

Crossing the Line: A Dynamical Systems Theory of Bank Runs

by
Lucas Brandt

The initial version of this paper was submitted in partial
fulfillment of the requirements for Mathematics 309:
Differential Equations, Nonlinear Dynamics, and Chaos;
it has since been iterated upon further

Williams College
Williamstown, Massachusetts

Spring / Summer 2026

1. Introduction

The collapse of Silicon Valley Bank (SVB) in March 2023 unfolded in just forty-eight hours. For months beforehand, the bank had accumulated unrealized losses on its securities portfolio as interest rates rose, a slow and visible decline in its fundamentals, yet deposits largely stayed put. Then, on 8 March, a routine capital-raise announcement shifted depositor expectations, sending withdrawal requests through the roof, and the bank failed within two days.¹ The speed and the suddenness of that collapse sit uncomfortably inside the standard framework economists use to think about bank runs.

The standard theory of bank runs, developed by Diamond and Dybvig in their landmark 1983 paper, gives a powerful and elegant explanation for why runs can happen at all.² Their insight was that banks are fragile because they do something valuable but structurally precarious: borrowing short and lending long. Depositors want to reach their money on demand and borrowers want long-term financing, and the role of a bank is to bridge that gap. However, Diamond and Dybvig argue that the bridge is unstable, as if depositors believe other depositors are about to withdraw, it becomes rational to withdraw first, even when the bank is perfectly solvent. The result is a self-fulfilling panic: the run happens because people expect a run.

Diamond and Dybvig show formally that this produces two equilibria, which they refer to quite creatively as “good” and “bad.” In the “good” one, only those who truly need cash early withdraw, the bank survives, and everyone ends up better off than they would be without a bank. However, in the “bad” one, everyone withdraws regardless of need, the bank fails, and everyone is worse off. Which one occurs depends, in the original model, on something left outside the model entirely: a rumor, a random signal, a coordination device.

¹ Board of Governors of the Federal Reserve System (2023). *Review of the Federal Reserve’s Supervision and Regulation of Silicon Valley Bank*. Washington, DC, April, which sets out the 8 March announcement, the 9 March run, and the 10 March closure.

² Diamond, D. and Dybvig, P. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3):401–419.

This is a beautiful and important result, but it only tells us that two equilibria exist, and says nothing about how a banking system moves between them, how long it can hover near the boundary, what the path of decline looks like before collapse, or why some shocks trigger runs while seemingly similar shocks do not. In effect, the model treats a bank collapse like a light switch: off or on, instantaneous, with nothing of interest in between.

However, real bank runs do not behave like light switches, but like SVB, with months of visible stress, partial withdrawals, managers scrambling, the balance sheet eroding, and then a sudden cascade. So, the interesting question is not only whether a run can happen, but why a system can absorb so much pressure for so long before collapsing all at once. That is a question about dynamics, not about equilibria, and dynamics calls for a different set of tools.

This paper applies the mathematics of nonlinear dynamical systems, the study of how interacting quantities change over time, to the bank-run problem. The core shift in viewpoint is simple: instead of asking where the system is (which equilibrium), we ask how the system moves (what are its dynamics). We treat depositor beliefs and bank fragility as two quantities that evolve continuously, governed by a coupled pair of equations. The equilibria of the Diamond-Dybvig model reappear naturally as the fixed points of this system, but the dynamical view delivers much more: the nullclines that organize the motion, the local stability of each fixed point, the separatrix that separates safety from collapse, and, most importantly, how the whole picture reshapes itself as the key background condition, the level of interest rates, changes. Box 1 collects the handful of terms someone who has not taken a course in non-linear dynamics needs; those comfortable with phase-plane analysis can skip it.

Our main results are threefold. First, the system undergoes a saddle-node bifurcation at a critical interest rate r^* . Below r^* , there is a single, globally stable resting point, meaning the bank is fundamentally safe and returns to calm after any disturbance. Above r^* , the system becomes bistable, with two stable fixed points, one good and one bad, separated by an unstable in-between point. The dividing ridge through that unstable point is the separatrix, or tipping line, and it splits the map into two basins of attraction. Start on

one side and the system flows to calm; start on the other and it flows to collapse. The creation of this tipping line is the moment a banking system becomes crisis-prone, which is not the same as the moment it fails.

Second, as the interest rate rises further above r^* , the system can undergo a Hopf bifurcation at $r^{**} > r^*$, where the good resting point loses stability and a limit cycle appears. The limit cycle is a closed loop in the state space, meaning the system neither calms down nor falls over, but instead oscillates persistently between near-health and near-crisis. We read this as the chronic-stress dynamics of a bank that is fundamentally impaired but has not yet collapsed, with repeated waves of concern and partial stabilization, each bringing it closer to the edge.

Third, we reinterpret a bank collapse as the crossing of a tipping line rather than as a choice between equilibria. In our framework, a run is not set off by a random coordination failure, but when an outside shock like a public announcement, a rating downgrade, or a large and visible withdrawal pushes the system across the separatrix into the bad resting points basin of attraction. The dynamics then carry it to collapse regardless of what anyone does next. This carries an important policy lesson: intervention must come before the line is crossed, not after.

We apply the framework to Silicon Valley Bank and argue that the Federal Reserve's aggressive 2022 to 2023 hiking cycle pushed SVB past the saddle-node threshold r^* , making the system bistable. The following months of on-and-off stress, partly stabilized deposit outflows, repeated reassurances from management, and a sliding share price, correspond to oscillatory dynamics near the tipping line. The capital-raise announcement of March 8th, 2023, by revealing the size of unrealized losses and signaling that management was under pressure, was the discrete shock to depositor beliefs that pushed the system across the line. Once that boundary was crossed, the outcome was no longer in doubt. The failure on March 10th, then, was not a surprise; it was the basin of attraction doing its work.

The paper also develops a parameter interpretation that ties the model directly to observable bank characteristics. The steepness of one curve in our belief equation, which we call k , governs how fast depositor beliefs respond to signals about fragility. We show that k rises as a depositor base becomes more uniform and more tightly connected, and in turn information travels faster between depositors. A concentrated, densely networked, sophisticated depositor base, precisely the venture-capital and start-up clientele SVB served, has a very high effective k : once beliefs begin to shift, they shift almost instantly. This formalizes the observation that SVB's depositor base was unusually dangerous from a stability standpoint, and it suggests that depositor concentration is a first-order risk factor that current regulation underweights.

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 presents the model and motivates each term from economic first principles. Section 4 carries out the phase-plane analysis, locating the nullclines and fixed points and examining their local stability. Section 5 presents the threshold (bifurcation) analysis, states and proves the two main results, and constructs the bifurcation diagram. Section 6 applies the framework to SVB. Section 7 draws out policy implications. Section 8 concludes.

Box 1. A Reader's Guide to the Key Ideas, in Plain Language

The argument uses a small number of ideas from the study of systems that change over time. None of them require advanced mathematics to understand. The following plain-language definitions are enough to follow the whole paper.

State and state space. At any moment the bank is described by two numbers: how worried depositors are, and how weak the balance sheet is. Think of these two numbers as a point on a map. The whole map of possible situations is the "state space".

Trajectory. As time passes, the point moves across the map, tracing a path. That path is the trajectory. It shows how the situation evolves from wherever it starts.

Fixed point. A spot on the map where the point stops moving. The bank can sit there indefinitely unless something disturbs it. Our model has a calm fixed point and, under some conditions, a crisis fixed point.

Stable vs. unstable. A resting point is stable if small nudges fade away and the system slides back to it, like a ball in a bowl. It is unstable if small nudges grow, like a ball balanced on a hilltop.

Basin of attraction. The region of the map from which all paths eventually flow into a given resting point. Each stable resting point has its own basin.

Separatrix (tipping line). The ridge line that divides one basin from another. On one side of it the bank recovers; on the other side it is doomed to fall into crisis. Crossing this line is the point of no return.

Nullclines. Two guide curves we draw on the map. On one, worry is momentarily steady; on the other, fragility is momentarily steady. Where the two cross, nothing is moving, so the crossings are exactly the resting points.

Bifurcation. A qualitative change in the shape of the map that happens when a background condition (here, the interest rate) passes a critical value. New resting points can appear, or existing ones can change character.

Saddle-node bifurcation. The moment when a calm-only map suddenly grows a second resting point (the crisis state) and a tipping line between the two. This is when a bank first becomes capable of a run.

Hopf bifurcation. The moment when a resting point stops being a quiet destination and instead becomes encircled by a repeating loop. The system can no longer settle; it cycles.

Limit cycle. A closed loop that the system circles indefinitely: a permanent oscillation between near-calm and near-crisis. We read it as a bank stuck in chronic stress, never settling and never quite failing, until a shock tips it over.

2. Related Literature

This paper sits at the meeting point of three literatures: the theory of bank runs and financial fragility, the global-games approach to equilibrium selection, and the use of nonlinear dynamical systems to study economic phenomena.

The foundational paper in the bank-run literature is Diamond and Dybvig (1983). Their model, which this paper directly extends, establishes the two-equilibrium structure of banking fragility described above. The insight that demand-deposit contracts create liquidity but also fragility has been enormously influential, and it was recognized with the 2022 Nobel Prize in Economics.³ Later work in this tradition has enriched the framework in many directions: adding asymmetric information (Chari and Jagannathan, 1988)⁴, incorporating interbank markets (Allen and Gale, 2000)⁵, and modeling how runs interact with asset prices (Shleifer and Vishny, 1997)⁶. Our contribution is methodological. We bring continuous-time dynamics and threshold analysis to a framework that has remained essentially static.

The global-games literature, begun by Carlsson and van Damme (1993)⁷ and applied to bank runs by Morris and Shin (2001)⁸, tackles the equilibrium-selection problem that Diamond-Dybvig leaves open. By adding small amounts of private information about fundamentals, Morris and Shin show that the multiplicity of equilibria collapses to a single equilibrium, fixed by a threshold in fundamental quality. This is an important advance, but it does not speak to the dynamic question we care about: how a system evolves over time, and what governs the path to collapse. Our approach is complementary rather than

³ Royal Swedish Academy of Sciences (2022). *The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2022*, awarded to Ben S. Bernanke, Douglas W. Diamond and Philip H. Dybvig for research on banks and financial crises.

⁴ Chari, V. and Jagannathan, R. (1988). Banking panics, information, and rational expectations equilibrium. *Journal of Finance*, 43(3):749–761.

⁵ Allen, F. and Gale, D. (2000). Financial contagion. *Journal of Political Economy*, 108(1):1–33.

⁶ Shleifer, A. and Vishny, R. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.

⁷ Carlsson, H. and van Damme, E. (1993). Global games and equilibrium selection. *Econometrica*, 61(5):989–1018.

⁸ Morris, S. and Shin, H. S. (2001). Rethinking multiple equilibria in macroeconomic modeling. *NBER Macroeconomics Annual*, 15:139–161.

competing. We give up the sharp uniqueness result in exchange for a richer description of the run-up to crisis.

Applying dynamical systems to economics has a long history, from the predator-prey models adapted to business cycles (Goodwin, 1967)⁹ to the Minsky financial-instability hypothesis, in which stability itself breeds instability through endogenous leverage cycles.¹⁰ More recently, Brunnermeier and Sannikov (2014)¹¹ have built a sophisticated continuous-time model of financial instability using stochastic differential equations, and Geanakoplos (2010)¹² has analyzed leverage cycles through a sequence of static models. Our contribution is to bring the specific tools of nonlinear dynamics, phase-plane geometry, guide-curve analysis, and threshold theory, directly to the bank-run problem in a tractable two-dimensional setting that yields clean analytical results and intuitive pictures.

The SVB collapse has already generated a substantial empirical literature. Jiang, Matvos, Piskorski, and Seru (2023)¹³ provide a comprehensive balance-sheet analysis showing that SVB was far from alone in carrying large interest-rate-induced unrealized losses, and that many regional banks faced similar vulnerabilities. Their work motivates our choice of the interest rate as the key threshold parameter. Cookson, Fox, Gil-Bazo, Imbet, and Schiller (2023)¹⁴ document the role of social media in accelerating the run, which supports our high- k reading of SVB's depositor base. We draw on both papers in the case study.

⁹ Goodwin, R. (1967). A growth cycle. In *Socialism, Capitalism and Economic Growth*. Cambridge University Press.

¹⁰ Minsky, H. P. (1992). The financial instability hypothesis. Working Paper No. 74, Jerome Levy Economics Institute of Bard College.

¹¹ Brunnermeier, M. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.

¹² Geanakoplos, J. (2010). The leverage cycle. *NBER Macroeconomics Annual*, 24:1–65.

¹³ Jiang, E., Matvos, G., Piskorski, T., and Seru, A. (2023). Monetary tightening and U.S. bank fragility in 2023: Mark-to-market losses and uninsured depositor runs. NBER Working Paper 31048.

¹⁴ Cookson, J. A., Fox, C., Gil-Bazo, J., Imbet, J. F., and Schiller, C. (2023). Social media as a bank run catalyst. Working paper.

3. The Model

3.1 Motivation and Overview

Before presenting the formal model, it is worth saying in plain terms what we are trying to capture and why these particular tools fit.

Diamond and Dybvig's great insight was that a bank run is a coordination problem.¹⁵ Whether a run happens depends not on whether the bank is solvent, but on what each depositor believes the others will do. If you expect others to run, your best move is to run; if you expect others to stay, your best move is to stay. Both beliefs are self-fulfilling. The model therefore has two resting points, calm and panic, with no account of how the system travels from one to the other.

Our extension starts from a simple observation: in reality, beliefs and fundamentals are not independent, but rather, they evolve together. As a balance sheet deteriorates, depositors grow anxious. As they grow anxious, they withdraw, which forces the bank to sell assets, which weakens the balance sheet further. This is a feedback loop, and feedback loops, when they are nonlinear, produce exactly the behavior we see in banking crises: long stretches of calm punctuated by sudden collapse.

To capture this, we track two quantities. The first, $b(t)$, is the fraction of patient depositors who expect a run at time t . This is the belief variable from Diamond-Dybvig, now allowed to vary continuously through time. The second, $f(t)$, is bank fragility, loosely the gap between the bank's liquid assets and its potential withdrawal obligations, with higher fragility meaning the bank is closer to insolvency if it is forced to liquidate. Both quantities evolve continuously according to equations that encode the feedback between them.

3.2 State Variables and Setup

¹⁵ Diamond and Dybvig (1983).

Let $b(t)$ in the interval from 0 to 1 denote the fraction of patient depositors who, at time t , expect a run. When $b = 0$, all patient depositors are calm and expect others to stay. When $b = 1$, all of them expect a run and will withdraw at once. Values in between represent partial panic: some depositors are alarmed, others are not.

Let $f(t)$, a non-negative number, denote bank fragility, defined as the degree to which the bank's liabilities exceed the market value of its liquid assets. When $f = 0$, the bank is perfectly liquid and could meet every withdrawal without distress. As f rises, the bank must sell ever more illiquid assets at worse prices to meet rising withdrawals. Very high f describes a bank that is effectively insolvent at current interest rates and asset prices.

Two outside conditions matter most. The first is r , the level of interest rates set by the central bank. Higher rates lower the market value of the bank's long-dated assets and so raise fragility directly. The second is k , the steepness of belief updating, discussed at length below. All other symbols are structural constants that set the speed and intensity of the feedback and are collected in Table 1.

3.3 The Belief-Dynamics Equation

The first equation describes how depositor beliefs change over time:

$$\frac{db}{dt} = \alpha \sigma(f - f_c) b(1 - b) - \delta b$$

It has three parts, each with a clear economic meaning.

The factor $\sigma(f - f_c)$ is a smooth switch (a logistic, or S-shaped, function) of fragility:

$$\sigma(f - f_c) = \frac{1}{1 + \exp(-k(f - f_c))}$$

Here f_c is a critical fragility level and $k > 0$ sets how abruptly the switch flips. The function is near zero for low fragility and near one for high fragility, with a rapid transition around f_c . This captures a key feature of depositor psychology: beliefs do not track fragility smoothly and proportionally. Below a

threshold, depositors largely ignore balance-sheet signals; above it, they become acutely sensitive to them.

Why should belief updating have this on-off character? Because depositors are not judging the balance sheet in isolation; they are trying to forecast what other depositors will do. At low fragility the bank's problems are not yet large enough to be coordination-relevant: even a worried depositor has little reason to think others are worried too, so withdrawing is not compelling. Once fragility crosses f_c , the problems become common knowledge, and the calculus flips as now each depositor knows that the others know, and the incentive to run switches on.

The parameter k sets the speed of that flip and has a direct economic reading. A small k describes a diffuse depositor base: retail customers who check on their bank rarely, react slowly to news, and do not compare notes. A large k describes a tightly networked, sophisticated base: venture funds and start-ups that talk constantly, monitor bank health in real time, and can coordinate withdrawals fast. SVB had an extremely high effective k , its depositors were so densely connected that once any significant actor signaled concern, that concern propagated almost instantly across the whole base. This is a central reason the SVB run was so much faster than a typical run: structurally, k was large. Figure 1 shows how the switch sharpens as k grows.

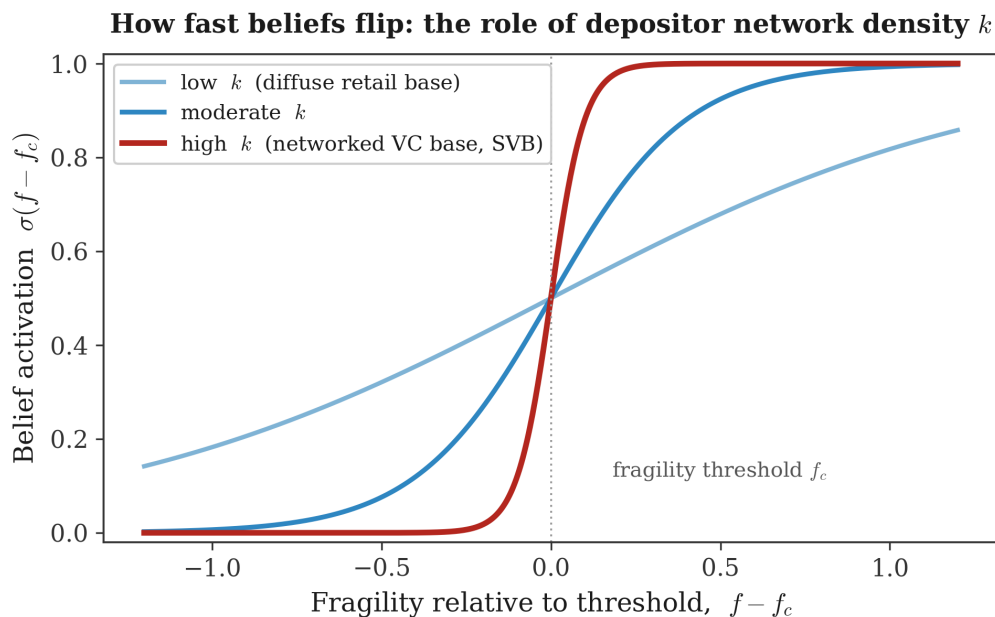


Figure 1. The belief-activation switch $\sigma(f - f_c)$ for three values of the steepness k . A low k (a diffuse retail base) gives a gentle, gradual response to rising fragility. A high k (a tightly networked base such as SVB's) makes the response a near step function: depositors flip from calm to alarm almost the instant fragility crosses the threshold f_c .

The factor $b(1 - b)$ captures how panic spreads and is borrowed from epidemic models, where it measures contact between an infected and a susceptible group.¹⁶ Here it reflects that panic spreads through interaction between alarmed and calm depositors. When b is near zero, very few are alarmed, so there are few carriers to transmit it, however when b is near one, almost everyone is alarmed, so there are few calm depositors left to convert. Transmission is fastest at intermediate b , exactly the region where a run is gathering momentum.

The final term, $-\delta b$, is mean reversion. Absent fresh coordination pressure, alarmed depositors gradually calm down as feared crises fail to arrive. The parameter $\delta > 0$ sets how fast panic decays when it is not reinforced, and it ensures that, at low fragility, the calm state is stable: the system returns to $b = 0$ after small disturbances.

¹⁶ The logistic contact term is standard in epidemic modeling; see Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London, Series A*, 115(772):700-721.

3.4 The fragility-dynamics equation

The second equation describes how fragility changes over time:

$$\frac{df}{dt} = \gamma r - \mu f + \beta b^n \quad (n > 1)$$

The term γr is the outside pressure from interest rates. When the central bank raises rates, the market value of the bank's fixed-income portfolio falls. This is mechanical accounting: a bond paying a fixed coupon is worth less once comparable new bonds offering higher yields. SVB had invested heavily in long-duration Treasuries and agency mortgage-backed securities bought at near-zero yields, so the Fed's 2022 to 2023 hiking cycle, one of the fastest on record, imposed enormous mark-to-market losses.¹⁷ The parameter $\gamma > 0$ measures how sensitive the balance sheet is to rate moves, which in turn reflects the portfolio's average duration, so longer-duration portfolios have higher γ .

The term $-\mu f$ is natural balance-sheet repair. Banks earn net interest income, retain earnings, and gradually rebuild capital, meaning that, left alone, a fragile bank slowly improves. The parameter $\mu > 0$ sets the speed of that repair. This stabilizing force grows with fragility itself: the weaker the bank, the harder it usually works to recover, by cutting dividends, lending less, or raising capital.

The term βb^n (with $n > 1$) is the heart of the feedback: the channel through which panic worsens the balance sheet. When depositors withdraw, the bank must pay out. With ample liquid reserves it does so at no cost; but as withdrawals scale up, it is forced to sell increasingly illiquid assets. The crucial point is that this damage is convex in the size of the run: a larger run does not merely cause proportionally larger damage, it causes disproportionately larger damage.

The mechanism is fire-sale dynamics. Selling a small quantity of bonds, a bank finds buyers near fair value and barely moves the price. Selling large quantities fast, the bank pushes prices down sharply. Other participants, seeing distress selling, widen their spreads and pull back, so the bank gets worse

¹⁷ On the scale of interest-rate-induced unrealized losses at SVB and across U.S. banks, see Jiang et al. (2023).

prices, books larger losses, and must sell still more to cover the remaining withdrawals. Therefore, each additional dollar withdrawn does more than a dollar of damage, otherwise known as convexity.¹⁸

We capture it with the power b^n , $n > 1$. For small b , withdrawals are modest and damage is limited; for large b , withdrawals are coordinated and damage is amplified. The exponent n sets how steeply fire-sale costs escalate with run intensity. For example, the value $n = 2$ gives quadratic damage, a natural baseline, whereas a higher n describes more illiquid portfolios, where forced selling moves prices more.

This convexity has a large qualitative consequence, developed below: it is precisely the ingredient that makes oscillatory behavior, the Hopf bifurcation, possible. A system with linear fragility accumulation cannot sustain oscillations; it simply converges or diverges. The nonlinear fire-sale amplification supplies the overshoot-and-correction that produces the wobble we associate with pre-crisis stress.

3.5 The Complete System and Parameter Summary

Collecting both equations, the complete dynamical system is:

$$\frac{db}{dt} = \alpha \sigma(f - f_c) b(1 - b) - \delta b$$

$$\frac{df}{dt} = \gamma r - \mu f + \beta b^n$$

where $\sigma(x) = \frac{1}{1 + \exp(-kx)}$

Table 1 summarizes every symbol, its economic meaning, its reading for SVB, and its qualitative effect on stability. From a policy standpoint the key parameters are r (the threshold parameter), k (depositor-network density), and n (fire-sale convexity). We return to all three in the policy section.

¹⁸ On fire-sale price dynamics and liquidity spirals, see Shleifer and Vishny (1997) and Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *The Review of Financial Studies*, 22(6):2201–2238.

Table 1. The model’s state variables and parameters, with their economic meaning, their interpretation in the SVB case, and their qualitative effect on stability.

Symbol	Economic meaning	Reading for SVB	Effect on stability
b	Fraction of depositors who expect a run (the belief state)	Largely uninsured, networked depositors	State variable
f	Bank fragility: how far liabilities exceed liquid asset value	Large unrealized losses on long-duration bonds	State variable
r	Interest-rate level set by the central bank	Fed funds rate, near zero to above 4% in a year	Bifurcation parameter; higher r raises fragility
k	Steepness of belief updating: how fast worry spreads	Very high: dense VC and start-up network	Higher k makes runs faster and sharper
n	Convexity of fire-sale damage from withdrawals	High: illiquid, long-duration holdings	Higher n enables oscillation
α	Contagion rate: strength of panic transmission	Elevated by tight information links	Higher α is destabilizing
δ	Mean reversion: rate at which panic fades on its own	Weak once concern becomes common knowledge	Higher δ is stabilizing
f_c	Fragility threshold at which depositors start reacting	Crossed as losses approached the equity cushion	Sets where the sigmoid switches on
γ	Sensitivity of the balance sheet to interest rates	High, due to long portfolio duration	Higher γ pushes the system toward crisis
μ	Speed of natural balance-sheet repair	Limited room to rebuild capital quickly	Higher μ is stabilizing
β	Intensity of withdrawal-driven damage (fire sales)	Large, given forced sales into a weak market	Higher β is destabilizing

4. Phase-Plane Analysis

To understand the geometry of the system, it helps to picture the two state variables as coordinates on a map (the phase plane), with b along one axis and f along the other. Every situation the bank can be in is a point on this map, and as time passes the point traces a path. Our task is to chart the map: where the points of no motion are, and which way the paths flow around them.¹⁹

4.1 Nullclines

We begin with the nullclines, the two guide curves along which one or other state variable is momentarily not changing. Where the two curves cross, neither variable is moving, so the crossings are the system's resting points.

The f -nullcline, where fragility is steady, is found by setting $df/dt = 0$:

$$\gamma r - \mu f + \beta b^n = 0 \implies f = \frac{\gamma r + \beta b^n}{\mu}$$

This curve rises and bends upward as b increases: for small b , the steady fragility level is about $\gamma r/\mu$, set almost entirely by interest rates; as b grows, the withdrawal term adds to it. The whole curve shifts upward as r rises, because higher rates require higher fragility for the balance sheet to hold steady. This upward shift is the geometric engine behind the saddle-node bifurcation proved in Section 5.

The b -nullcline, where beliefs are steady, is found by setting $db/dt = 0$:

$$\alpha \sigma (f - f_c) b (1 - b) - \delta b = 0 \implies b [\alpha \sigma (f - f_c) (1 - b) - \delta] = 0$$

¹⁹ The phase-plane apparatus used here (nullclines, the Jacobian, and linear stability analysis) is standard; see Strogatz, S. H. (1994). *Nonlinear Dynamics and Chaos*. Addison-Wesley; and Guckenheimer, J. and Holmes, P. (1983). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag.

This factors into two branches. The first is $b = 0$, the horizontal axis: if nobody expects a run, no contagion occurs and beliefs stay at zero. The second is an interior curve, found by solving the bracketed term:

$$\alpha \sigma(f - f_c)(1 - b) = \delta \implies b = 1 - \frac{\delta}{\alpha \sigma(f - f_c)}$$

This branch exists only when $\sigma(f - f_c)$ exceeds δ/α , that is, only when fragility is high enough to sustain a positive level of run expectation. For f near zero the switch is essentially off and no interior branch exists. As f rises past f_c the switch turns on and the branch appears, starting at high b and then decreasing as f climbs further. Geometrically it is a falling, S-shaped curve that emerges above a threshold level of fragility.

4.2 Resting Points and their Stability

The resting points are the crossings of the nullclines. Depending on the parameters there are either one or three of them, and that multiplicity is the heart of the model.

Resting point E₁: the good (calm) state. The $b = 0$ branch always crosses the f -nullcline at

$$(b^*, f^*) \approx \left(0, \frac{\gamma r}{\mu}\right)$$

This is the calm state: no run expectations, fragility set entirely by interest rates and the bank's structural position. To judge whether it is stable, we examine how small disturbances behave nearby. The standard tool is the Jacobian, a small table of slopes that records how each rate of change responds to a small nudge in each variable. Writing $g = db/dt$ and $h = df/dt$, it is

$$\mathbf{J} = \begin{bmatrix} \partial g/\partial b & \partial g/\partial f \\ \partial h/\partial b & \partial h/\partial f \end{bmatrix}$$

with entries

$$\partial g / \partial b = \alpha \sigma (f - f_c)(1 - 2b) - \delta$$

$$\partial g / \partial f = \alpha k \sigma (1 - \sigma) b(1 - b)$$

$$\partial h / \partial b = \beta n b^{n-1}$$

$$\partial h / \partial f = -\mu$$

The signs of these slopes (more precisely, the growth or decay rates they imply, the eigenvalues) tell us whether a small disturbance grows or fades. At E_1 , where b is essentially zero, the cross terms nearly vanish: the contagion term disappears because $b \approx 0$, and the withdrawal term disappears because $b^{n-1} \approx 0$ for $n > 1$. The two diagonal slopes are about $\alpha \sigma (\gamma r / \mu - f_c) - \delta$ and $-\mu$. At low r the switch is barely on, the first slope is negative, and both decay rates are negative, so disturbances fade and E_1 is stable. In short, the calm state is stable when interest rates are low.

Resting point E_2 : the saddle (tipping point). As r rises and the f -nullcline shifts up, it can begin to cross the interior b -branch. The first such crossing, when it exists, lands on the rising part of that branch. There the Jacobian implies one growing direction and one fading direction, the signature of a saddle: paths approach along one line and are flung away along another. The approaching line, the stable side of the saddle, is the separatrix, or the tipping line of the model.

Resting point E_3 : the bad (run) state. The second crossing, when it exists, sits at high b and high f . This is the run state: persistent panic and severe impairment. Here disturbances fade in both directions, so E_3 is stable; and depending on the parameters the approach may be a direct settling (a node) or a spiraling-in (a spiral). When it spirals, the system overshoots the run state before settling, an early hint of the oscillations that become central in Section 5. Whether E_3 is a node or a spiral depends on the strength of the convex fire-sale term.

4.3 The Separatrix as an Economic Tipping Point

The saddle E_2 deserves extended discussion, because it corresponds to a concept of great practical importance: the point beyond which a run becomes inevitable.

A saddle has two special directions. Along one, paths approach it; along the other, paths leave it. The approaching set, the separatrix, is a curve that divides the map into two regions, two basins of attraction. Paths starting on one side of it flow to the good state E_1 ; paths starting on the other side flow to the bad state E_3 . Figure 2 shows the full picture in the crisis-prone regime: the two nullclines, the three resting points, the separatrix, and a sample of trajectories colored by where they end up.

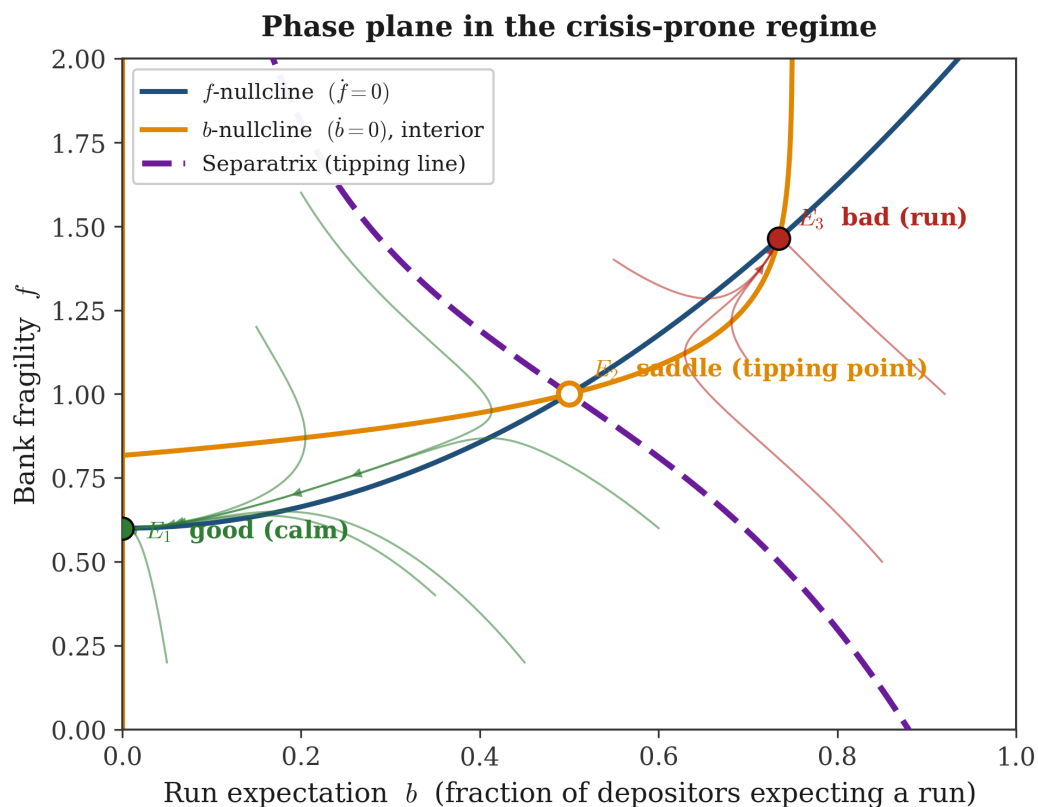


Figure 2. The phase plane in the crisis-prone (bistable) regime. The blue curve is the f -nullcline and the amber curve the interior b -nullcline; their crossings give the three resting points. E_1 (green) is the good, calm state and E_3 (red) the bad, run state, both stable. E_2 (open circle) is the unstable saddle. The dashed purple curve is the separatrix, or tipping line: trajectories starting to its left (green) return to calm, while those to its right (red) are carried to collapse. The same shock can therefore have opposite outcomes depending on which side of the line the bank sits.

The economic content of this picture is profound. A bank can lie in either basin, and which one it occupies depends not on its fragility or its belief level alone, but on the joint position of the pair (b, f)

relative to the separatrix. A highly fragile bank can still sit in the good basin if run expectations are low enough; a only moderately fragile bank can sit in the bad basin if expectations are already elevated. This explains why seemingly similar banks can fare very differently after the same shock: what matters is not the shock alone, but where it pushes the system relative to the line.

It also reframes the policy question. Standard regulation watches balance-sheet indicators, capital ratios, liquidity coverage, stress-test results, which are essentially measures of f . Our model suggests the relevant question is whether the joint state (b, f) is on the safe side of the separatrix, and that belief dynamics can carry the system across that line faster than any balance-sheet repair can pull it back. Attention to depositor beliefs, concentration, and network structure (the k parameter) may therefore matter as much as attention to traditional balance-sheet metrics.

5. Bifurcation Analysis

5.1 The Interest Rate as the Threshold Parameter

We now ask how the qualitative shape of the map changes as the interest rate r varies. Recall that raising r shifts the f -nullcline upward everywhere while leaving the b -nullclines unchanged. So as rates climb, the f -nullcline can gain new crossings with the interior b -branch, creating resting points, or lose existing ones, destroying them. These qualitative changes are bifurcations: moments when the structure of the map, not just its numbers, is reorganized.

5.2 The Saddle-Node Bifurcation

Our first result concerns the birth of the crisis-prone (bistable) regime.

Theorem 1 (Saddle-Node Bifurcation). There exists a critical interest rate r^* such that:

- (i) For $r < r^*$, the system has a single resting point E_1 , which is globally stable; the f -nullcline does not meet the interior b -branch.
- (ii) At $r = r^*$, the f -nullcline just touches the interior b -branch. Two resting points are born together: a saddle E_2 and a stable state E_3 .
- (iii) For $r > r^*$, there are three resting points, E_1 (stable), E_2 (saddle), and E_3 (stable). The system is bistable.

The proof follows from the geometry of the nullclines. The f -nullcline is an increasing, upward-bending function of b ; the interior b -branch is a decreasing function of f . For the two to meet, the f -nullcline must be high enough, which requires r to be large enough. At $r = r^*$ the curves are tangent, touching at exactly one point, which is the simultaneous birth of E_2 and E_3 . For $r > r^*$ they cross at two distinct points, creating the saddle-and-stable pair.

The economic interpretation is perhaps the single most important sentence in the paper: the saddle-node bifurcation at r^* is the moment a banking system changes from fundamentally safe to crisis-prone. Below r^* , no matter how bad the news or how alarmed some depositors become, the system always returns to calm, because E_1 is the only destination. Above r^* , the separatrix exists, and a large enough shock to belief or fragility can push the system into irreversible collapse. The Federal Reserve's hiking cycle did not merely move SVB to a worse spot within a fixed landscape; it changed the shape of the landscape itself.

5.3 The Hopf bifurcation

Our second result concerns the emergence of oscillation.

Theorem 2 (Hopf bifurcation). Under conditions on the parameters (specifically, when the convexity n and the contagion rate α are large enough relative to the damping parameters μ and δ), there exists $r^{**} > r^*$ such that, as r rises through r^{**} , the calm state E_1 loses stability and a stable limit cycle appears around it.

To see the mechanism, recall the condition for a Hopf bifurcation: the local growth and decay rates must turn from fading into a pure oscillation, which happens when the sum of the diagonal slopes of the Jacobian (its trace) passes through zero while the product-based quantity that guards against divergence (its determinant) stays positive. The trace at E_1 is

$$\text{tr}(\mathbf{J})|_{E_1} = [\alpha \sigma(f_1 - f_c) - \delta] + (-\mu)$$

where $f_l = \gamma r / \mu$ is the fragility level at E_1 . At low r the switch term is small, the bracket is negative, and the trace is negative, so E_1 is stable. As r rises, f_l rises, pushing the switch on. When r reaches r^{**} , the switch has turned on enough that

$$\alpha \sigma(f_1 - f_c) = \delta + \mu$$

and the trace passes through zero. At that moment the Hopf bifurcation occurs.

What does the limit cycle mean economically? For r between r^* and r^{**} , the calm state, though stable, is approached in a spiral: paths circle inward with shrinking swings. Once r exceeds r^{**} , the calm state itself becomes unstable: nearby paths spiral outward instead of inward and settle onto the limit cycle. The bank is no longer at rest but in permanent oscillation, repeatedly building up fragility and worry, partly stabilizing, then building up again.

This is, we argue, exactly the qualitative behavior of a bank in chronic pre-crisis stress. Consider a bank that has passed r^* but not yet suffered a decisive run. It is neither in the good state nor in the bad one; it cycles. A bad-news event triggers partial outflows; management responds with reassurances and some asset sales; calm partly returns; another shock arrives. Each loop brings the trajectory closer to the separatrix, and the limit cycle is the formal structure behind this drawn-out instability.

The Hopf bifurcation is made possible by the convexity of the fragility equation, the b^n term. With linear fragility accumulation ($n = 1$), the system generically only converges or diverges; it cannot generate the overshoot a cycle requires. It is precisely the fire-sale convexity, the feature that makes large coordinated runs qualitatively different from small ones, that creates the oscillatory instability. The same mechanism that makes runs catastrophic when they occur also makes the pre-crisis dynamics wobble and hard to read.

5.4 The bifurcation diagram

Figure 3 draws the bifurcation diagram, plotting the equilibrium level of fragility f^* against the threshold parameter r . It is the paper's central visual result, summarizing all three regimes in one picture.

Bifurcation diagram: three regimes of bank stability




Figure 3. Bifurcation diagram: equilibrium fragility f^* as a function of the interest rate r . For $r < r^*$ there is a single stable branch (the good state E_1 , green), so the bank is safe. At r^* a saddle-node bifurcation creates two new branches: the stable run state E_3 (red, upper) and the unstable saddle E_2 (amber, dashed). At r^{**} the calm branch loses stability through a Hopf bifurcation, and a band of oscillation (the limit-cycle amplitude, shaded) opens up. The three shaded regions are the monostable (safe), bistable (crisis-prone), and oscillatory (unstable) regimes.

For $r < r^*$, a single stable branch sits at low fragility, the good state, edging up with r as higher rates mechanically raise steady-state fragility. At $r = r^*$, the branch undergoes a saddle-node bifurcation: a new pair of branches appears, the upper one the stable run state E_3 and the middle one the unstable saddle E_2 , in the characteristic fold shape. At $r = r^{**}$, the lower branch loses stability through a Hopf bifurcation, and the diagram opens into a band whose width is the amplitude of the resulting limit cycle, growing as r increases.

The diagram compresses the three qualitative regimes of banking stability into one figure: monostable and safe for $r < r^*$, bistable and crisis-prone for $r^* < r < r^{**}$, and oscillatorily unstable for $r > r^{**}$. SVB’s path through 2022 and 2023 can be read directly as a journey rightward along the r -axis, passing through both

thresholds. Figure 4 shows the same three regimes as paths through time, the form in which they would actually be observed.




Figure 4. The three regimes seen as paths through time. (a) In the safe regime a disturbance fades and both run expectation b and fragility f return to calm. (b) In the chronic-stress regime the system settles onto a limit cycle, oscillating indefinitely between near-health and near-crisis without ever resting. (c) Once the tipping line is crossed, run expectation and fragility climb together and the system is carried to the run state: the rapid, irreversible cascade that, for SVB, played out in roughly forty-eight hours.

6. Application: Silicon Valley Bank

6.1 Background

At the time of its failure on March 10th, 2023, Silicon Valley Bank (SVB) was the sixteenth largest bank in the United States, with roughly \$209 billion in total assets.²⁰ It occupied a distinctive niche, serving as the primary banker for a large part of the venture-capital ecosystem and holding deposits from start-ups, technology companies, and the funds that backed them. That concentration was both the source of the bank's success and, as our model predicts, a structural weakness.

SVB's balance sheet reflected its unusual business, the bank held a disproportionately large securities portfolio, roughly \$120 billion in available-for-sale and held-to-maturity securities as of late 2022, nearly 60 percent of total assets. Much of it consisted of long-duration fixed-income instruments bought when interest rates were near zero. When the Federal Reserve began raising rates in March 2022, lifting the federal funds rate from near zero to above 4 percent within a year, the market value of those securities fell sharply. By late 2022 the bank had accumulated unrealized losses of roughly \$15 billion on its held-to-maturity portfolio alone, a figure that nearly matched its total equity capital.²¹

6.2 Mapping SVB to the model

We now read SVB's trajectory through the lens of the model. The path runs through all three regimes of the bifurcation diagram, moving rightward along the interest-rate axis until two thresholds have been crossed and a single shock finishes the job. Table 2 sets the timeline beside the model interpretation; the text that follows expands on it.

Phase 1 (pre-2022, $r < r^*$). Before the hiking cycle, SVB sat in the monostable regime. Its securities were marked near par, fragility was low, and depositor beliefs were calm. Small disturbances, occasional

²⁰ Board of Governors of the Federal Reserve System (2023); the firm grew to over \$211 billion in assets by 2021 and was the sixteenth largest U.S. bank at failure.

²¹ SVB balance-sheet figures are drawn from Jiang et al. (2023).

outflows from start-ups burning cash, minor market swings, were absorbed without consequence. The calm state E_1 was the only destination.

Phase 2 (2022, r crosses r^*). As the Fed raised rates through 2022, SVB's unrealized losses mounted. In the model's terms, the f -nullcline shifted upward as γr grew, and the system crossed the saddle-node threshold r^* . It became bistable, and a tipping line now existed in the state space. SVB had not yet run, and deposits were not yet collapsing; the system remained near E_1 , but the structural vulnerability had been created: the bank now lived in a landscape where a large enough shock could carry it across the line.

Phase 3 (late 2022 to early 2023, r approaches r^{}).** As rates kept rising, the system approached, and may have crossed, the Hopf threshold r^{**} . This is the period of intermittent stress: deposit outflows speeding up and then easing, the share price sliding in waves, management repeatedly reassuring the market. In the model these are oscillatory paths near the tipping line. The system is not converging to calm but circling near it, each loop bringing the state closer to the edge.

Phase 4 (8–9 March 2023, crossing the line). On March 8th, SVB announced that it had sold \$21 billion of its available-for-sale securities at a \$1.8 billion after-tax loss and was seeking to raise \$2.25 billion in new equity.²² This was the discrete belief shock, as it revealed at once that unrealized losses were being realized, that management was under pressure, and that the equity cushion was thinner than it had appeared. In the model, the announcement was a sudden upward jump in run expectation b , and that jump pushed the system across the separatrix.

Phase 5 (9–10 March 2023, flow to E_3). Once the line was crossed, the outcome was no longer in doubt. SVB's depositors, overwhelmingly large, uninsured, and sophisticated, transmitted alarm at near-instant speed. Withdrawal requests on 9 March reached roughly \$42 billion, about a quarter of the deposit base in

²² Details of the March 8th 2023 securities sale and capital raise are reported in Jiang et al. (2023).

a single day. The bank could not meet them, regulators were notified, and the FDIC placed it in receivership on 10 March.²³ The system had reached the run state E_3 .

Table 2. Silicon Valley Bank’s collapse mapped to the five dynamical phases of the model, from the safe regime through the crossing of the tipping line to irreversible failure.

Period	What happened	Model interpretation
Phase 1: pre-2022	Securities marked near par, fragility low, depositors calm. Routine outflows from cash-burning start-ups were absorbed easily.	Monostable (safe) regime, $r < r^*$. The calm state E_1 is the only destination.
Phase 2: through 2022	The Fed raised rates aggressively; unrealized losses on the bond portfolio mounted.	The interest rate crosses the saddle-node threshold r^* . A tipping line appears and the system becomes bistable.
Phase 3: late 2022 to early 2023	Deposit outflows accelerated then moderated; the share price fell in waves; management repeatedly reassured investors.	The rate approaches the Hopf threshold r^{**} . On-and-off stress reads as oscillation near the tipping line.
Phase 4: 8–9 March 2023	SVB sold \$21B of securities at a \$1.8B after-tax loss and sought \$2.25B in new equity, revealing the scale of losses.	A discrete upward jump in run expectation b . The shock pushes the state across the separatrix.
Phase 5: 9–10 March 2023	Withdrawal requests on 9 March reached roughly \$42B, about a quarter of deposits in one day. FDIC receivership followed on 10 March.	Irreversible flow into the run state E_3 . The basin of attraction carries the system to collapse.

6.3 The k Parameter and Depositor-Network Structure

The case study also brings out the role of the steepness parameter k , with direct regulatory implications. SVB’s depositor base was, by any measure, unusually concentrated and tightly networked. Venture funds and their portfolio companies made up a large majority of deposits, and these actors communicated constantly through shared investors, board members, and industry channels. When prominent venture capitalists began publicly urging their portfolio companies to pull funds on March 9th, by social media,

²³ Figures on the uninsured depositor withdrawals of 9 March 2023 are drawn from Jiang et al. (2023).

group chats, and direct calls, the message spread almost instantly across the whole relevant population.²⁴

There was no gradual erosion of confidence, there was a step change.

This is exactly what a high k implies in the model, the belief switch becomes extremely steep: depositors who were calm one moment flipped to full alarm the next as the fragility threshold was crossed. The practical consequence is that SVB's run was faster and more complete than a comparable bank with a diffused, retail-oriented base would have suffered. The network structure of the depositor base is therefore a first-order driver of crisis dynamics, not a second-order detail.

²⁴ On the role of social media in accelerating the run, see Cookson et al. (2023).

7. Policy Implications

The dynamical view yields several policy lessons that differ in emphasis from those of the static Diamond-Dybvig model.

1. The threshold r^* is the right regulatory target, not individual balance sheets. Standard regulation works to keep each bank's capital and liquidity buffers adequate. In the model's terms, that is regulation of fragility f , but the structurally decisive event is the moment the interest rate crosses r^* and the system becomes bistable, not the moment any one bank's fragility crosses some line. Macroprudential policy should therefore track where the banking system as a whole sits relative to the saddle-node threshold.

Doing so means estimating how sensitive bank balance sheets are to the rate cycle across the entire sector.

2. Deposit insurance shifts the threshold by changing belief dynamics. Insurance lowers the effective k and α : if depositors know their funds are guaranteed, they have less reason to monitor the bank or to run preemptively. Geometrically, this flattens the interior belief curve and makes the crisis intersection harder to reach. The Diamond-Dybvig conclusion that insurance eliminates runs holds in the static model²⁵; the dynamical view adds nuance, in that insurance raises the effective threshold r^* , buying the system more room before it turns bistable. SVB shows the limit of this protection: its base was overwhelmingly uninsured, which effectively switched the protection off and set k to its uninsured maximum.

3. Depositor concentration is an unregulated first-order risk. The parameter k , which governs how fast belief contagion spreads, is at present almost entirely unregulated. Banks face capital rules, liquidity rules, and stress tests, but no requirement regarding the diversity or network structure of their depositors. The model implies that a bank with a concentrated, tightly networked base is structurally more fragile, at any given level of balance-sheet health, than one with a diffuse retail base, even holding every traditional

²⁵ Diamond and Dybvig (1983), who show that deposit insurance can eliminate the run equilibrium.

risk metric constant. That argues for treating depositor-concentration measures as a complement to conventional prudential regulation.

4. Intervention must come before the line is crossed. Perhaps the most practical lesson is about timing. Once the system crosses the tipping line, the dynamics are irreversible short of extraordinary intervention. Emergency liquidity, deposit guarantees, and capital injections can move the line itself (by changing parameters) or move the system relative to it (by injecting liquidity or capital). But once paths are flowing strongly toward the run state, the force needed to pull the system back is very large. The intervention window is defined by the tipping line, not by failure indicators: regulators should be asking where the system stands relative to the line, rather than whether it has failed yet

8. Conclusion

This paper has argued that the canonical Diamond-Dybvig model of bank runs, foundational as it is, is limited by its static character. By treating a collapse as a choice between two fixed states, it cannot account for the rich run-up to crisis, the prolonged metastability, the oscillatory stress, and the sudden discontinuous collapse that real banking crises display.

We have proposed a dynamical extension in which depositor beliefs and bank fragility evolve together as a coupled nonlinear system. The saddle-node bifurcation at r^* marks the interest rate at which a banking system turns from structurally safe to crisis-prone, the moment the tipping line is created and the system becomes bistable. The Hopf bifurcation at r^{**} marks a further threshold at which the calm state loses stability and a self-sustaining oscillation appears, formalizing the idea of a bank in chronic pre-crisis stress. A collapse is reinterpreted as the crossing of a tipping line, which makes clear that the window for useful intervention closes when the line is crossed, before any observable failure.

The Silicon Valley Bank collapse offers a clean application. The Federal Reserve's 2022 to 2023 hiking cycle pushed the bank past both of these thresholds. The intervening months of on-and-off stress correspond to the oscillatory regime. The capital-raise announcement of March 8th was the discrete belief shock that crossed the line, and the dynamics then carried the institution to failure in forty-eight hours.

Several extensions suggest themselves. First, adding random noise to the system would allow analysis of noise-induced transitions, cases where the system crosses the tipping line not through a single discrete shock but through accumulated random fluctuations; this connects to the literature on stochastic resonance and could yield expressions for the expected time to collapse as a function of noise size and distance from the line.²⁶ Second, extending the framework to a network of banks, where each bank's fragility feeds into the others' belief dynamics, would allow analysis of systemic contagion, and the phase

²⁶ On noise-induced transitions and stochastic resonance, see Horsthemke, W. and Lefever, R. (1984). *Noise-Induced Transitions*. Springer-Verlag; and Gammaitoni, L., Hänggi, P., Jung, P., and Marchesoni, F. (1998). Stochastic resonance. *Reviews of Modern Physics*, 70(1):223–287.

transitions familiar from network science would likely appear naturally.²⁷ Third, the threshold parameter r could be made endogenous in a general-equilibrium model in which monetary policy responds to financial-stability indicators, creating a feedback loop between the central bank's reaction function and the stability landscape of the banking sector.

More broadly, the paper is an argument for bringing the toolkit of nonlinear dynamics more systematically into the study of financial fragility. The features that make financial crises so costly and so hard to prevent, discontinuous collapse, path dependence, threshold effects, oscillatory instability, are precisely the features that nonlinear dynamics is built to analyze.

²⁷ On cascades and phase transitions in networks, see Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771.

References

- Allen, F. and Gale, D. (2000). Financial contagion. *Journal of Political Economy*, 108(1):1–33.
- Board of Governors of the Federal Reserve System (2023). *Review of the Federal Reserve’s Supervision and Regulation of Silicon Valley Bank*. Washington, DC, April.
- Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *The Review of Financial Studies*, 22(6):2201–2238.
- Brunnermeier, M. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.
- Carlsson, H. and van Damme, E. (1993). Global games and equilibrium selection. *Econometrica*, 61(5):989–1018.
- Chari, V. and Jagannathan, R. (1988). Banking panics, information, and rational expectations equilibrium. *Journal of Finance*, 43(3):749–761.
- Cookson, J. A., Fox, C., Gil-Bazo, J., Imbet, J. F., and Schiller, C. (2023). Social media as a bank run catalyst. Working paper.
- Diamond, D. and Dybvig, P. (1983). Bank runs, deposit insurance, and liquidity. *Journal of Political Economy*, 91(3):401–419.
- Gammaitoni, L., Hänggi, P., Jung, P., and Marchesoni, F. (1998). Stochastic resonance. *Reviews of Modern Physics*, 70(1):223–287.
- Geanakoplos, J. (2010). The leverage cycle. *NBER Macroeconomics Annual*, 24:1–65.
- Goodwin, R. (1967). A growth cycle. In *Socialism, Capitalism and Economic Growth*. Cambridge University Press.
- Guckenheimer, J. and Holmes, P. (1983). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, New York.
- Horsthemke, W. and Lefever, R. (1984). *Noise-Induced Transitions: Theory and Applications in Physics, Chemistry, and Biology*. Springer-Verlag, Berlin.
- Jiang, E., Matvos, G., Piskorski, T., and Seru, A. (2023). Monetary tightening and U.S. bank fragility in 2023: Mark-to-market losses and uninsured depositor runs. NBER Working Paper 31048.

- Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London, Series A*, 115(772):700–721.
- Merton, R. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29(2):449–470.
- Minsky, H. P. (1992). The financial instability hypothesis. Working Paper No. 74, Jerome Levy Economics Institute of Bard College.
- Morris, S. and Shin, H. S. (2001). Rethinking multiple equilibria in macroeconomic modeling. *NBER Macroeconomics Annual*, 15:139–161.
- Royal Swedish Academy of Sciences (2022). *The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2022*. Press release, 10 October. Awarded to Ben S. Bernanke, Douglas W. Diamond and Philip H. Dybvig.
- Shleifer, A. and Vishny, R. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.
- Strogatz, S. H. (1994). *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Addison-Wesley, Reading, MA.
- Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771.