors remain unpaid and, because they do not know of one another, may remain silent. The following subsections formalize the two components of this process, the allocation of payments and the production of testimony.

4.1 Allocation of scarce cash

Let creditor i have probability q_i of publicly complaining if unpaid, suppose the complaint would impose reputational damage η_i on the borrower, and let c_i be the cost of satisfying creditor i . Given cash B , a strategic borrower seeking to minimize expected reputational damage chooses a set S of creditors to satisfy, subject to $\sum_{i \in S} c_i \leq B$, so as to maximize $\sum_{i \in S} q_i \eta_i$. The *priority index* of creditor i is

$$\mu_i = \frac{q_i \eta_i}{c_i}.$$

Proposition 4.1 (Priority rule in the fractional relaxation). *Suppose partial payment of creditor i in fraction $x_i \in [0, 1]$ reduces the expected damage by $x_i q_i \eta_i$. Then the optimal allocation of cash B satisfies the rule: order creditors by decreasing μ_i , pay in full along this order, and pay the creditor at which cash is exhausted in proportion to the remaining cash.*

Proof. The problem is the linear program of maximizing $\sum_i x_i q_i \eta_i$ over $x \in [0, 1]^n$ subject to $\sum_i x_i c_i \leq B$. Suppose an optimal x has $x_j < 1$ and $x_k > 0$ with $\mu_j > \mu_k$. Moving a small amount ϵ of cash from k to j changes the objective by $\epsilon(\mu_j - \mu_k) > 0$, contradicting optimality. Hence in an optimal solution no creditor of higher index is underpaid while a creditor of lower index receives anything, which is the stated rule. $\square$

The 0–1 version of the problem, in which each creditor is either paid in full or not, is a knapsack problem, and the ratio rule is then a heuristic rather than an optimum. The point is not that real borrowers solve this program. The point is that any borrower, whether calculating or merely responsive to pressure, will tend to satisfy the loudest, best connected, or most threatening creditors first, and that this tendency is precisely what shapes the testimony available to prospective lenders.

4.2 Conservation of recovery and the manufacture of satisfied creditors

Selective repayment does not create resources. If assets total A and claims total nd with $A < nd$, the aggregate recovery across all creditors is at most A whatever the allocation. What the allocation changes is the distribution of recovery across creditors, and with it the number of creditors who have been made whole.

Proposition 4.2 (Same total, different testimony). *Let n creditors hold equal claims of size d and let the borrower have cash $A < nd$. Under pro rata distribution every creditor receives the fraction $A/(nd) < 1$ of the claim and no creditor is made whole. Under selective full payment, exactly $k = \lfloor A/d \rfloor$ creditors are made whole. The aggregate paid is at most A in both cases.*

Proof. Pro rata gives each creditor $A/n < d$. Selective payment of k creditors in full costs $kd \leq A$, and $k + 1$ would cost more than A . Both allocations distribute at most A by construction. $\square$

The same aggregate payment therefore generates opposite local evidence. Under pro rata distribution, every creditor can truthfully report a shortfall. Under selective payment, a fraction k/n of creditors can truthfully report full repayment, while the remainder bear the entire shortfall.

4.3 Reputational survivorship

Prospective lenders do not observe the full creditor population. They observe whatever testimony reaches them. Let k of n creditors be paid, let paid creditors report their experience with probability g_p , and let unpaid creditors report with probability g_u , independently across creditors. The expected number of positive reports is kg_p , the expected number of negative

reports is $(n - k)g_u$, and the share of positive reports among all reports, approximated by the ratio of expectations, is

$$\varphi = \frac{k g_p}{k g_p + (n - k) g_u}.$$

Proposition 4.3 (Upward bias in reported repayment). *For $0 < k < n$, the share φ of positive reports exceeds the true repayment share k/n if and only if $g_p > g_u$. Moreover, as $g_u/g_p \rightarrow 0$, $\varphi \rightarrow 1$ for every fixed $k \geq 1$.*

Proof. We have $\varphi > k/n$ exactly when $n g_p > k g_p + (n - k) g_u$, that is, when $(n - k)g_p > (n - k)g_u$. Since $n - k > 0$, this is equivalent to $g_p > g_u$. For the limit, divide numerator and denominator of φ by g_p to obtain $\varphi = k/[k + (n - k)(g_u/g_p)]$, which tends to 1 as $g_u/g_p \rightarrow 0$. $\square$

The extreme case is instructive. If a single creditor out of one hundred is paid and the unpaid creditors are silent, every report a prospective lender receives is positive, although the true repayment share is one percent. Nothing in the reports is false. The distortion arises entirely from which experiences become testimony. The effect is the reputational analogue of survivorship bias, and it is amplified by the dynamic fact that positive testimony recruits additional creditors from the same pool of prospective lenders, who in turn replenish the cash used to satisfy the next round of loud creditors. A prospective lender who treats received reports as a random sample of the creditor population will overestimate reliability by exactly the factor identified in Proposition 4.3.

5 Debt Spirals Without Fraud

The preceding sections allowed for strategic behavior. This section removes it. A household or small enterprise that services several credit accounts, drawing on one to pay another, can arrive at insolvency by a path that involves no intention to deceive anyone. The mathematics is the mathematics of compounding.

5.1 The accounting identity

Let D_t be total outstanding debt at date t , r_t the effective rate, $C_t \geq 0$ new borrowing for purposes other than repayment, $B_t \geq 0$ new borrowing used to repay existing creditors, and $P_t \geq 0$ payments out of non-borrowed resources. Gross payments to creditors are $X_t = P_t + B_t$. Then

$$D_{t+1} = (1 + r_t)D_t - X_t + B_t + C_t = (1 + r_t)D_t + C_t - P_t.$$

Proposition 5.1 (Refinancing neutrality). *The trajectory of aggregate debt does not depend on B_t : it is determined by r_t , C_t , and P_t alone. Borrowing from one lender to repay another leaves aggregate liabilities unchanged, while the gross payments $X_t = P_t + B_t$ visible to creditors may exceed P_t by an arbitrary amount.*

Proof. Substituting $X_t = P_t + B_t$ into the accounting identity cancels B_t exactly. $\square$

The proposition has a direct evidentiary consequence. A creditor who sees payments X_t arriving on schedule cannot infer P_t , because B_t is invisible. Timely payment to every account is entirely consistent with $P_t = 0$.

5.2 The constant-parameter solution

Specialize to constant $r > 0$, constant non-borrowed payment $P > 0$, and $C_t = 0$. The recursion is

$$D_{t+1} = (1 + r)D_t - P,$$

with fixed point $D^* = P/r$ (since $D^* = (1 + r)D^* - P$ gives $rD^* = P$). The solution is

$$D_t = (1 + r)^t \left(D_0 - \frac{P}{r} \right) + \frac{P}{r}.$$

This is verified by substitution: $(1 + r)D_t - P = (1 + r)^{t+1}(D_0 - P/r) + (1 + r)P/r - P = (1 + r)^{t+1}(D_0 - P/r) + P/r = D_{t+1}$.

Proposition 5.2 (Three regimes). *In the constant-parameter model:*

(a) *if $D_0 < P/r$, the debt is extinguished in finite time, at the first t satisfying*

$$t \geq \frac{\ln \left(\frac{P/r}{P/r - D_0} \right)}{\ln(1 + r)};$$

(b) *if $D_0 = P/r$, the debt is constant and the borrower pays interest forever without reducing principal;*

(c) *if $D_0 > P/r$, the debt grows without bound, $D_t \rightarrow \infty$, even though the borrower makes the payment P in every period.*

Proof. Set $D_t = 0$ in the closed form to obtain $(1 + r)^t(P/r - D_0) = P/r$, which has the stated solution in case (a). In case (b) the coefficient of $(1 + r)^t$ vanishes. In case (c) the coefficient is positive and $(1 + r)^t \rightarrow \infty$. $\square$

The important condition is not whether the borrower is paying something but whether repayment capacity is sufficient to prevent the outstanding obligation from growing. In regime (c) the sequence of payments is regular, the account statements are current, and the aggregate position deteriorates geometrically.

Real constraints interrupt the trajectory. Let L be the aggregate credit limit across all accounts. In regime (c), with $L > D_0 > P/r$, debt reaches the limit at

$$t_L = \frac{\ln\left(\frac{L-P/r}{D_0-P/r}\right)}{\ln(1+r)}.$$

Up to t_L , every creditor observes a borrower who pays on time. After t_L , further refinancing is impossible and the borrower defaults on whichever obligations cannot be met. The interruption can be financially devastating without implying that the borrower ever intended to defraud anyone, and its timing is determined by the aggregate limit and the rate, neither of which any single lender observes.

5.3 The no-Ponzi condition

Dividing the identity $D_{t+1} = (1+r)D_t - P_t$ by $(1+r)^{t+1}$ gives $\frac{D_{t+1}}{(1+r)^{t+1}} = \frac{D_t}{(1+r)^t} - \frac{P_t}{(1+r)^{t+1}}$, and therefore

$$\frac{D_T}{(1+r)^T} = D_0 - \sum_{s=0}^{T-1} \frac{P_s}{(1+r)^{s+1}}.$$

Definition 5.3 (Sustainability). The obligation network is *sustainable* if

$$\limsup_{T \rightarrow \infty} \frac{D_T}{(1+r)^T} \leq 0, \quad \text{equivalently} \quad D_0 \leq \sum_{s=0}^{\infty} \frac{P_s}{(1+r)^{s+1}}.$$

The right-hand side is the present value of repayment capacity from non-borrowed resources. Sustainability is therefore the statement that the initial aggregate debt does not exceed the present value of the resources that can be devoted to it. It makes no reference to the amount of gross payments X_t , and a borrower can satisfy every payment in the record while violating it. For constant P the present value is P/r , and the condition reduces to regime (a) or (b) of Proposition 5.2.

5.4 Minsky's taxonomy and the legal notion of a Ponzi scheme

Minsky's financial instability hypothesis classifies borrowers by the relation between cash flows and payment commitments [14]. Let $I_t = rD_t$ be interest due and let A_t denote scheduled principal amortization. A *hedge* unit has $P_t \geq I_t + A_t$. A *speculative* unit has $I_t \leq P_t < I_t + A_t$, so that interest is covered but principal must be refinanced. A *Ponzi* unit has $P_t < I_t$, so that even interest must be financed by new borrowing. With no further borrowing for other purposes, $D_{t+1} - D_t = rD_t - P_t$, so a Ponzi unit has strictly growing debt, and since rD_t is then increasing, a Ponzi unit remains Ponzi unless its payments rise. In the constant-parameter model, regime (c) is exactly the Ponzi regime.

It is essential not to conflate Minsky's notion with the legal and colloquial notion associated with a particular fraudulent scheme [3]. In the latter, investors are promised returns and paid from the contributions of later investors, and the operator knows that no underlying activity generates the returns. Minsky's notion describes a cash-flow relation that can

arise without any deception, for instance when a firm's operating income falls below its interest burden. The two share a funding structure but differ in knowledge and intent. The mathematics of the next proposition applies to both.

Proposition 5.4 (Collapse of a pure rollover scheme). *Consider a scheme that promises each existing creditor a return $\rho > 0$ per period on outstanding principal, has no external income, and pays every return out of new contributions. Let D_t be principal outstanding and K the total amount that the pool of prospective contributors is willing to contribute over the scheme's life. Then $D_{t+1} = (1 + \rho)D_t$, every promised payment is made on time for $t < t^*$, where*

$$t^* = \frac{\ln(1 + K/D_0)}{\ln(1 + \rho)},$$

and the scheme cannot meet its promises at the first date after t^* .

Proof. New contributions at date t must equal the return due, $N_t = \rho D_t$, and they add to principal, so $D_{t+1} = D_t + N_t = (1 + \rho)D_t$ and $D_t = (1 + \rho)^t D_0$. Cumulative contributions through date $t - 1$ equal $D_t - D_0$ and cannot exceed K . Hence $(1 + \rho)^t D_0 \leq D_0 + K$, which holds exactly for $t \leq t^*$. $\square$

During the entire interval $t < t^*$ the observed repayment rate is one hundred percent for every creditor. The record of on-time payment is perfect and says nothing about the date at which the arithmetic of recruitment will fail.

6 Authentic Evidence and Fabricated Credibility

A creditor's confidence rests on evidence, and the previous sections show that the evidence is local. This section examines the logic of evidence itself, in particular the separation between the authenticity of an item and the truth of the claim it is offered to support. A photograph may be genuine, but its claimed provenance may be false. A bank statement may accurately show a payment without establishing where the payment originated.

6.1 Five kinds of evidence

Consider five propositions about a borrower or claimant: that a payment was made, that the claimant has the identity asserted, that the claimant has the need asserted, that the claimant is solvent, and that the representation is honest. Encode a possible world as $\omega \in \{0, 1\}^5$ with these five coordinates. Evidence of type k is a function $e_k(\omega) = \omega_k$, possibly noisy.

Lemma 6.1 (Evidence passes only through dependence). *Let E be evidence whose distribution depends on the world only through a coordinate Z , and let H be another coordinate. Then*

$$\mathbb{P}(H \mid E) = \sum_z \mathbb{P}(H \mid Z = z) \mathbb{P}(Z = z \mid E).$$

In particular, if H and Z are independent under the prior, then $\mathbb{P}(H \mid E) = \mathbb{P}(H)$ for every E .

Proof. Conditional independence of E and H given Z yields $\mathbb{P}(H \mid E) = \sum_z \mathbb{P}(H, Z = z \mid E) = \sum_z \mathbb{P}(H \mid Z = z, E) \mathbb{P}(Z = z \mid E) = \sum_z \mathbb{P}(H \mid Z = z) \mathbb{P}(Z = z \mid E)$. If H and Z are independent, $\mathbb{P}(H \mid Z = z) = \mathbb{P}(H)$ for each z , and the sum collapses to $\mathbb{P}(H)$. $\square$

Taking Z to be payment and H to be solvency, the lemma says that evidence of payment updates belief about solvency only to the extent that payment and solvency are correlated in the relevant population. Proposition 3.4 identifies the source of the weak correlation: access to credit lets the insolvent pay. The likelihood ratio for solvency from a record of payment is

$$\Lambda = \frac{\mathbb{P}(\text{paid} \mid S = 1)}{\mathbb{P}(\text{paid} \mid S = 0)},$$

and it equals one for as long as the insolvent borrower can refinance. Repeated payments multiply this ratio, so $\Lambda^t = 1$ throughout.

6.2 Reused evidence and provenance

Suppose several purportedly independent claimants submit the same photograph in support of separate appeals. Let I denote the hypothesis that the claimants are genuinely independent and each photograph originates from the claimant's own event, and let C denote the hypothesis that the images share a common source. Let $\epsilon = \mathbb{P}(\text{identical image} \mid I)$ and $\kappa = \mathbb{P}(\text{identical image} \mid C)$.

Proposition 6.2 (Posterior odds of a common source). *Upon observing that two supposedly independent images are identical, the posterior odds of a common source are*

$$\frac{\mathbb{P}(C \mid \text{dup})}{\mathbb{P}(I \mid \text{dup})} = \frac{\kappa}{\epsilon} \cdot \frac{\mathbb{P}(C)}{\mathbb{P}(I)}.$$

Proof. This is Bayes' rule in odds form, the posterior odds being the product of the likelihood ratio and the prior odds. $\square$

When ϵ is small, the likelihood ratio κ/ϵ is large even if the prior probability of C is small. The inference is nonetheless restricted. The hypothesis C is a hypothesis about *provenance*, not about *coordination*. A common source is consistent with a single operator running many appeals, with several independent opportunists who copied the same stock or scraped image, with a template circulated through a community, and with a legitimate shared event photographed once and honestly redistributed. The observation undermines the claimed provenance of the images without establishing which of these explanations holds. The posterior is also sensitive to ϵ , which depends on the context, since some genuine events do generate widely shared images. A careful evaluation reports the dependence on ϵ rather than a single number.

6.3 Why the five kinds of evidence do not entail one another

For each observed subset $K \subseteq \{1, \dots, 5\}$ of coordinates, the set of worlds consistent with noiseless evidence on K has $2^{5-|K|}$ elements and contains worlds with both values of every unobserved coordinate. Evidence of payment is consistent with insolvency and with dishonest representation, evidence of identity is consistent with the absence of need, and so on. The practical content is that a document verifying one coordinate is routinely treated as if it verified several. A bank statement verifies payment, a photograph verifies that a photograph exists, and a favorable credit history verifies a particular history of obligations. None independently establishes the borrower's future financial position, and Lemma 6.1 says how much each does: exactly as much as its coordinate is statistically tied to the one in question.

7 Distributed Informal Credit Origination

The structures considered so far assumed that the borrower is a household, firm, or institution with some institutional footprint. Modern communication systems permit a different configuration, in which a single actor approaches a very large number of financially independent people through private messages, each of whom is asked for a sum small enough that investigation seems unwarranted. The configuration can be called *distributed informal credit origination*. The label is a convenience, since not every transfer is technically a loan. Some are gifts, some are advances, some are fraudulent solicitations, and some are genuine requests for emergency assistance. What they share is that the solicitation mechanism permits one actor to approach many people without making the aggregate solicitation visible to any of them.

7.1 The arithmetic of micro-solicitation

Let N be the number of recipients, p the probability that any one recipient responds, and a the requested amount. Under independent responses the expected receipts and the variance are

$$\mathbb{E}[R] = Npa, \quad \text{Var}(R) = Np(1-p)a^2.$$

With $N = 5,000$ and $a = \$5$, the maximum possible receipts are \$25,000, attained if every recipient responds. If sending an additional request costs c_m , the expected net return from the campaign is $N(pa - c_m)$, which is positive whenever $pa > c_m$. Because c_m is close to zero for messaging platforms, the condition is met for very small p , so the borrower does not need a convincing story for every recipient, only a sufficiently persuasive one for a small fraction.

The structure inverts that of banking. A conventional bank aggregates funds from many depositors and extends credit under a framework of underwriting, capital requirements, accounting, and supervision. In the private-message configuration the borrower aggregates funds from many individuals while the individuals independently perform the role of lenders without knowing that they constitute a creditor network. The qualification worth preserving

is that banks do not have perfect knowledge of borrowers' obligations, and deposit insurance generally protects eligible deposits and not the bank's own lending losses. Regulated lenders nevertheless have access to credit reporting, risk models, contractual disclosures, and legal recovery procedures that an ordinary person responding to a private message generally lacks.

The structural feature can be measured. Define the *visibility ratio* of a participant as the fraction of aggregate exposure that the participant can observe. A recipient who contributes a out of aggregate contributions of expected size Npa sees a fraction

$$\nu = \frac{a}{Npa} = \frac{1}{Np}$$

of the exposure. The aggregation occurs at the recipient of the money, while the information remains fragmented among the people providing it.

7.2 Collective verification failure

The second feature is that verification itself has a cost. Suppose the appeal is identical for all recipients and is either legitimate or fraudulent, with $\mathbb{P}(\text{fraudulent}) = \phi$ common to all. Each recipient independently would be willing to send a with probability p . A recipient who verifies at personal cost c_i detects fraud with probability d and withholds if fraud is detected. For a would-be sender, the expected loss avoided by verifying is $d\phi a$.

Proposition 7.1 (Collective verification failure). *(i) A would-be sender verifies privately if and only if $c_i < d\phi a$.*

(ii) If a single verification can be shared with all other recipients at total sharing cost s , the expected loss avoided across the group is $Npd\phi a$, so sharing is collectively worthwhile if and only if $c + s < Npd\phi a$.

(iii) Consequently, for every verification cost c satisfying

$$d\phi a \leq c < Npd\phi a - s,$$

no recipient verifies although verification is collectively profitable. The interval is non-empty if and only if $(Np - 1)d\phi a > s$.

Proof. Part (i) is the comparison of the cost of verification with its expected benefit. For (ii), each of the expected Np would-be senders avoids an expected loss $d\phi a$ if the result of one verification is shared with them, and the total cost is the verification cost c plus the sharing cost s . Part (iii) combines the two: individual verification fails when $c \geq d\phi a$, while collective verification succeeds when $c < Npd\phi a - s$. The interval is non-empty exactly when the lower end is below the upper end. $\square$

The proposition is a statement of the public-goods structure identified by Olson [16]. It also connects to the classical observation that rational ignorance can be the individually

correct response to small stakes [5] and to the economics of limited attention [17]. Someone might reasonably decide that investigating a five-dollar appeal for half an hour is not worthwhile. But five thousand people making the same decision collectively transfer a large sum without meaningful scrutiny, and the ratio between the collective and individual benefit of verification is Np . For Np in the hundreds or thousands, the interval of costs for which verification is collectively worthwhile and individually unjustified spans two to three orders of magnitude. The interval is widened further by the solicitor’s control of the channel. Private messages keep recipients disjoint, which raises the sharing cost s , since recipients must first discover that other recipients exist.

7.3 Messaging platforms, payment rails, and the observability gap

Two distinct kinds of service participate in the configuration. Messaging platforms supply the communication layer in which the solicitation occurs. Payment and remittance services supply the rails on which the transfer occurs. We should be precise about responsibility. Neither kind of provider is inherently facilitating fraud, and the providers publish fraud-prevention guidance concerning emergency appeals, impersonation, and relationships that culminate in requests for money. The structural point is informational. A payment provider has visibility into transactions on its own network without necessarily seeing the conversation that motivated them, while a messaging provider hosts the solicitation without seeing the complete financial activity across payment services.

Let $\mathbf{M}$ denote the set of messages, $\mathbf{T}$ the set of transfers, and $\Gamma \subseteq \mathbf{M} \times \mathbf{T}$ the unobserved relation recording which transfers were prompted by which messages. The messaging provider observes the projection $\pi_{\mathbf{M}}(\mathbf{M}, \mathbf{T}, \Gamma) = \mathbf{M}$, the payment provider observes $\pi_{\mathbf{T}}(\mathbf{M}, \mathbf{T}, \Gamma) = \mathbf{T}$, and the participants observe their own incident messages and transfers.

Proposition 7.2 (Necessity of a join). *The number of transfers to a payee s that were prompted by solicitations from s is not a function of $(\pi_{\mathbf{M}}, \pi_{\mathbf{T}})$ alone. There exist two worlds with identical $\mathbf{M}$ and $\mathbf{T}$ in which this number equals 0 and $|\mathbf{T}_s|$ respectively, where $\mathbf{T}_s$ is the set of transfers to s .*

Proof. Take $\mathbf{M}$ to contain a large set of messages from s to many recipients and $\mathbf{T}$ to contain many small transfers to s . In the first world let $\Gamma = \emptyset$, so the transfers are unrelated to the messages, as they would be if s were simply a person who receives small gifts for independent reasons. In the second let Γ link every transfer to a message. The projections $\pi_{\mathbf{M}}$ and $\pi_{\mathbf{T}}$ coincide across the two worlds, while the attributed count differs. $\square$

The proposition does not say that providers lack fraud controls. It says that detecting a pattern of repeated solicitation can require information held by different organizations and by private individuals, and that a provider’s controls must therefore partly infer Γ from correlates such as timing and account behavior. The empirical scale of the underlying problem is considerable. The Federal Trade Commission reported that in 2025 consumers reported

losses of \$2.1 billion to scams that began on social media, with scams originating on Facebook accounting for the largest losses among identified platforms [9]. That figure measures social-media fraud generally, and it counts only reported losses. It is not a measurement of the particular micro-solicitation mechanism analyzed here, and the proposal is not that every request for small sums is a scam. It is the more general and more defensible statement that *modern communication systems permit the aggregation of financial claims without the aggregation of financial knowledge*.

8 The Sale of Doubtful Claims

Institutional lenders confront the same informational problem at larger scale, and one of their responses is to sell claims they expect to collect only in part. Suppose a creditor holds \$10 million of loans but expects to collect only a fraction. The creditor may sell the claims to a debt purchaser at a substantial discount. The purchaser does not need every borrower to repay. It needs recoveries, net of collection expenses, to exceed the acquisition price and the required return.

8.1 Four valuations of one claim

A portfolio carries at least four numbers, which institutions often treat as interchangeable. The *face value* F is the legal claim. The *expected recovery* is qF , where $q \in [0, 1]$ is the expected recovery fraction. The *market price* P is what a purchaser pays. The *carrying value* B_v is the figure on the holder's books. Let K denote collection costs. The purchaser's expected profit is

$$\mathbb{E}[\Pi] = qF - K - P.$$

For example, a portfolio with \$10 million face value might sell for \$800,000. If the purchaser eventually recovers \$1.5 million and spends \$400,000 collecting it, the profit is $1,500,000 - 400,000 - 800,000 = \$300,000$. This is not intrinsically fraudulent or evidence of a Ponzi scheme. Distressed-debt purchasing can be an economically legitimate transfer of risk. A \$10,000 obligation remains a \$10,000 legal claim even if the probability-weighted recoverable amount is \$2,000, and the discrepancy $F - qF$, the *nominal-economic wedge*, matters whenever face values, estimated recoveries, market prices, and accounting values are treated as interchangeable.

8.2 Conservation of loss

The sale of a claim moves the loss between parties without changing its size. Let the originator have advanced principal $\mathcal{L}$ in total. Suppose the originator sells at price P , and the purchaser later recovers X at collection cost K .

Proposition 8.1 (Conservation of loss). *The originator’s net result is $P - \mathcal{L}$, the purchaser’s net result is $X - K - P$, and their sum is*

$$X - K - \mathcal{L},$$

which does not depend on the price P . The price determines only how the aggregate loss $\mathcal{L} - X + K$ is divided.

Proof. Adding the two net results cancels P . □

The aggregate loss is determined by the borrowers’ repayment capacity, represented by X , and by collection technology, represented by K . Neither the original valuation of the portfolio nor the price at which it changes hands alters either. This framing answers the question that the earlier sections raise at institutional scale: who absorbed the write-down? If the originator’s carrying value B_v exceeds P , the difference $B_v - P$ is a loss that was latent in the books and becomes realized at the sale. A creditor of the originator who sees B_v and not P observes the asset at a value that no marketplace supports.

8.3 When trade occurs

Let the originator value the portfolio at $q_o F - K_o$ from its own expected recovery and collection costs, and let it attach a value $\psi \geq 0$ to relief from capital or administrative burden. Let the purchaser value it at $q_p F - K_p$. A mutually acceptable price exists if and only if

$$q_o F - K_o - \psi \leq P \leq q_p F - K_p,$$

and the interval is non-empty exactly when $q_p F - K_p \geq q_o F - K_o - \psi$. Trade therefore requires that the purchaser have lower collection cost, a higher estimate of recovery, or that the originator value the relief. None of these involves a deception. The informational point is that the market price is informative about the purchaser’s estimate q_p , and the existence of a large gap between price and carrying value is information about the extent to which the nominal valuation of the claims was ever an estimate of recovery. The wedge between nominal and economic value points toward a theory of claims whose nominal values survive after their economic values have deteriorated. The mortgage chain of the next section is the case in which the claims are numerous, repackaged, and used as collateral.

9 Mortgages, Securitization, and the Crisis of 2008

The financial crisis of 2008 was not simply banks selling bad debts to collection agencies. The more consequential mechanisms included mortgage origination, securitization, leveraged holdings of mortgage-related securities, dependence on short-term financing, and cascading losses of confidence. Official retrospectives discuss opacity, leverage, interconnectedness, and

deteriorating mortgage-backed securities as amplifiers [7, 8], and Bernanke’s testimony emphasizes runs on short-term funding within the shadow banking system [1, 11]. This section does not attempt a history. It isolates three mathematical features of the chain that are separable from the history: conservation of underlying cash flow, the role of correlation in the senior claim, and the second layer of fragility introduced by short-term funding.

9.1 Successive representations of one household cash flow

A mortgage can be followed through its successive representations. It is first a household obligation, then an asset held by a lender, then a claim transferred into a mortgage pool, then a security sold to investors, and potentially collateral supporting further borrowing. The same underlying household cash flow can support a complicated structure of contractual claims. Transferring or repackaging those claims does not create additional household repayment capacity.

Lemma 9.1 (Pathwise conservation). *Let a pool have random total recovery $X \geq 0$ and let tranches partition the interval $[0, F]$ by attachment points $0 = a_0 < a_1 < \dots < a_m = F$, tranche k receiving $\min(X, a_k) - \min(X, a_{k-1})$. Then the tranche payoffs sum to $\min(X, F) \leq X$ in every state of the world. In the two-tranche case with senior promise S , the senior tranche receives $\min(X, S)$ and the junior $(X - S)^+$, and their sum is X when $X \leq F = S + (\text{junior promise})$ and is capped at F otherwise.*

Proof. The sum telescopes: $\sum_{k=1}^m [\min(X, a_k) - \min(X, a_{k-1})] = \min(X, a_m) - \min(X, a_0) = \min(X, F)$. For two tranches, $\min(X, S) + (X - S)^+ = X$. $\square$

More generally, if each layer of a chain pays a nondecreasing function of the previous layer’s payoff that never exceeds it, so that $Y_{k+1} = \varphi_k(Y_k)$ with $0 \leq \varphi_k(y) \leq y$, then $Y_k \leq X$ for every k . Each layer is a claim on a share of the same cash flow. The total *notional* of the claims across the layers can exceed the notional of the underlying pool, since a security may itself be pledged as collateral, but no individual layer can pay more than the underlying produces. The fragility arises when the same underlying value is pledged to several claimants in a way that makes their joint payoffs exceed it in aggregate, and in the correlated behavior of the layers in a bad state.

9.2 The senior tranche and the systemic factor

Consider a pool of n loans of unit face value, each paying nothing if it defaults and one if it does not, so that the pool recovers $X = n - K_d$, where K_d is the number of defaults. A junior tranche of size δn absorbs the first δn units of loss, and the senior tranche is promised $S = (1 - \delta)n$. The senior tranche loses if and only if $K_d > \delta n$.

If defaults are independent with common probability $p < \delta$, Hoeffding’s inequality [12] gives

$$\mathbb{P}(K_d > \delta n) \leq \exp(-2n(\delta - p)^2),$$

which falls geometrically in n . This is the standard statement of why a senior claim on a large diversified pool appears safe. Now allow a systemic factor. Let $\Theta \in \{0, 1\}$ with $\mathbb{P}(\Theta = 1) = \theta > 0$, and suppose that conditional on $\Theta = j$ defaults are independent with probability p_j , where $p_0 < \delta < p_1$.

Theorem 9.2 (A floor on senior-tranche loss). *Under the mixture model above,*

$$\lim_{n \rightarrow \infty} \mathbb{P}(K_d > \delta n) = \theta.$$

In particular, the senior-tranche loss probability does not tend to zero with pool size, and an analysis that assumes independence reports a loss probability that falls exponentially in n while the true loss probability is bounded below by approximately θ .

Proof. Conditional on $\Theta = j$, K_d/n is a mean of n independent Bernoulli variables with mean p_j . By Chebyshev's inequality, $K_d/n \rightarrow p_j$ in probability, so $\mathbb{P}(K_d > \delta n \mid \Theta = 0) \rightarrow 0$ because $p_0 < \delta$, and $\mathbb{P}(K_d > \delta n \mid \Theta = 1) \rightarrow 1$ because $p_1 > \delta$. Averaging over Θ gives the limit $(1 - \theta) \cdot 0 + \theta \cdot 1 = \theta$. $\square$

The theorem contains the substance of the claim that supposedly well-protected claims can deteriorate together. A larger pool increases the precision of the estimate of the conditional default rate while leaving the uncertainty about the state of the systemic factor untouched. Confidence in the diversification is therefore increased exactly when the exposure to the factor is unchanged. The result is the structured-finance counterpart of the earlier observation about repayment records: more observations of a local quantity produce more confidence about a global one, and the confidence is misplaced when the global quantity depends on a variable the local observations do not reveal. A discussion of the empirical evidence on structured-finance ratings may be found in [2].

9.3 The funding layer

The crucial extension is that financial institutions may also finance their holdings with short-term debt, which introduces a second layer of fragility. An asset can retain substantial eventual recovery value while its holder runs out of immediately available cash. Let an institution hold assets with eventual value V and fire-sale value ℓV with $\ell < 1$, financed by short-term liabilities L with $\ell V < L \leq V$. The institution is solvent in the sense that eventual assets cover liabilities, but it cannot meet a simultaneous demand for repayment of L without liquidating at the discount. Section 11 analyzes the resulting game. Section 10 analyzes the dynamic in which liquidation itself reduces asset values.

10 Leverage, Forced Sales, and the Institutional Spiral

Let a financial institution have assets A_t , liabilities L_t , and equity $E_t = A_t - L_t$. Its asset-to-equity leverage is $\lambda_t = A_t/E_t$. An institution holding \$100 billion in assets against \$95

billion in liabilities has equity of \$5 billion and leverage 20. Table 2 shows the effect of a decline in asset values with liabilities initially unchanged.

	Initial	After 1% loss	After 5% loss
Assets	100	99	95
Liabilities	95	95	95
Equity	5	4	0
Leverage A/E	20	24.75	undefined

Table 2: An illustrative balance sheet, in billions. A five percent decline in asset value eliminates the equity cushion.

10.1 Leverage targeting and forced sales

The static arithmetic understates the problem, because an institution that targets leverage must respond to a loss by selling assets. Suppose institution j holds assets $A_j = \sum_k h_{jk}$ across asset classes k , has leverage $\lambda_j = A_j/E_j$, and keeps leverage fixed. Let x_k be the fractional price decline of class k . The mark-to-market loss is $\ell_j = \sum_k h_{jk}x_k$, which reduces assets and equity by the same amount.

Lemma 10.1 (Sales needed to restore leverage). *To return to leverage λ_j after a loss ℓ_j , with the proceeds of sales used to repay debt, the institution must sell assets of value*

$$Q_j = (\lambda_j - 1)\ell_j.$$

Proof. After the loss, assets are $A_j - \ell_j$ and equity is $E_j - \ell_j$. Selling Q and repaying debt reduces assets by Q and leaves equity unchanged, so the target requires $(A_j - \ell_j - Q)/(E_j - \ell_j) = \lambda_j$. Solving, $Q = A_j - \ell_j - \lambda_j E_j + \lambda_j \ell_j = (\lambda_j - 1)\ell_j$, using $A_j = \lambda_j E_j$. $\square$

In the example of Table 2, a decline of one percent in asset values produces a loss of \$1 billion and forced sales of $(20 - 1) \times 1 = \$19$ billion, nineteen percent of the balance sheet, in response to a one percent price change. The sales are the mechanism by which a small fundamental loss becomes a large flow of supply.

10.2 Price impact and the feedback operator

Suppose the sales of institution j are spread across asset classes in proportion to holdings, and that the price decline in class k induced by total sales of that class is, to first order, κ_k times the value sold, where $1/\kappa_k$ is the market depth. Let $x(t) \in \mathbb{R}_{\geq 0}^K$ be the vector of incremental price declines in round t . The sales of class k in round t are $\sum_j Q_j(t)h_{jk}/A_j$, and the declines in the next round are

$$x_k(t+1) = \kappa_k \sum_j (\lambda_j - 1) \frac{h_{jk}}{A_j} \sum_{k'} h_{jk'} x_{k'}(t),$$

which is the linear system

$$x(t+1) = Mx(t), \quad M_{kk'} = \kappa_k \sum_j w_j h_{jk} h_{jk'}, \quad w_j = \frac{\lambda_j - 1}{A_j} \geq 0,$$

with an exogenous initial shock $x(0) = u$. The system is a first-order linearization, valid when price declines are small relative to asset values, and it is the stylized form of the feedback loop in which declining equity induces liquidation which causes further valuation losses. In the notation of a scalar model, $A_{t+1} = A_t - \ell_t - \alpha Q_t$ with Q_t increasing as equity falls.

Lemma 10.2 (Spectrum of the feedback operator). *Write $M = \text{diag}(\kappa)H$ with $H = \sum_j w_j h_j h_j^\top$ symmetric and positive semidefinite, $h_j = (h_{jk})_k$. Then M is similar to the symmetric positive semidefinite matrix $\text{diag}(\kappa)^{1/2} H \text{diag}(\kappa)^{1/2}$, so its eigenvalues are real and nonnegative and $\rho(M)$ equals the largest of them.*

Proof. Let $\mathcal{K} = \text{diag}(\kappa)$ with positive entries. Then $\mathcal{K}^{-1/2} M \mathcal{K}^{1/2} = \mathcal{K}^{1/2} H \mathcal{K}^{1/2}$, which is symmetric and positive semidefinite because $H = \sum_j w_j h_j h_j^\top$ with $w_j \geq 0$. $\square$

Theorem 10.3 (Amplification and instability). *Let $u \geq 0$ be an initial shock vector.*

(i) *If $\rho(M) < 1$, the cumulative price decline $\sum_{t \geq 0} x(t) = (I - M)^{-1}u$ is finite and satisfies $(I - M)^{-1}u \geq u$ componentwise, with amplification given by the Neumann series $\sum_{t \geq 0} M^t u$.*

(ii) *If M is irreducible, $\rho(M) \geq 1$, and $u \neq 0$, then $x(t)$ does not converge to zero.*

Proof. (i) Since $\rho(M) < 1$ the Neumann series $\sum_t M^t$ converges to $(I - M)^{-1}$, and since $M \geq 0$ each term is nonnegative with the zeroth term equal to the identity, so $(I - M)^{-1}u \geq u$. (ii) By the Perron–Frobenius theorem for irreducible nonnegative matrices [13], there is a strictly positive left eigenvector w with $w^\top M = \rho(M)w^\top$. Then $w^\top x(t) = \rho(M)^t w^\top u \geq w^\top u > 0$, which is bounded away from zero, so $x(t) \not\rightarrow 0$. $\square$

For a single asset class held by m identical institutions of size A and leverage λ , the operator reduces to the scalar

$$g = \kappa(\lambda - 1)s, \quad s = mA,$$

where s is the total value held by leveraged institutions. The system is stable if and only if $g < 1$, that is, if and only if leverage is below the critical value

$$\lambda^* = 1 + \frac{1}{\kappa s}.$$

The cumulative decline is $u/(1 - g)$ for a shock u , so a shock of one percent produces a total decline of $1/(1 - g)$ percent. The equity of an institution is exhausted when the cumulative

decline reaches $1/\lambda$, so the smallest initial shock that causes failure is

$$u^* = \frac{1-g}{\lambda}.$$

When $g = 0$ this is the static threshold of five percent for leverage 20. As $g \rightarrow 1$, the threshold tends to zero. A financial institution can therefore experience a collapse even when the initial losses seem relatively modest, and the failure is produced by the interaction among leverage, funding maturity, creditor behavior, and market liquidity, none of which is visible on a single balance sheet. In particular κ and s are properties of the market, not of any one institution, and the stability of one institution depends on the leverage of the others. A creditor assessing a single institution from its own statements observes λ and does not observe g .

When $g > 1$ the linear system diverges until nonlinearities intervene: institutions fail, sales cease because there is nothing left to sell, or authorities intervene. The linearization is informative about the existence and direction of the instability and uninformative about its terminal size.

11 Runs and the Collective-Action Problem

At the interpersonal scale, an isolated creditor may be unable to establish the borrower's financial position. At the institutional scale, creditors face uncertainty about the position of the institution and also about the behavior of the other creditors. This introduces a distinct failure mode. A solvent institution holding long-term assets may be unable to meet simultaneous short-term withdrawals without liquidating assets at distressed prices, and the beliefs that cause the withdrawals can be self-fulfilling.

11.1 A minimal model

Following the structure of Diamond and Dybvig [4], consider a continuum of depositors, each holding a claim of one unit that is payable on demand. The bank has invested one unit per depositor in a long asset that pays $R > 1$ at maturity but yields only $\ell \in (0, 1)$ per unit if liquidated early. The bank is solvent: its eventual assets exceed its liabilities. Let $w \in [0, 1]$ be the fraction of depositors who demand repayment early. Depositors served early are chosen at random, and the bank liquidates only as needed.

If $w \leq \ell$, the bank meets all withdrawals by liquidating w/ℓ units of the asset and holds the remainder to maturity, so a depositor who waits receives

$$V_{\text{wait}}(w) = \frac{R(1 - w/\ell)}{1 - w},$$

while a depositor who withdraws receives 1. If $w > \ell$, the bank liquidates everything and can pay in total only ℓ , so a withdrawer receives ℓ/w in expectation and a waiter receives 0.

Theorem 11.1 (Coexistence of calm and run). *Assume $0 < \ell < 1 < R$. Then:*

(i) V_{wait} is strictly decreasing in w on $[0, \ell]$, with $V_{\text{wait}}(0) = R > 1$ and $V_{\text{wait}}(\ell) = 0$, and it equals the withdrawal payoff 1 at exactly

$$w^* = \frac{\ell(R-1)}{R-\ell} \in (0, \ell);$$

(ii) both $w = 0$ and $w = 1$ are equilibria, the first a calm equilibrium in which everyone waits and the second a run equilibrium in which everyone withdraws;

(iii) w^* is increasing in ℓ , so a deeper liquidation discount (smaller ℓ) lowers the threshold above which waiting ceases to be the best response.

Proof. (i) Differentiating, $\frac{d}{dw} \frac{1-w/\ell}{1-w} = \frac{-(1-w)/\ell + (1-w/\ell)}{(1-w)^2} = \frac{1-1/\ell}{(1-w)^2} < 0$ since $\ell < 1$. The endpoint values follow by substitution. Setting $R(1-w/\ell) = 1-w$ gives $R-1 = w(R/\ell-1)$ and hence the stated w^* , which lies below ℓ because $R-1 < R-\ell$. (ii) At $w = 0$ a single (measure-zero) deviator who withdraws receives $1 < R = V_{\text{wait}}(0)$, so waiting is a best response. At $w = 1$ a withdrawer receives $\ell > 0$ while a waiter receives 0, so withdrawing is a best response. (iii) Differentiating w^* gives $\frac{\partial w^*}{\partial \ell} = \frac{(R-1)(R-\ell) + \ell(R-1)}{(R-\ell)^2} = \frac{R(R-1)}{(R-\ell)^2} > 0$. $\square$

For $w < w^*$ waiting yields more than withdrawing, and for $w > w^*$ withdrawing yields more, so under best-response dynamics w^* separates the basin of the calm equilibrium from the basin of the run. An interior equilibrium exists at w^* with depositors indifferent, but it is unstable. The model has the property noted in the earlier sections: the institution is the same in both equilibria. Nothing about its assets or liabilities differs between calm and run, and the difference is entirely in the beliefs of creditors about one another. Adding noisy private signals of fundamentals selects a unique equilibrium with a threshold in the fundamentals in the global-games literature [15, 10], so the multiplicity here should not be read as a prediction of arbitrary switching, but it illustrates why individually rational withdrawal can contribute to a collectively destructive liquidation. Bernanke's testimony and the report of the Financial Crisis Inquiry Commission both describe the 2008 episode in terms of runs on short-term funding markets, especially within the shadow banking system [1, 8], and Gorton and Metrick analyze the repo market as a run on securitized collateral [11].

11.2 Opacity cuts both ways

The comparison with the interpersonal case is instructive. In an interpersonal fraud, concealment protects a failing arrangement from scrutiny, and creditors' beliefs about solvency are too optimistic. In a banking panic, widespread suspicion can destroy an arrangement that might otherwise have remained viable, and creditors' beliefs about one another are too pessimistic or at least self-confirming. The common problem is incomplete knowledge of the whole financial structure, but the consequences of the incomplete knowledge are not the same.

A formal summary is that strategic complementarity, the property that one creditor's incentive to withdraw rises with the number of others who withdraw, produces multiple equilibria in the banking case, while in the fraud case the relevant externality is informational and the equilibrium is unique but wrong.

12 A General Theory of Financial Opacity

The preceding sections used a common vocabulary of hidden liabilities and local observation. This section states the vocabulary generally.

12.1 Networks and projections

Represent the financial system as a directed weighted graph $G = (V, E, W)$, where vertices V are economic actors, $W_{jk} \geq 0$ is the liability of actor j to actor k , and each vertex j has external assets $e_j \geq 0$. Each participant observes only a projection $\pi_i(G)$. A natural choice is the *incident projection*, consisting of the weights W_{ij} and W_{ji} on edges incident to i , the flows of payments on those edges, and the participant's own assets, together with any public aggregates. The fundamental question is whether properties of the complete system can be inferred from the observations available to its participants.

Theorem 12.1 (Hidden-liability indistinguishability). *Let G be a network, $b \in V$ a borrower, and $S \subseteq V \setminus \{b\}$ a set of observers. For every $\Delta > 0$ there is a network G' obtained by adding a vertex $z \notin S$ and an edge of weight $W_{bz} = \Delta$ such that $\pi_i(G') = \pi_i(G)$ for every $i \in S$, while the aggregate liability of b rises by exactly Δ . If b has net worth $\nu = e_b + \sum_k W_{kb} - \sum_k W_{bk} \geq 0$ in G and $\Delta > \nu$, then b is solvent in G and insolvent in G' .*

Proof. Edges incident to $i \in S$ are unchanged, since the new edge is incident to b and z only and $b, z \notin S$. If hidden payments from z to b are needed to reproduce payment flows, they are incident to b and z only. Hence incident projections of members of S are equal. The aggregate liability of b is the sum of its out-weights, which increases by Δ , and the net worth falls by Δ , which changes its sign under the stated condition. $\square$

The theorem is the graph form of Theorem 3.1. Whether two networks can produce identical observations for every creditor while differing substantially in their resilience to a funding shock is answered affirmatively under restricted observation. We now isolate three results that organize the essay.

Proposition 12.2 (Summary of the three propositions). *(1) Local repayment non-identifiability. Successful repayment to an individual creditor is insufficient to determine aggregate solvency when undisclosed liabilities and external borrowing remain unconstrained (see Theorem 3.1, Proposition 3.4, and Theorem 12.1).*

(2) Selective repayment and distorted reputation. *If a borrower's allocation of payments affects which creditors provide public testimony, the observed repayment rate can systematically exceed the repayment rate across all creditors (Proposition 4.3), subject to explicit assumptions about reporting and selection, namely $g_p > g_u$.*

(3) Network amplification. *Under a specified price-impact and forced-liquidation model, an initial loss generates secondary losses exceeding the original shock, and instability occurs when the relevant feedback operator has spectral radius at least one (Theorem 10.3), in the form $x_{t+1} = Mx_t + u_t$ with stability if and only if $\rho(M) < 1$ in the irreducible case.*

12.2 Clearing and the dilution of the visible creditor

The clearing framework of Eisenberg and Noe [6] makes the effect of a hidden edge on recoveries explicit. Let $\bar{p}_j = \sum_k W_{jk}$ be total obligations of j and $\Pi_{jk} = W_{jk}/\bar{p}_j$. A clearing payment vector satisfies

$$p_j = \min\left(\bar{p}_j, e_j + \sum_k \Pi_{kj} p_k\right),$$

so that each node pays in proportion to its claims whenever it cannot pay in full. Consider a borrower b owing 1 to a visible creditor c and, in the network G' , a further 1 to a hidden creditor z , with external assets $e_b \in [0, 2]$ and no other receipts. In G , creditor c recovers $\min(1, e_b)$. In G' , c recovers $e_b/2$ for $e_b < 2$. When $e_b = 1$, creditor c recovers 1 in G and $1/2$ in G' , and the incident projection of c is identical in the two networks. The hidden creditor dilutes the visible one, and the dilution is invisible to the diluted party.

This also clarifies the contrast with selective repayment. Under proportional sharing, no creditor is made whole when $e_b < 2$ and testimony from both creditors is poor. If the borrower pays c in full first, c recovers 1 and reports a successful experience, while z bears the shortfall alone. The aggregate recovery e_b is the same in the two regimes, consistent with Proposition 4.2.

12.3 Projection and closure

The general question is what properties of a system survive when the system is observed through a restricted interface. Solvency does not, and neither does the criticality of the leverage operator, and neither does the equilibrium that a population of creditors occupies. What survives are quantities that the interface contains by construction: the claim of one party on another, the payment made in one period, the price at which a transaction cleared. The remaining claims are conclusions that require additional structure, whether in the form of registries, caps, verification, or models of the hidden part, and the content of those additional assumptions is where the real risk resides. To say that a repayment record is informative is to say something about the strength of that additional structure, and Propositions 3.3 and 3.4 measure that strength.

13 Trust Without Omniscience

The conclusion should avoid the simplistic position that nobody should borrow money or that financial trust is impossible. Credit is indispensable precisely because people need to make commitments across time. A mortgage finances a house before the borrower has accumulated its full purchase price. A business loan can finance productive capacity that generates future income. Even borrowing between friends can bridge a temporary interruption in cash flow. The problem arises when trust is grounded in observations that cannot establish what people believe they establish.

A completed payment is evidence of a completed payment. A photograph is evidence that a photograph exists. A favorable credit history is evidence about a particular history of obligations. None independently establishes the entire future financial position of the person or institution. The results of this essay place numbers on that sentence. A repayment record certifies own resources of at most $(X - \bar{B})^+$ and nothing when hidden credit is open. It raises the odds of solvency by exactly $1/G(t)$, which is one when the insolvent can refinance indefinitely. Reported repayment overstates true repayment by an amount fixed by who speaks. A group of recipients can fail to verify an appeal that is worth verifying by a factor of Np . A senior claim on a diversified pool inherits the probability of the systemic state however large the pool. A leveraged system is stable only below a threshold that no single balance sheet reveals, and a solvent bank can be in a calm or a run equilibrium without any change in its assets.

The common structure is that *a global structure is projected onto local observations, and the properties that decide the outcome are not recoverable from the projections*. The cases differ in what follows from this. In personal fraud, concealment sustains the arrangement. In a debt spiral, arithmetic does the work without anyone concealing anything. In the micro-solicitation configuration, the cost of checking exceeds its private benefit while remaining far below its collective benefit. In structured finance, a modeling assumption conceals a correlated exposure. In a run, suspicion sustains itself. The differences among these cases are as important as their shared form.

The conclusion can be stated in the following way.

A debt can be paid without being extinguished from the system. A creditor can be satisfied without the borrower becoming solvent. A bank can report profits while accumulating fragility. A financial system can preserve the appearance of trust long after the conditions supporting that trust have deteriorated.

The central danger is not borrowing itself but the substitution of local evidence for knowledge of aggregate obligations. The organizing question was how much a successful repayment can tell us about a financial system whose other obligations remain invisible. The answer, under explicit mathematical assumptions, is remarkably little, and the amount is exactly what the assumptions about hidden credit, reporting, correlation, and funding allow.

References

- [1] B. S. Bernanke, *Causes of the Recent Financial and Economic Crisis*, testimony before the Financial Crisis Inquiry Commission, September 2, 2010.
- [2] J. Coval, J. Jurek, and E. Stafford, The economics of structured finance, *Journal of Economic Perspectives* 23(1) (2009), 3–25.
- [3] *Cunningham v. Brown*, 265 U.S. 1 (1924).
- [4] D. W. Diamond and P. H. Dybvig, Bank runs, deposit insurance, and liquidity, *Journal of Political Economy* 91(3) (1983), 401–419.
- [5] A. Downs, *An Economic Theory of Democracy*, Harper & Row, 1957.
- [6] L. Eisenberg and T. H. Noe, Systemic risk in financial systems, *Management Science* 47(2) (2001), 236–249.
- [7] Federal Deposit Insurance Corporation, *Crisis and Response: An FDIC History, 2008–2013*, 2017.
- [8] Financial Crisis Inquiry Commission, *The Financial Crisis Inquiry Report*, U.S. Government Printing Office, 2011.
- [9] Federal Trade Commission, New FTC data show people have lost billions to social media scams, news release and Data Spotlight, April 2026.
- [10] I. Goldstein and A. Pautner, Demand-deposit contracts and the probability of bank runs, *Journal of Finance* 60(3) (2005), 1293–1327.
- [11] G. Gorton and A. Metrick, Securitized banking and the run on repo, *Journal of Financial Economics* 104(3) (2012), 425–451.
- [12] W. Hoeffding, Probability inequalities for sums of bounded random variables, *Journal of the American Statistical Association* 58(301) (1963), 13–30.
- [13] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed., Cambridge University Press, 2013.
- [14] H. P. Minsky, *Stabilizing an Unstable Economy*, Yale University Press, 1986.
- [15] S. Morris and H. S. Shin, Unique equilibrium in a model of self-fulfilling currency attacks, *American Economic Review* 88(3) (1998), 587–597.
- [16] M. Olson, *The Logic of Collective Action*, Harvard University Press, 1965.
- [17] C. A. Sims, Implications of rational inattention, *Journal of Monetary Economics* 50(3) (2003), 665–690.