Finite-Size Effects, Heavy Tails and Transient Impact in Absorbing Margin-Liquidation Cascades

Eric Bonabeau[1, 2]

[1] *Discontinuity Capital, Berkeley, California, USA*
[2] *School of Complex Adaptive Systems, Arizona State University, Tempe, Arizona, USA**

Margin liquidation can turn a smooth price decline into an abrupt loss of market exposure. We study a one-asset absorbing threshold model in which positions are weighted by dollar exposure and forced sales feed back through price impact. The local instability is controlled by the exposure density near the maintenance threshold, the sensitivity of that threshold to price, and marginal impact. For finite cohorts, the classical Daniels-Smith fiber-bundle exponents remain the reference scaling laws. The new result is the weighted variance, including normalization fluctuations, and the resulting parameter-free prefactors for nonlinear price-mediated load sharing. Heavy-tailed exposures form a less settled regime in which rare large, conservatively levered holders can delay the loss of stability. An event-independent background-loss statistic agrees closely with the corresponding mean-field dilution prediction. Path-resolved simulations with partial deleveraging and transient impact show a finite cascade duration and stalled ignition when price impact decays too quickly.

I. INTRODUCTION

A leveraged market can fail abruptly even when the external shock itself changes smoothly. The mechanism is simple: a price decline pushes some positions through their maintenance thresholds, and forced sales then lower the same price that determines the solvency of the remaining positions. The question addressed here is when this feedback removes a macroscopic fraction of exposure at once. The answer is a local instability condition: the product of the dollar exposure density near the moving maintenance threshold, the sensitivity of that threshold to price, and the marginal price impact reaches unity. At that point the low-liquidation state disappears and the system jumps to a high-liquidation state. The jump is irreversible: reducing the external shock afterward does not restore positions that have already been sold.

Several results are derived in this article: (1) The instability is controlled by exposure, so the joint distribution of position size and leverage can move the collapse boundary substantially. (2) The abrupt transition does not require a square-root impact singularity, or even strict concavity: linear, logarithmic, saturating, and linear-to-square-root response laws can all produce the same loss of stability when their marginal response is large enough. (3) Finite cohorts have nontrivial and quantitatively predictable corrections. With finite exposure variance, the collapse threshold fluctuates as $N^{-\frac{1}{2}}$, its mean lies above the infinite-population threshold by order $N^{-\frac{2}{3}}$, and the state from which the macroscopic jump begins fluctuates as $N^{-\frac{1}{3}}$. Event-driven simulations recover the predicted prefactors without fitted parameters. Heavy-tailed position sizes break this ordinary finite-size picture: typical fluctuations shrink more slowly, while rare dominant holders can delay or reorganize the main event. In a path-resolved model, transient impact introduces a second competition, between the time needed for a cascade to recruit new sellers and the time over which its own price impact decays. When impact memory is short, the cascade is delayed and fragmented, and weak first waves can die out before a later macroscopic ignition.

These effects connect familiar financial feedback mechanisms to an absorbing threshold transition. Funding-liquidity spirals, leverage cycles, portfolio insurance, fire sales, and related models all contain versions of price-mediated deleveraging [1–9]. Price-mediated contagion with capital constraints and downward-sloping demand appears in early fire-sale models and in modern stress-testing frameworks [10–14]. Threshold models more broadly are a standard route to collective cascades [15–17], while smooth positive-feedback systems can develop folds and cusps [18–20]. The present setting differs in one important respect: liquidation is absorbing. The reversible fixed-point equation is useful for locating the loss of stability, but its lower fold is not a physical recovery path: the cusp does not represent equilibrium coexistence or reversible hysteresis.

The closest finite-size analog is the global-load-sharing fiber bundle [21, 22], for which failure occurs when a common load pushes heterogeneous elements beyond their thresholds. Fiber-bundle ideas have previously been used in financial cascade models [23–25]. Here the common variable is market price rather than a conserved mechanical load: failed positions carry unequal dollar exposures, exposure can correlate with leverage, price impact can be nonlinear, and the impact itself can decay. These changes generate the weighted finite-cohort theory, the heavy-tail effects, and the stalled-ignition regime developed below. Institutional data motivate nonlinear impact, including a crossover from linear to square-root response [26], while latent-liquidity models provide related mechanisms for nonlinear price response [27, 28]. Kyle' s linear model is used only as a benchmark [29].

A different structural mechanism is finite carrying capacity. Caravelli, Sindoni, Caccioli, and Ududec [30] study

* ebonabeau@asu.edu

stochastic growth when available resources impose a finite carrying capacity. In the present setting, the analogous extension would make the market' s remaining demand-side absorption capacity a state variable that is depleted by forced sales and replenished over time. That mechanism is distinct from the exogenous impact functions used here, because the capacity to absorb the next sale would depend on the history of previous sales. It is not included in the present model, but it is a natural extension for markets in which the pool of marginal buyers is itself finite.

Kavya [31] independently develops a threshold-density model of stop-loss and margin-liquidation cascades with a local density-impact instability criterion, shock-driven discontinuous transitions, and a Galton-Watson description of precursor cascades. Those results overlap with the equal-exposure instability analysis in Secs. III-IV and with the account-count branching construction in Sec. IX. Accordingly, the novelty claimed here does not rest on the existence of the fold, the local density-impact criterion, or the Galton-Watson mechanism. It rests on unequal dollar exposures and their correlation with leverage, the weighted finite-cohort calculation and its prefactors, the heavy-tail dominant-holder regime, and the competition between intrinsic cascade duration and finite impact memory. Empirical studies of crypto-perpetual liquidations likewise find event-heterogeneous early-warning behavior [32] and subcritical within-venue branching during an abrupt market-wide transition [33].

The paper proceeds from the minimal threshold model to progressively less idealized dynamics. Sections II–VI derive the instability boundary and show how it moves with the initiating drawdown, maintenance requirement, exposure-leverage dependence, and impact law. Sections VII–IX develop and test the finite-cohort, heavy-tail, and precursor-avalanche predictions using event-driven simulations and fiber-bundle methods [21, 34–39]. Section X then resolves partial deleveraging, execution delays, rate limits, and transient impact. The resulting picture is mechanical rather than calibrated to a specific asset: it identifies which combinations of balance-sheet thresholds, exposure concentration, and market impact are sufficient for an abrupt liquidation cascade, and which finite-size and time-scale effects can move or split that transition.

II. EXPOSURE-WEIGHTED THRESHOLD MODEL

A. Maintenance threshold

Let the initial price be 1. A long position with leverage $\updownarrow_i \geq 1$ has initial equity $1/\updownarrow_i$ per unit asset and debt $1 - 1/\updownarrow_i$. After a fractional drawdown d, the asset value is $1 - d$ while debt is unchanged. Requiring the equity ratio to remain above a maintenance fraction m gives a breach when

$$\frac{1/\updownarrow_i - d}{1-d} \leq m, \qquad \updownarrow_i \geq \updownarrow_{th}(d) \equiv \frac{1}{m + (1-m)d}. \quad (1)$$

The derivative

$$\frac{d\updownarrow_{th}}{dd} = -(1-m)\,\updownarrow_{th}^2 < 0 \quad (2)$$

enters the feedback gain. Leverage is parameterized by a latent standard normal variable Z,

$$\updownarrow_i = max\left[1,\ \bar{\updownarrow}\,(1 + \sigma Z_i)\right], \quad (3)$$

where $\bar{\updownarrow}$ is the central leverage and σ the relative cross-sectional dispersion. The lower censoring is irrelevant for the upper-tail calculations reported below.

B. Exposure weighting

A sale moves price in proportion to exposure, not to the number of accounts. Let Q_i be the initial position size and define the exposure-weighted breach-response function

$$F(d) = \frac{\mathbb{E}\left[Q\,\mathbf{1}\{\updownarrow \geq \updownarrow_{th}(d)\}\right]}{\mathbb{E}[Q]}. \quad (4)$$

The unweighted account fraction is recovered only when Q is constant or independent of the threshold variable. A tractable joint model is

$$Q = Q_0 \exp\left[s_q\left(\rho_{qz} Z + \sqrt{1-\rho_{qz}^2}\,U\right) - \frac{s_q^2}{2}\right], \qquad U \sim N(0,1), \quad (5)$$

with U independent of Z. The exponential tilt gives the closed form (Appendix A)

$$F(d) = \Phi\left(\chi + \frac{1 - \updownarrow_{th}(d)/\bar{\updownarrow}}{\sigma}\right), \qquad \chi \equiv \rho_{qz} s_q, \quad (6)$$

where Φ is the standard normal cumulative distribution. The joint distribution of position size and leverage enters through the exposure tilt χ. Positive χ means that larger positions are concentrated in the higher-leverage tail.

For an arbitrary exposure-weighted leverage distribution F_Q with density f_Q, Eq. (4) is $F(d) = 1 - F_Q\left(\updownarrow_{th}\right)$, and the local feedback gain contains

$$-f_Q\left(\updownarrow_{th}\right)\,\frac{d\updownarrow_{th}}{dd} > 0, \quad (7)$$

with the minus sign required because the maintenance threshold falls as the price falls.

C. Impact law and effective feedback scale

Let $f(x)$ be a normalized impact function with $f(0) = 0$, $f(1) = 1$, $f'(x) > 0$, and, unless stated otherwise, $f''(x) \leq 0$. The reversible static response associated with a liquidated exposure fraction B is

$$d = d_0 + \Lambda f(B), \tag{8}$$

up to a price floor that is kept inactive in the analytic results. Here d_0 is the exogenous drawdown that would occur without forced sales and Λ is an effective feedback amplitude. For the square-root metaorder law $I = Y v_H \sqrt{q/V_H}$, with v_H the volatility and V_H the traded volume on horizon H, selling a fraction B of a vulnerable exposure Q_v gives

$$\Lambda = Y v_H \sqrt{Q_v/V_H}, \qquad f(B) = \sqrt{B}. \tag{9}$$

This mapping is meaningful only when V_H and the participation regime are defined on the same horizon as the impact estimate. Λ is not an empirically measured quantity in the simulations below and Eq. (9) does not extrapolate to arbitrary Q_v/V_H. If both Q_v and V_H are fixed fractions of free float, free float cancels from their ratio. Float is therefore not an independent control variable of Eq. (8). Any float effect must operate through Q_v, V_H, volatility, or a change in the impact function itself.

III. ABSORBING DYNAMICS

The static equality map is

$$B = G(B) \equiv F\left[d_0 + \Lambda f(B)\right]. \tag{10}$$

Its local gain is

$$R(B) = G'(B) = \left[-f_Q\left(\updownarrow_{th}\right) \frac{d\updownarrow_{th}}{dd}\right] \Lambda f'(B), \tag{11}$$

which factorizes into the exposure density near the threshold, the balance-sheet sensitivity to price, and the marginal market impact. A fold of the equality map occurs at $R = 1$.

Liquidation itself is irreversible. Let L_t be the cumulative fraction of initial vulnerable exposure already liquidated. For complete liquidation and permanent impact the physical update is

$$L_{t+1} = max\left[L_t,\ G\left(L_t\right)\right]. \tag{12}$$

A stationary state needs only $L \geq G(L)$. It need not satisfy equality. Because G is monotone, iteration from $L_0 = 0$ converges to the least fixed point, analogous to monotone fixed-point constructions in payment networks, which admit least and greatest clearing vectors [40]. The equality map is therefore a bifurcation diagram, and only the loss of its low fixed point controls the onset of a monotone cascade. The descending branch and the lower fold are not a recovery path for exposure that has already been sold.

For nondecreasing driving and a monotone response map, iteration from zero is already nondecreasing, so the max operator in Eq. (12) is inactive along the approach to collapse. It becomes essential when the external field is later reduced, because it prevents already liquidated exposure from reappearing.

Figure 1 illustrates the distinction at $\bar{\updownarrow} = 2$, $\sigma = 0.16$, $m = 0.25$, $\chi = 0.30$, $\Lambda = 0.30$, and square-root impact. The equality map written as $d_0(B)$ has folds at

$$B_{up} = 0.0293, \quad d_{0,up} = 0.10893; \qquad B_{low} = 0.4128, \quad d_{0,low} = 0.08 \tag{13}$$

Under the absorbing update, increasing d_0 produces a single macroscopic jump at the upper-fold field $d_0 = 0.10893$, from $B = 0.0293$ to 0.7113, a jump of 0.682. Decreasing d_0 afterward does not recreate the sold exposure. The cusp of the equality map organizes the collapse boundary but should not be read as equilibrium coexistence or reversible hysteresis.

IV. ANALYTIC PHASE BOUNDARY

For the power impact family $f(B) = B^\alpha$, $0 < \alpha \leq 1$, define

$$z = \Phi^{-1}(B), \qquad A = 1 - \sigma(z - \chi), \qquad c = 1 - m, \qquad b_0 = m + c\, d_0. \tag{14}$$

In Eq. (14), the exposure shift is inside the parentheses before multiplication by the leverage dispersion.

Solving Eq. (6) for $\bar{\updownarrow}$ gives the equation of state

$$\bar{\updownarrow}(B; \Lambda) = \frac{1}{A\left(b_0 + c\Lambda B^\alpha\right)}. \tag{15}$$

The fold condition $\partial\bar{\updownarrow}/\partial B = 0$ yields (Appendix B)

$$\Lambda_{sp}(B) = \frac{\sigma b_0 B^{1-\alpha}}{c\left[\alpha A \varphi(z) - \sigma B\right]}, \tag{16}$$

where φ is the standard normal density and the denominator must be positive. The cusp is the minimum of Eq. (16) over the physical interval. Two dependencies follow. For fixed (σ, χ, α) the cusp fraction B^* is independent of m and d_0, while

$$\Lambda_c \propto \frac{b_0}{1 - m}, \qquad \bar{\updownarrow}^* \propto \frac{1}{b_0}. \tag{17}$$

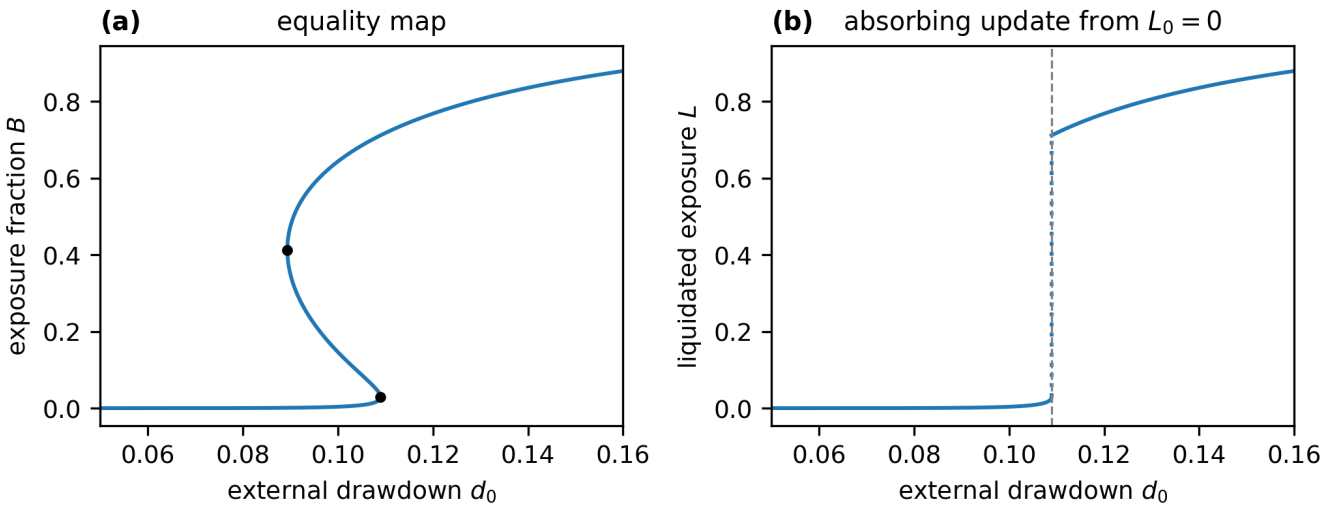


FIG. 1. Reversible bifurcation diagram versus absorbing dynamics for $\bar{\updownarrow} = 2$, $\sigma = 0.16$, $m = 0.25$, $\chi = 0.30$, $\Lambda = 0.30$, and square-root impact. (a) The equality map has two folds (dots). (b) The absorbing update started from zero follows the least fixed point until the upper fold is lost (dashed line), then jumps to the high-liquidation state. There is no lower-fold recovery of sold exposure.

A large initiating drawdown therefore moves the cusp to lower leverage. For $\sigma = 0.16$, $\chi = 0$, and $\alpha = 1/2$, $B^* = 0.14447$. At $m = 0.25$ the cusp $(\Lambda_c, \bar{\updownarrow}^*)$ is (0.212, 2.457) for $d_0 = 0.05$, (0.240, 2.173) for $d_0 = 0.10$, (0.295, 1.766) for $d_0 = 0.20$, and (0.389, 1.339) for $d_0 = 0.37$. Fragility at near-unit leverage requires a very large exogenous decline. For smaller shocks the cusp persists at higher leverage (Fig. 2).

At the cusp of the equality map, with $H\left(B; \bar{\updownarrow}, \Lambda\right) = G(B) - B$, the scalar conditions are

$$H = H_B = H_{BB} = 0, \qquad H_{BBB} \neq 0, \tag{18}$$

and a two-parameter unfolding additionally requires

$$\det \begin{pmatrix} H_{\bar{\updownarrow}} & H_\Lambda \\ H_{B\bar{\updownarrow}} & H_{B\Lambda} \end{pmatrix} \neq 0. \tag{19}$$

Direct substitution verifies Eq. (19) for the parameter sets used here.

V. POSITION-SIZE AND LEVERAGE DEPENDENCE

Equation (6) shows why account fractions can be misleading. The entire instability boundary shifts when large positions are concentrated in a different part of the leverage distribution. Figure 3 varies the exposure tilt at $d_0 = 0.10$, $m = 0.25$, $\sigma = 0.16$, and $\alpha = 1/2$. As χ rises from -0.8 to $+0.8$, Λ_c falls from 0.276 to 0.212 and the cusp leverage falls from 2.378 to 2.000. The equal-position case $\chi = 0$ lies between, with $(\Lambda_c, \bar{\updownarrow}^*) = (0.240, 2.173)$.

This result is not limited to the lognormal-normal construction. For a general joint distribution, the relevant quantity is the exposure-weighted density f_Q in Eq. (11). The scalar threshold map is unchanged in form, but the location and the existence of a fold depend on the amount of *dollar exposure* near the moving maintenance threshold.

VI. DEPENDENCE ON THE IMPACT LAW

A square-root law has a divergent derivative at zero, so it is necessary to determine whether the fold relies on that feature. For a general normalized f and fixed $\bar{\updownarrow}$, the equation of state in the external field is

$$d_0(B) = t(B) - \Lambda f(B), \qquad t(B) \equiv \frac{1/\left(\bar{\updownarrow} A\right) - m}{1 - m}, \tag{20}$$

where $t(B)$ is the exposure-weighted quantile of the individual threshold fields (Appendix C), and a fold satisfies

$$\Lambda_{sp}(B) = \frac{t'(B)}{f'(B)}, \qquad t'(B) = \frac{\sigma}{(1-m)\bar{\updownarrow} A^2 \varphi(z)}. \tag{21}$$

Any increasing impact law can create a fold if its marginal response is large enough relative to the weighted threshold density and the balance-sheet slope. Five functions with $f(1) = 1$ are considered:

$$\sqrt{x}, \qquad \frac{\sqrt{1 + x/x_c} - 1}{\sqrt{1 + 1/x_c} - 1}, \qquad x, \qquad \frac{\ln(1 + kx)}{\ln(1 + k)}, \qquad \frac{1 - e^{-kx}}{1 - e^{-k}}, \tag{22}$$

The five laws are square-root, linear-to-square-root crossover, linear, logarithmic, and saturating. With

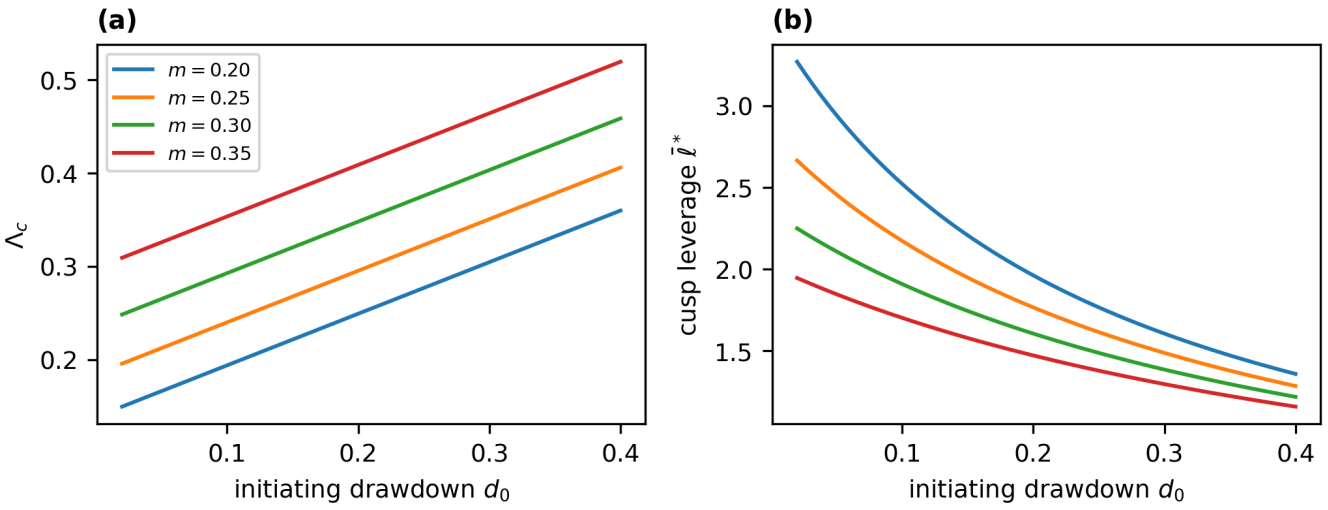


FIG. 2. Dependence of the cusp on the initiating drawdown and maintenance requirement for $\sigma = 0.16$, $\chi = 0$, and square-root impact. (a) Critical feedback amplitude Λ_c. (b) Central leverage $\bar{\ell}^*$ at the cusp. Smaller initiating drawdowns move the cusp to higher leverage without eliminating it.

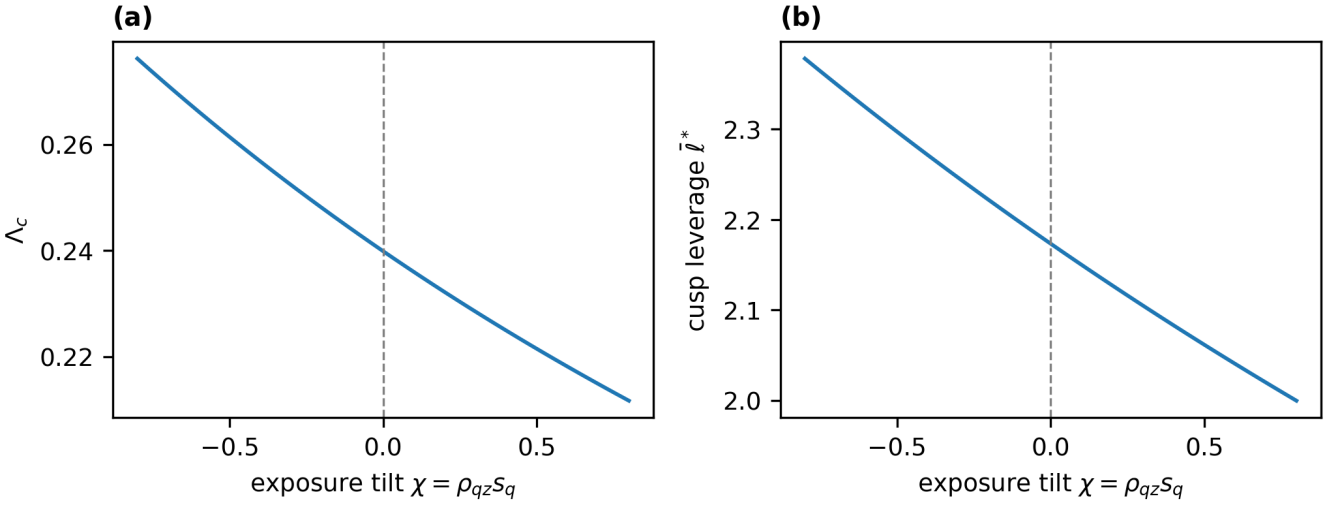


FIG. 3. Effect of exposure weighting at $d_0 = 0.10$, $m = 0.25$, $\sigma = 0.16$, and square-root impact. The tilt $\chi = \rho_{qz} s_q$ is positive when larger positions tend to have higher leverage. Positive tilt lowers both (a) the feedback amplitude needed for a cusp and (b) the central leverage at the cusp. The dashed line is the uncorrelated case.

$x_c = 0.02$ for the linear-to-square-root crossover, $k = 10$ for the logarithmic law, and $k = 4$ for the saturating law. For $\bar{\updownarrow} = 2$, $\sigma = 0.16$, $\chi = 0.30$, and $m = 0.25$, all five admit folds at finite feedback strength. The onset values of Λ, the minima of Eq. (21), are 0.241, 0.221, 0.233, 0.184, and 0.138, respectively (Fig. 4). These values are not empirically comparable because the functions differ in shape despite common endpoint normalization. The relevant result is that neither the square-root singularity nor strict concavity is necessary for a threshold-induced fold.

VII. FINITE-COHORT THEORY OF THE COLLAPSE FIELD

A. Collapse as the maximum of a weighted empirical process

The finite-size exponents used below are the classical Daniels-Smith results for global-load-sharing fiber bundles, not new critical exponents. The extension here is to unequal exposure weights and nonlinear price-mediated load sharing. The weighted empirical process contains fluctuations of both the breached-exposure numerator

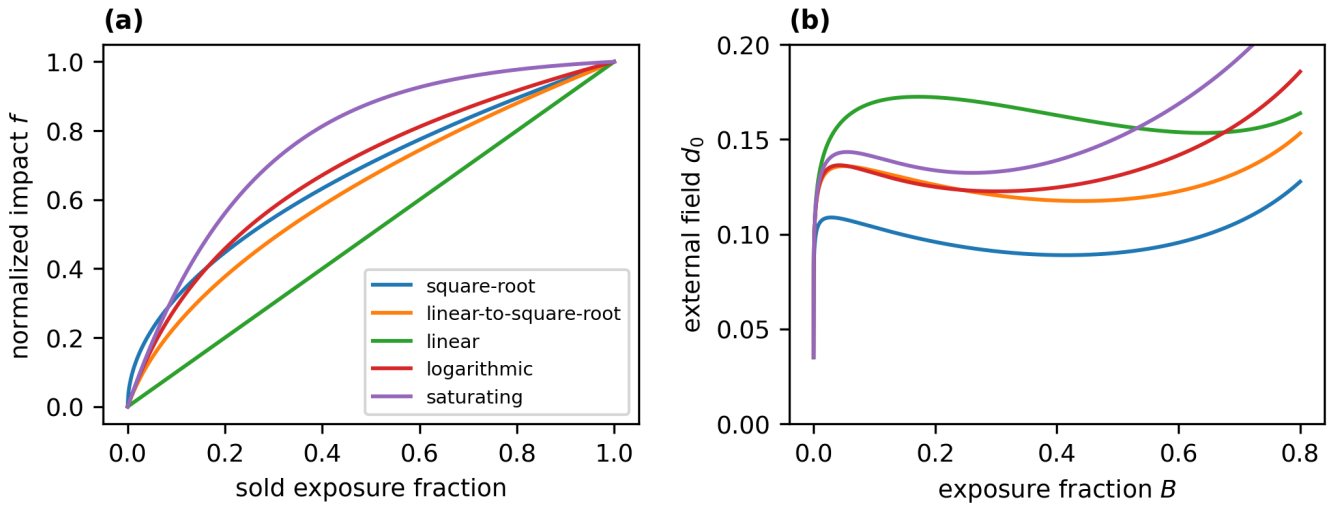


FIG. 4. Impact-law dependence. (a) The five normalized response functions of Eq. (22). The linear-to-square-root form reflects the crossover observed in institutional trading data [26]. (b) External-field equations of state at 1.25 times each law's own fold-onset feedback. All display a two-fold structure.

and its random normalization, which changes the variance entering the collapse-field calculation and fixes the prefactors tested below. A scalar-noise expansion of the fold condition is not sufficient because collapse is a first-passage event for a noise process that varies along the state axis. The appropriate asymptotic tool is the Daniels theory developed for fiber bundles [21, 34–37].

Consider N accounts with threshold fields t_i (Appendix C) and weights $w_i = Q_i/\sum_j Q_j$. Sort the accounts by increasing t_i and let $T_N(B)$ be the weighted empirical quantile, the threshold of the account at which the cumulative weight first reaches B. With complete liquidation and permanent impact, the external field required to breach that account when the exposure B before it has been sold is

$$D_N(B) = T_N(B) - \Lambda f(B), \qquad d_{0,+}^{(N)} = \max_{B \text{ near } B_+} D_N(B). \tag{23}$$

The low-liquidation state survives until d_0 exceeds the local maximum of D_N, so the second relation defines the finite-cohort collapse field. As $N \to \infty$, $D_N(B) \to d_0(B)$ of Eq. (20) and the maximum is the upper fold.

To leading order $T_N(B) - t(B) \simeq -t'(B)\,[F_N(t(B)) - B]$, where F_N is the weighted empirical distribution of the t_i. The process $\sqrt{N}\,(F_N - B)$ converges to a Gaussian process with variance

$$V(B) = (1-B)^2\, m_2(B) + B^2\,[m_2(1) - m_2(B)]\,, \qquad m_2(B) = \frac{\mathbb{E}\left[Q^2 \mathbf{1}\{t \le t(B)\}\right]}{\mathbb{E}[Q]^2}, \tag{24}$$

which includes the fluctuation of the normalization $\sum_j Q_j$, and with locally Brownian increments of variance rate $q_2(B) = dm_2/dB$. For the lognormal-normal model of Eq. (5),

$$m_2(B) = e^{s_q^2}\,\Phi(z+\chi), \qquad q_2(B) = exp\left(s_q^2 - \chi z - \chi^2/2\right). \tag{25}$$

Near the fold, with $u = B - B_+$ and $\kappa = -d_0''(B_+) > 0$,

$$D_N(B_+ + u) \simeq d_{0,+} - \frac{\kappa}{2}u^2 + \frac{t'(B_+)}{\sqrt{N}}\left[\sqrt{V(B_+)}\,Y + \sqrt{q_2(B_+)}\,W(u)\right] \tag{26}$$

where Y is a standard normal variable and W is a two-sided standard Brownian motion with $W(0) = 0$. The first noise term shifts the whole parabola and gives the leading dispersion. The second term makes the location of the maximum random. Brownian scaling [37] gives $\max_u\left[\epsilon W(u) - \kappa u^2/2\right] = \left(\epsilon^4/\kappa\right)^{1/3} M$ in distribution, with $M = \max_v\left[W(v) - v^2/2\right]$, and the maximizer scales as $(\epsilon/\kappa)^{2/3}$. With $\epsilon_0 = t'(B_+)\sqrt{q_2(B_+)}$, the predictions are

$$s.d.\left[d_{0,+}^{(N)}\right] \simeq t'(B_+)\sqrt{V(B_+)}\;N^{-1/2}, \tag{27}$$

$$\left\langle d_{0,+}^{(N)}\right\rangle - d_{0,+} \simeq \mathbb{E}[M]\left(\frac{\epsilon_0^4}{\kappa}\right)^{1/3} N^{-2/3}, \qquad s.d.\left[B_{takeoff}^{(N)}\right] \simeq s_a\left(\frac{\epsilon_0}{\kappa}\right) \tag{28}$$

with $\mathbb{E}[M] = 0.99619$ [37] and $s_a = 2^{2/3}\sqrt{0.26356} = 0.812$ the standard deviation of the maximizer, obtained from the variance of Chernoff's distribution [41]. The mean shift is positive: a finite cohort collapses later than the mean-field fold. The displacement of the state at takeoff is of order $N^{-1/3}$, not the $N^{-1/4}$ that a scalar-noise

saddle-node expansion would suggest. Equal exposures are recovered with $m_2(B) = B$, for which $V = B(1-B)$ and $q_2 = 1$, the classical bundle result.

B. Numerical test

Eqs. (27) and (28) were tested with the event-driven quasistatic protocol of Appendix C, which has no control grid. Cohorts have $\overline{\updownarrow} = 2$, $\sigma = 0.16$, $m = 0.25$, $\Lambda = 0.30$, $s_q = 0.75$, and $\rho_{qz} = 0.4$, so that $\chi = 0.30$ and the mean-field fold is Eq. (13). The theory inputs are $t'(B_+) = 0.8761$, $\kappa = 6.765$, $V(B_+) = 0.0937$, and $q_2(B_+) = 2.959$, which give the coefficients 0.2682 for the dispersion, 0.9100 for the mean shift, and 0.2995 for the takeoff state. Table I and Fig. 5 compare these predictions with simulations for $N = 250$ to $256,000$.

For $N \geq 4000$ the mean shift agrees with Eq. (28) within statistical error. The simulated dispersion of the collapse field lies 4% to 7% below the leading-order prediction, and the ratio approaches one slowly, as expected for corrections of relative order $N^{-1/6}$. The takeoff-state dispersion agrees within 6% for $N \geq 16,000$. Fitted over the full range, the exponents are -0.48 for the collapse-field dispersion and -0.67 for the mean shift. The mean macroscopic jump converges to the mean-field value 0.682 from above (0.684 at $N = 256,000$). The dispersion of the jump size decays with an effective exponent of -0.39 over the full range and -0.41 over the four largest sizes. This is a crossover: the jump is the difference between a landing state, whose fluctuations are of order $N^{-1/2}$, and the takeoff state, whose fluctuations are of order $N^{-1/3}$ and dominate asymptotically.

VIII. HEAVY-TAILED EXPOSURES

The theory developed in the previous section requires $\mathbb{E}\left[Q^2\right] < \infty$. Economic size distributions and trading-related quantities are often heavy tailed [42], and granular effects of this kind are known to survive aggregation in other economic settings [43]. Assume Q is Pareto distributed,

$$\Pr(Q > x) = x^{-\mu}, \qquad x \geq 1, \quad \mu > 1, \tag{29}$$

independent of leverage, so that $\chi = 0$ and the mean-field map is unchanged by the weighting. For $\overline{\updownarrow} = 2$, $\sigma = 0.16$, $m = 0.25$, $\Lambda = 0.30$, and square-root impact the folds are $(B_+, d_{0,+}) = (0.0415,\ 0.12746)$ and $(B_-, d_{0,-}) = (0.3435,\ 0.11713)$. For $1 < \mu < 2$ the sum $\sum_i Q_i (\mathbf{1}_i - B)$ that controls $F_N - B$ is attracted to a stable law and fluctuates on the scale $N^{1/\mu}$, so that the typical spread of the collapse field is

$$IQR\left[d_{0,+}^{(N)}\right] \sim N^{-(1-1/\mu)}, \qquad 1 < \mu < 2, \tag{30}$$

in place of $N^{-1/2}$. There is a second scale. The largest weight in a cohort is of order $N^{1/\mu - 1}$, and the probability that a single account holds more than a fixed fraction ω of the total exposure is of order $N^{1-\mu}$. Such an account changes the collapse field by an amount of order one. If its threshold lies above the fold field, which is the likely case because only a fraction B_+ of accounts lie below it, the account dilutes the weights of all others: the remaining cohort behaves as if the feedback were reduced to $\Lambda f((1-\omega)B)$, the fold moves to a higher field or disappears, and the collapse is late. A dominant account below the fold triggers an early collapse instead, with a probability smaller by a factor of order B_+. Hence

$$\Pr\left[d_{0,+}^{(N)} > d_{0,+} + \delta\right] \sim N^{1-\mu}, \qquad s.d.\left[d_{0,+}^{(N)}\right] \sim N^{-(\mu-1)/2}, \tag{31}$$

for fixed δ of order one in units of the threshold dispersion. Since $(\mu-1)/2 < 1 - 1/\mu$ for $1 < \mu < 2$, the variance is controlled by the rare late collapses, not by the core.

Simulations were run with $\mu = 1.25$, 1.5, 1.75, 2.5, and 4 for $N = 10^3$ to 1.024×10^6, with 20,000 cohorts at the smallest size and 250 to 400 at the largest (Fig. 6). The fitted exponents of the interquartile range are 0.23, 0.31, 0.37, 0.47, and 0.50, against the predictions 0.20, 0.33, 0.43, 0.50, and 0.50. The deviations for $\mu = 1.75$ and 2.5 are in the direction of the slow crossover expected near $\mu = 2$. For $\mu = 1.25$ the probability that the collapse field exceeds $d_{0,+} + 0.04$ falls from 0.088 at $N = 10^3$ to 0.015 at $N \approx 10^6$, an exponent of 0.26 against the predicted $\mu - 1 = 0.25$. Early collapses by more than 0.02 occur in fewer than 0.1% of cohorts at every size, which confirms the asymmetry. The fitted exponents of the standard deviation are 0.12, 0.30, and 0.45 for $\mu = 1.25$, 1.5, and 1.75, against $(\mu-1)/2 = 0.125$, 0.25, and 0.375. These estimates depend on a small number of rare cohorts and are less reliable than the interquartile range. At $N \approx 10^6$ and $\mu = 1.25$ the standard deviation of the collapse field is 0.017, seven times the interquartile range, and the account that initiates the macroscopic avalanche has a median weight 2.7 times the mean weight, against a population median of 0.35.

These fits support the predicted ordering and broad scaling trend, but they are not a precise asymptotic confirmation. The clearest discrepancy is for tail index 1.75, where the fitted interquartile-range exponent is 0.37 compared with the predicted 0.43. The standard-deviation exponents are especially sensitive to a few rare cohorts and should be interpreted qualitatively.

The first relation in Eq. (31) follows from a one-big-jump estimate. For a fixed exposure share, a cohort has as many chances as accounts for one Pareto draw to be comparable with the total exposure, which gives the exponent shown in Eq. (31). The standard-deviation relation needs one additional assumption: a dominant-holder event must shift the collapse field by an amount that does not vanish with cohort size. The second relation

TABLE I. Finite-cohort collapse statistics for lognormal exposures, with simulation and parameter-free theory predictions.

		Mean shift (sim.)	Mean shift (theory)	S.d. field (sim.)	S.d. field (theory)	S.d. takeoff (sim.)	S.d. takeoff (theory)
250	4000	0.02428	0.02293	0.01405	0.01696	0.0327	0.0475
1000	4000	0.01022	0.00910	0.00793	0.00848	0.0248	0.0299
4000	4000	0.00365	0.00361	0.00395	0.00424	0.0173	0.0189
16 000	3000	0.00142	0.00143	0.00203	0.00212	0.0112	0.0119
64 000	2000	0.00056	0.00057	0.00101	0.00106	0.0073	0.0075
256 000	1000	0.00026	0.00023	0.00051	0.00053	0.0046	0.0047

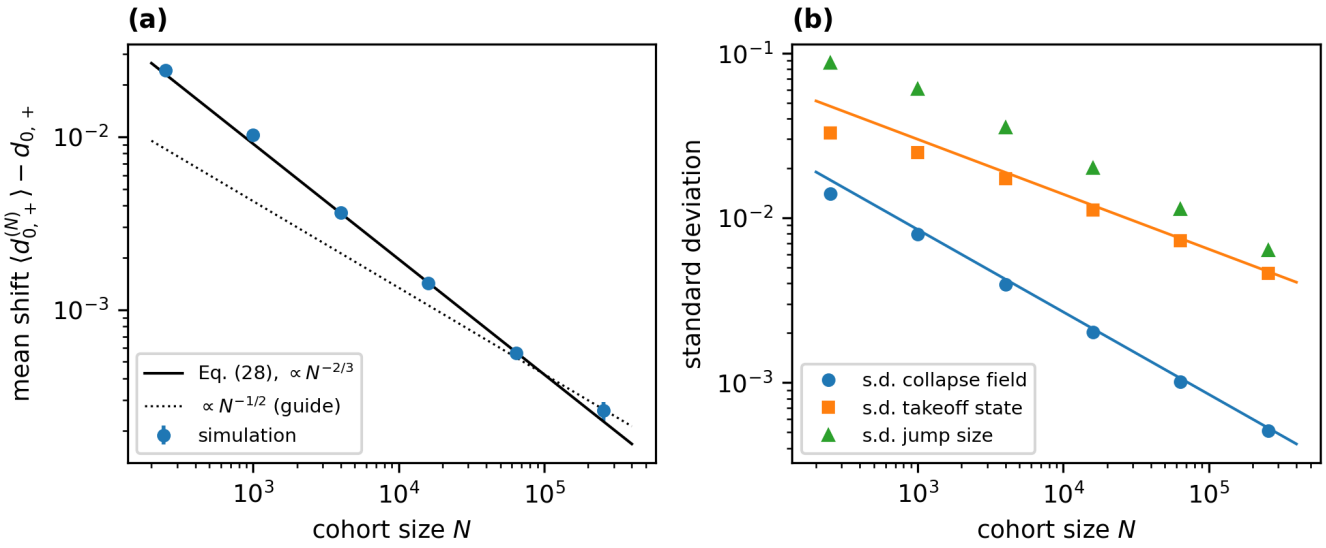


FIG. 5. Finite-cohort statistics at the collapse boundary for lognormal exposures. (a) Mean shift of the collapse field above the mean-field fold. Error bars are two standard errors and the solid line is Eq. (28) with no fitted parameter. (b) Standard deviations of the collapse field (circles), the takeoff state (squares), and the jump size (triangles). Solid lines are Eqs. (27) and (28).

in Eq. (31) is therefore treated as a conditional asymptotic prediction.

The dilution mechanism was therefore tested with an event-definition-independent statistic. While a conservative dominant holder remains intact, the mean-field background map is obtained by applying impact to the fraction of total exposure sold by the remaining cohort. For holder shares of 0%, 10%, and 20%, the corresponding upper-fold fields are 0.12746, 0.13083, and 0.13509. In the same controlled cohort of 100,000 background accounts, the fields at which the background first lost 25% of its own exposure were 0.12784, 0.13136, and 0.13561. The simulation therefore tracks the mean-field dilution prediction to within 0.00053 in all three cases. This check does not depend on identifying the largest avalanche and shows directly that a sufficiently large conservative holder delays the loss of stability of the remaining cohort.

This result isolates the robust part of the dominant-holder effect. A conservative holder reduces the fraction of total exposure that participates in the background feedback, so a larger external shock is required before the remaining cohort loses stability. The much larger shift produced by the largest-avalanche statistic in an earlier diagnostic is not used as evidence here because, for an extreme holder, the event definition itself can change which avalanche is labeled as the collapse. The heavy-tail scaling of the standard deviation remains conditional on rare dominant-holder events producing an order-one shift, and the simulations should be read as support for the mechanism and scaling order rather than a precise asymptotic confirmation.

IX. PRECURSOR AVALANCHES

In the quasistatic protocol, each increase of the external field releases an avalanche of S accounts. For equal exposures, the avalanche started at state B is a Galton-Watson process with Poisson offspring of mean $R(B)$, Eq. (11), so its size has the Borel distribution. The number of avalanches started in dB is proportional to $[1 - R(B)]\, dB$, and the size distribution integrated over the ramp up to the collapse is

$$D(S) = \frac{1}{\mathcal{N}} \int_0^{B_+} [1 - R(B)] \, \frac{e^{-R(B)S} \, [R(B)S]^{S-1}}{S!} \, dB, \qquad \mathcal{N} = \int_0^{B_+} \tag{32}$$

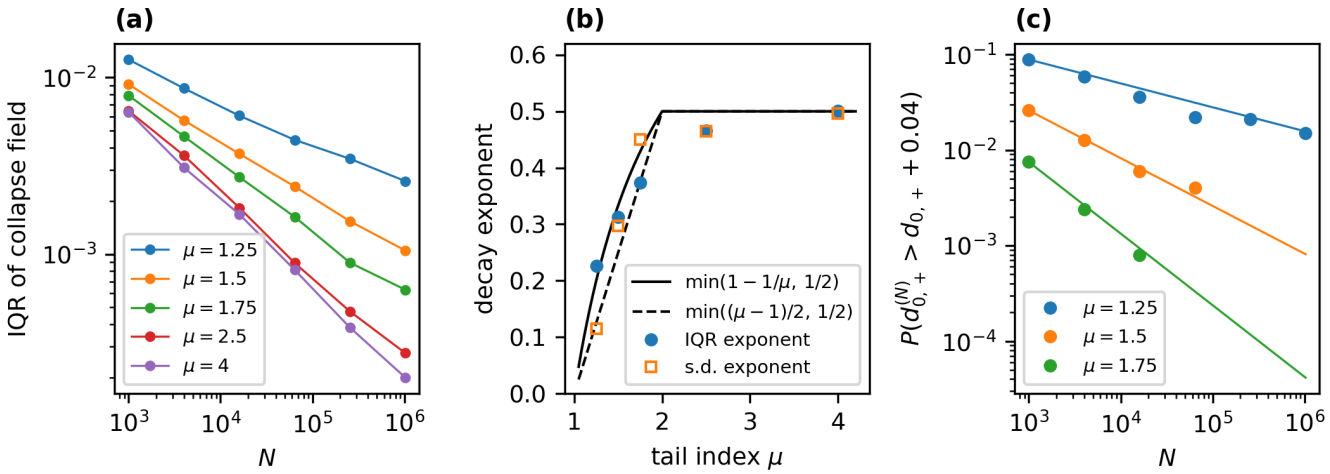


FIG. 6. Heavy-tailed exposures independent of leverage. (a) Interquartile range of the collapse field versus cohort size for five Pareto tail indices. (b) Fitted decay exponents of the interquartile range (filled circles) and of the standard deviation (open squares) compared with Eqs. (30) and (31). (c) Probability of a late collapse, $d_{0,+}^{(N)} > d_{0,+} + 0.04$. Lines have the predicted slope $1 - \mu$ and pass through the first point. Points with fewer than three events are omitted.

Because $1 - R$ vanishes linearly at the fold, $D(S) \sim S^{-5/2}$ asymptotically, as in fiber bundles [38, 39]. Figure 7 compares Eq. (32) with simulations of 40 cohorts of $N = 2 \times 10^6$ accounts at the parameters of Sec. VIII for equal and Pareto exposures and of Sec. VII for lognormal exposures. For equal exposures the prediction, which has no fitted parameter, agrees with the simulation within 2% for $S < 200$. At larger sizes the simulation exceeds the prediction by 20% to 40%. A separate size sweep tests whether this excess decreases with cohort size. The approach to the asymptotic exponent is slow: the local logarithmic slope of Eq. (32) is -2.0 to -2.1 for S between 10 and 25 and is still -2.3 at $S \approx 300$, because for square-root impact the gain rises as $\sqrt{B}$ at small B and a large share of avalanches start far from the fold.

For avalanche sizes between 200 and 1000 accounts, the geometric mean simulation-to-theory ratio fell from 1.25 at 125,000 accounts, to 1.20 at 500,000 accounts, and to 1.08 at 2,000,000 accounts. The mean absolute relative discrepancy over the same range fell from 25.3% to 20.0% and then to 11.7%. This size dependence supports finite-cohort rounding as the main source of the excess, although the largest avalanches remain too sparse for a precise asymptotic test.

For lognormal exposures ($s_q = 0.75$, $\rho_{qz} = 0.4$) and for Pareto exposures with $\mu = 1.5$ the integrated distributions are close to the equal-exposure curve over the accessible range. For heavy-tailed offspring a critical branching process has a size exponent $1 + 1/\mu$ in place of $3/2$, which would give an integrated exponent $3 - 1/\mu$. The difference between 2.33 and 2.5 cannot be resolved given the slow crossover. The practical implication is that a measured precursor exponent near 2 is compatible with the mean-field theory at realistic sizes and should not be read as evidence against it.

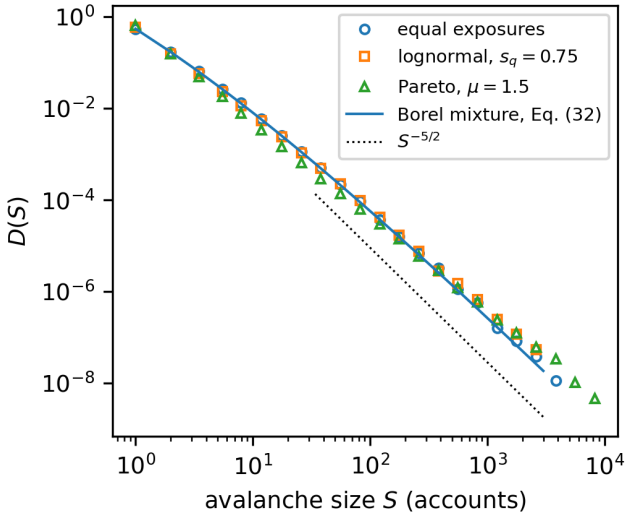


FIG. 7. Size distribution of precursor avalanches, counted in accounts, integrated over the quasistatic ramp up to the macroscopic collapse. Symbols: simulations with $N = 2 \times 10^6$ and 40 cohorts for equal, lognormal, and Pareto exposures. Solid line: Eq. (32) for equal exposures, with no fitted parameter. Dotted line: slope $-5/2$.

Equation (32) and Fig. 7 count avalanche size in accounts because that is the observable for which the equal-exposure Borel prediction is derived. For lognormal and Pareto cohorts, the exposure-weighted avalanche size is the more direct market-impact observable. The present comparison does not claim that the Borel law applies unchanged to that exposure-weighted distribution.

X. PATH-RESOLVED PARTIAL DELEVERAGING

A. Model

It is useful to resolve the path beyond the binary model and allow partial liquidation. Account i holds remaining shares $s_i(t)$ and debt $D_i(t)$ at price $P(t) = 1 - d(t)$. Current assets and equity are

$$A_i = Ps_i, \qquad E_i = A_i - D_i. \tag{33}$$

If $E_i/A_i < m_i$ and the breach has persisted for longer than the account's execution delay, the dollar sale required to restore the maintenance ratio is

$$x_i^* = max\left(0,\ A_i - E_i/m_i\right), \tag{34}$$

with $x_i^* = A_i$ when equity is negative. Only a fraction can be executed in one step,

$$x_i(t) = min\left(x_i^*,\ r_i A_i\right), \tag{35}$$

where r_i is a heterogeneous sale-rate limit. Sale proceeds reduce debt one for one. Let $u(t)$ be the shares sold in the step divided by the initial vulnerable shares. The unabsorbed impact stock evolves as

$$U(t+1) = \delta_I U(t) + u(t), \qquad \delta_I = 2^{-1/h_I}, \tag{36}$$

where h_I is the impact half-life and $h_I = \infty$ gives permanent impact. The drawdown relaxes toward the contemporaneous target,

$$d(t+1) = d(t) + \gamma\left[d_0(t) + \Lambda f\left(U(t+1)\right) - d(t)\right], \tag{37}$$

with f the linear-to-square-root law of Eq. (22). Position sizes and leverage follow Eqs. (3) and (5), maintenance ratios are normal and truncated at four standard deviations, execution delays are Poisson, and sale-rate limits are beta distributed (Table II). The external field is ramped linearly over T steps and then held for 500 steps. The time step is dimensionless. Only ratios among the price-relaxation, execution-delay, impact-memory, and driving times are meaningful. The values are not a calibration. $\Lambda = 0.50$ is chosen so that a transition lies inside the numerical window. Unless stated otherwise, $h_I = 80$ steps.

Binary mean-field reference. For the Table II population, using the same linear-to-square-root impact law and the same truncated distribution of maintenance ratios but complete instantaneous liquidation and permanent impact, the upper mean-field fold occurs at external field 0.10469. The corresponding total drawdown is 0.13683 and the breached exposure at the fold is 0.01891. The path-resolved permanent-impact simulations, which add partial deleveraging, heterogeneous execution delays, sale-rate limits, and finite price relaxation, have large-cohort median transition fields between 0.19537 and 0.19565. With an 80-step impact half-life, the corresponding range is 0.23210 to 0.23523. For these parameters, the path-resolved balance-sheet and execution rules account for the larger first shift away from the binary fold, while impact decay produces a further delay. The decomposition is descriptive rather than additive because the mechanisms interact nonlinearly.

B. Intrinsic cascade duration

Figure 8(c,d) shows the dependence on the ramp duration for $N = 8000$ and 32 cohorts, with the same cohorts used at every T. The 10–90% width of the sold-exposure transition, measured in steps, converges to a constant for slow ramps: it is 32.2 ± 2.7 at $T = 4000$ and 31.8 ± 2.6 at $T = 64,000$ (mean and standard deviation over cohorts), and the average over the five slowest ramps is $\tau_c = 32.4$ steps. The width in field units is therefore

$$\Delta d_0^{10\text{–}90} = \tau_c \, \frac{d_0^{\max} - d_0^{\min}}{T}, \tag{38}$$

which falls from 0.0296 ± 0.0007 at $T = 250$ to $(1.4 \pm 0.1) \times 10^{-4}$ at $T = 64,000$. The observed narrowing is consistent with a discontinuous quasistatic limit, and the observed $1/T$ sharpening is the signature of a cascade with a finite intrinsic duration set by the execution delay, the sale-rate limit, and the price relaxation. For fast ramps ($T \leq 1000$) the width in steps is smaller and the final sold fraction larger (0.915 at $T = 250$ against 0.84 for slow ramps), because the external field continues to rise during the cascade.

C. Cohort-size dependence and stalled ignition

With permanent impact the cohort-to-cohort standard deviation of the transition field, defined as the field at which half of the final sold exposure has been sold, decreases from 0.0043 at $N = 1000$ to 0.0007 at $N = 64,000$ (48 cohorts per size), a fitted exponent of -0.45 [Fig. 8(b)]. With $h_I = 80$ it does not decrease: it is 0.0064 at $N = 2000$, 0.0070 at $N = 8000$, and 0.0051 at $N = 64,000$. The mean transition field (0.232 to 0.237), the final sold fraction (0.83 to 0.85), and the cascade duration (32 to 34 steps) are stable across sizes.

This pattern is a form of stalled ignition. Near the fold the cascade grows slowly, so under transient impact the first recruitment wave can lose price impact before it recruits enough additional sellers. In a subset of cohorts the first wave dies after selling only a few percent of exposure and the macroscopic cascade ignites later. The resulting transition-field distribution is bimodal, and its

width is set by the separation of the two modes rather than by ordinary sampling noise. Figure 8(c) shows one example in which the same cohort transitions at different fields under different ramp durations. No convergence to a unique transition field is established for transient impact over the cohort sizes studied here, so no deterministic large-cohort limit is asserted.

The proposed mechanism was tested directly in a matched 128-cohort control with transient impact. 24 cohorts, or 18.8%, were classified as late ignitions because their transition field lay more than 0.008 above the cohort median. These late cohorts sold 9.6% of exposure on average in the first recruitment wave, compared with 36.9% for the other cohorts. Between that first wave and the later macroscopic ignition, the unabsorbed impact stock fell by a median 69%. A decay greater than 25% occurred in 79% of late cohorts but in only 8.7% of the others. The median quiet interval after the first wave was 243 steps in the late group and zero in the other group. These measurements directly support the interpretation that a weak first wave can lose its own price impact before it recruits the full cascade.

D. Impact memory

To separate suppression from delay, the ramp is extended to $d_0 = 0.45$ at the same rate ($T = 4600$) and results compared with a no-feedback baseline, $\Lambda = 0$, in which 44% of exposure is sold gradually by the end of the ramp and the largest sale within any 200-step window is $J = 0.046$. Figure 9 shows 22 half-lives from 3 steps to permanent impact, three mean execution delays, and 32 cohorts of $N = 8000$ per point.

The largest 200-step event decreases smoothly with decreasing memory, from $J = 0.857 \pm 0.004$ for permanent impact to 0.805 ± 0.017 at $h_I = 80$, 0.613 ± 0.094 at $h_I = 20$, and 0.231 ± 0.009 at $h_I = 3$ (mean delay 5). Even the shortest memory gives an event five times the baseline. The field at which the largest event occurs rises from 0.194 to 0.39. Below a crossover half-life the deleveraging fragments: the number of separate events with $J > 0.1$ per cohort is exactly one for $h_I \geq 28$ and approaches two for $h_I \leq 4$. The crossover half-life is 24, 28, and 40 steps for mean delays of 1, 5, and 15 steps, comparable to the intrinsic cascade duration τ_c. Apart from this shift the curves for the three delays nearly coincide, so the execution delay is not the controlling time scale. Cohort-to-cohort variability is largest for $10 \leq h_I \leq 24$, where the sold fraction at $d_0 = 0.30$ has a standard deviation of 0.16 to 0.20 because some cohorts collapse before that field and some after.

Transient impact therefore does not eliminate the cascade. It delays it to a larger external shock, reduces the size of the largest event, and splits the deleveraging into several events when the half-life falls below the cascade duration. An apparent suppression can be an artifact of a driving window that ends before the delayed event: with the ramp stopped at $d_0 = 0.30$ and $h_I = 20$, some cohorts show only partial deleveraging, and the same cohorts collapse when the ramp continues.

XI. RELATION TO PRIOR WORK AND SCOPE

Threshold-cascade models provide the broad mathematical setting [15–20], and global-load-sharing fiber bundles provide the closest finite-size analog [21, 22]. Financial uses of this analogy include the common threshold framework of Lorenz et al. [23], the finite systemic-risk analysis of Burkholz et al.[24], and the bank-failure model of Biswas and Chakrabarti [25]. In the present model the shared variable is asset price rather than a conserved mechanical load. Failed positions carry unequal dollar exposures, those exposures can correlate with leverage thresholds, price impact is nonlinear, and impact can decay. These features produce the weighted finite-size, heavy-tail, and stalled-ignition results in Secs. VII–X.

Other financial fire-sale models emphasize overlapping portfolios, balance-sheet networks, or recurrent leverage dynamics [8, 10–14, 44]. This one-asset model removes network topology and couples positions through a common price. The local instability is then controlled by exposure density near the threshold, the sensitivity of the maintenance threshold to price, and marginal price impact. Onaga et al. [14] model fire sales on overlapping portfolios. Fosset et al. [45] model liquidity crises generated by feedback among price trends, volatility, and liquidity provision. Kavya [31] already contains the local threshold-density instability criterion, shock-driven discontinuous transitions, and branching-process precursors. The distinct results here are the exposure weighting and leverage correlation, the weighted finite-size variance and prefactors for nonlinear price-mediated load sharing, the heavy-tail dominant-holder regime, and the transient-impact dynamics.

Finite carrying capacity provides a different extension. Caravelli et al. [30] study stochastic growth under finite carrying capacity. For a liquidation model, an analogous demand-side variable would represent the remaining resources of marginal buyers or liquidity providers that can absorb forced sales. Unlike the fixed impact law used here, that capacity would be depleted by past absorption and could later recover. Such a state variable could be especially relevant when nominal market size is large but the effective pool of marginal buyers is limited. Coupling that mechanism to the present absorbing threshold model is left for future work.

The model also has clear empirical limits. Ordinary margined cash-equity positions satisfy Eq. (1). Options, swaps, portfolio-margin accounts, and leveraged exchange-traded products do not share a single maintenance equation and should not be pooled into Eq. (4) without channel-specific state variables and sale weights. Markets with mechanical liquidation rules and observable position-level data, such as perpetual futures and on-

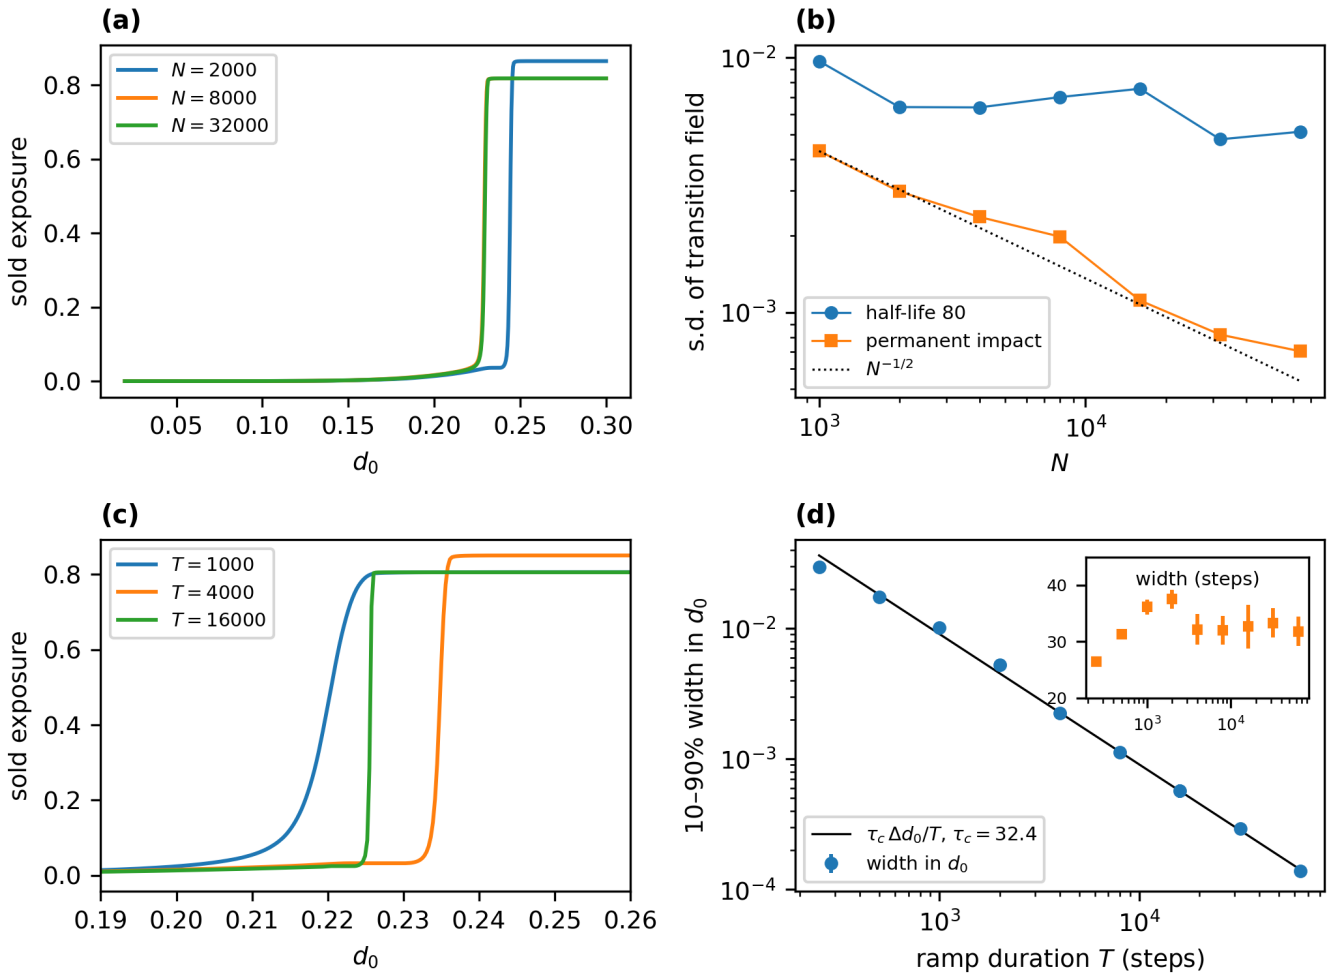


FIG. 8. Path-resolved partial deleveraging with heterogeneous positions, maintenance ratios, execution delays, sale-rate limits, and linear-to-square-root impact with half-life $h_I = 80$ unless noted. (a) Representative sold-exposure paths for three cohort sizes at $T = 3000$. (b) Cohort-to-cohort standard deviation of the transition field for transient and permanent impact. 48 cohorts per point. (c) One cohort at three ramp durations. The late transition at $T = 4000$ is a stalled ignition. (d) 10–90% transition width in field units versus ramp duration with Eq. (38). The inset shows the same width in steps. Error bars are standard deviations over 32 cohorts.

chain collateralized lending, provide direct settings for testing the model. A credible empirical test would estimate vulnerable exposure, its joint distribution with leverage, the maintenance rules, the impact function on a specified horizon, and the initiating shock before the event. It would then classify collapses out of sample rather than infer an unobserved parameter from the occurrence of a crash. If position sizes are heavy tailed, the largest holders must be measured individually rather than summarized only through an aggregate distribution.

The population is also long-only and static. Positions do not enter or leave during the experiment, investors do not reduce risk voluntarily before a maintenance breach, liquidity providers do not respond to the developing order flow, and the impact law does not change with market state. Pre-emptive selling can matter in real deleveraging episodes. It can reduce the exposure left for a later mechanical cascade, but it also adds selling pressure before the formal margin threshold is reached. The net effect cannot be signed without a behavioral response rule. The transition studied here should therefore be interpreted as the mechanically triggered component of a broader deleveraging process, not as a complete model of a market event.

XII. CONCLUSION

A one-asset population of margined positions coupled through price undergoes a macroscopic absorbing liquidation cascade when the product of exposure density near the maintenance threshold, threshold sensitivity to price, and marginal impact reaches one. Exposure weighting, the initiating drawdown, and the maintenance require-

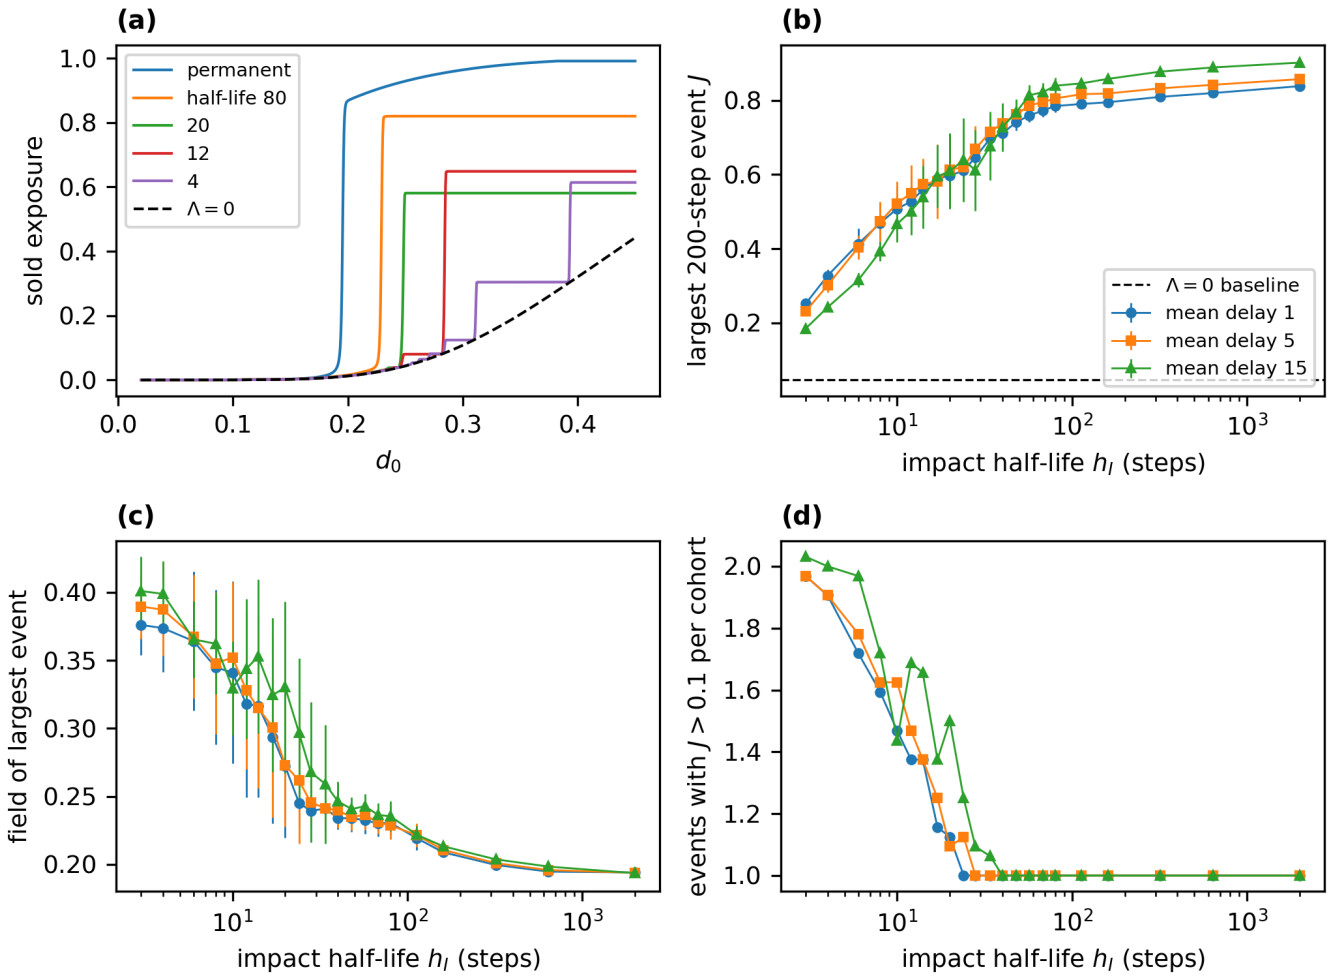


FIG. 9. Effect of impact memory with the ramp extended to $d_0 = 0.45$. $N = 8000$, 32 cohorts per point, error bars are standard deviations over cohorts. The rightmost point in (b) to (d) is permanent impact. (a) Representative paths for mean delay 5. The dashed line is the no-feedback baseline. (b) Largest sale within a 200-step window. (c) External field at which the largest event occurs. (d) Number of separate events with $J > 0.1$ per cohort.

ment shift the boundary analytically, and the fold does not depend on the square-root singularity or on strict concavity of impact. The reversible cusp organizes the onset but does not imply reversible liquidation.

The finite-cohort calculation builds on the classical Daniels-Smith fiber-bundle exponents. The new finite-size result is the exposure-weighted variance, including normalization fluctuations, together with parameter-free prefactors for nonlinear price-mediated load sharing. For heavy-tailed exposures, scaling arguments and simulations indicate a more granular regime in which rare large holders can delay the background instability. The standard-deviation scaling remains a conditional asymptotic prediction because it requires those rare dominant-holder events to produce an order-one shift. The event-independent background-loss check supports the dilution mechanism, but the available simulations do not constitute a precise asymptotic confirmation of the heavy-tail standard-deviation exponent. Precursor account-count avalanches follow the expected Borel-mixture crossover, while the exposure-weighted avalanche distribution remains a separate observable.

The path-resolved simulations qualify the mean-field picture in two ways. The cascade has a finite intrinsic duration, which explains the $1/T$ sharpening under slower driving. Transient impact delays, shrinks, and fragments the cascade when its half-life falls below that duration, and stalled ignitions make the transition field bimodal across cohorts.

No empirical market is calibrated in this study. The numerical parameter sets are theoretical examples, not forecasts for a specific asset. An empirical application would require an estimated threshold distribution and impact response. These inputs determine whether the local gain in Eq. (11) can reach unity, where the low-liquidation state disappears, and how far a finite, granular cohort can deviate from that point.

Appendix A: Derivation of the exposure tilt

Let $z_{th} = \left(\updownarrow_{th}/\bar{\updownarrow} - 1\right)/\sigma$. With Eq. (5), integration over U leaves

$$F(d) = \mathbb{E}\left[e^{\chi Z - \chi^2/2}\,\mathbf{1}\{Z \geq z_{th}\}\right] = 1 - \Phi\left(z_{th} - \chi\right) = \Phi\left(\chi - z_{th}\right), \tag{A1}$$

which is Eq. (6). The same calculation with Q^2 in place of Q gives a tilt of 2χ and a prefactor $\mathbb{E}\left[Q^2\right]/\mathbb{E}[Q]^2 = e^{s_q^2}$, which is Eq. (25).

Appendix B: Fold derivation for power impact

On the equality map, $B = F(d)$. From Eq. (6), $z = \Phi^{-1}(B)$ implies

$$\bar{\updownarrow} = \frac{\updownarrow_{th}}{1 - \sigma(z - \chi)}. \tag{A2}$$

For $d = d_0 + \Lambda B^{\alpha}$,

$$\updownarrow_{th} = \frac{1}{b_0 + c\Lambda B^{\alpha}}, \tag{A3}$$

which yields Eq. (15). Differentiating $\ln \bar{\updownarrow}$ and using $dz/dB = 1/\varphi(z)$ gives

$$\frac{\sigma}{A\varphi(z)} - \frac{c\Lambda\alpha B^{\alpha-1}}{b_0 + c\Lambda B^{\alpha}} = 0, \tag{A4}$$

from which Eq. (16) follows.

Appendix C: Event-driven quasistatic protocol

For each finite cohort the maintenance condition can be written as an individual external-field threshold

$$t_i = \frac{1/\updownarrow_i - m}{1 - m}. \tag{A5}$$

The t_i are sorted once with the associated weights w_i. If a cumulative weighted exposure L has already been liquidated, the next threshold t_{k+1} is reached at

$$d_0 = t_{k+1} - \Lambda f(L), \tag{A6}$$

which is Eq. (23). At that same d_0, L is updated and any additional thresholds satisfying $t_i \leq d_0 + \Lambda f(L)$ are recruited recursively. Equivalently, avalanches start at the running records of D_N and end just before the next record. The avalanche with the largest liquidated exposure defines the finite-cohort collapse field, the takeoff state, and the jump size. This construction eliminates control-grid broadening.

Appendix D: Path-resolved simulation parameters

Cohort-size tests use $N = 1000$ to $64,000$ at $T = 3000$ with 48 cohorts per size. Ramp-rate tests use $T = 250$ to $64,000$ at $N = 8000$ with 32 cohorts, the same cohorts at every T. Impact-memory tests use 22 half-lives at $N = 8000$ and $T = 4600$ with 32 cohorts, the same cohorts at every half-life. The transition field is the field at which half of the final sold exposure has been sold, and the 10–90% width is the interval between the first times at which 10% and 90% of the final sold exposure have been sold.

Appendix E: Code and data availability

No empirical data are used. The Python source code, fixed random seeds, generated numerical tables, the event-independent dominant-holder check, the Section X mean-field reference calculation, and the figure files used in this manuscript are available at https://github.com/bonabeau-gif/Absorbing_Margin_Liquidation_Cascades.

APPENDIX F: Acknowledgments and disclosures

This work received no external funding.

The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

During preparation of this work, the author used Claude Code Opus 5.5 to support simulation coding. ChatGPT 5.6 Sol was used to evaluate and improve the manuscript. The author reviewed and edited all such output and takes full responsibility for the published content.

[1] M. K. Brunnermeier and L. H. Pedersen, Market liquidity and funding liquidity, Rev. Financ. Stud. **22**, 2201 (2009).
[2] J. Geanakoplos, The leverage cycle, NBER Macroecon. Annu. **24**, 1 (2010).
[3] G. Gennotte and H. Leland, Market liquidity, hedging, and crashes, Am. Econ. Rev. **80**, 999 (1990).
[4] A. Shleifer and R. W. Vishny, The limits of arbitrage, J. Finance **52**, 35 (1997).
[5] S. Thurner, J. D. Farmer, and J. Geanakoplos, Leverage causes fat tails and clustered volatility, Quant. Finance **12**, 695 (2012).
[6] F. Caccioli, J.-P. Bouchaud, and J. D. Farmer, A pro-

TABLE II. Reference parameters for the path-resolved simulations. Time is dimensionless.

Quantity	Reference value
Central leverage	2.0
Leverage dispersion	0.16
Log-position dispersion	0.75
Position-leverage correlation	0.40
Maintenance ratio	normal, mean 0.25, s.d. 0.025, truncated at 4 s.d.
Feedback amplitude	0.50
Price relaxation	0.22
Execution delay	Poisson, mean 5 steps (1 and 15 in Sec. X D)
Sale-rate limit	beta, mean 0.25, s.d. 0.10
External field	0.02 to 0.30 (0.02 to 0.45 in Sec. X D), then held 500 steps
Impact law	linear-to-square-root,
Impact half-life	80 steps unless stated
Price floor	

posal for impact-adjusted valuation: Critical leverage and execution risk, Risk 25(12), 74–77 (2012).
[7] R. Cont and L. Wagalath, Running for the exit: Distressed selling and endogenous correlation in financial markets, Math. Finance **23**, 718 (2013).
[8] C. Aymanns and J. D. Farmer, The dynamics of the leverage cycle, J. Econ. Dyn. Control **50**, 155 (2015).
[9] S. Poledna, S. Thurner, J. D. Farmer, and J. Geanakoplos, Leverage-induced systemic risk under Basle II and other credit risk policies, J. Bank. Finance **42**, 199 (2014).
[10] R. Cifuentes, G. Ferrucci, and H. S. Shin, Liquidity risk and contagion, J. Eur. Econ. Assoc. **3**, 556 (2005).
[11] R. Greenwood, A. Landier, and D. Thesmar, Vulnerable banks, J. Financ. Econ. **115**, 471 (2015).
[12] R. Cont and E. Schaanning, Monitoring indirect contagion, J. Bank. Finance 104, 85 (2019).
[13] F. Duarte and T. M. Eisenbach, Fire-sale spillovers and systemic risk, J. Finance **76**, 1251 (2021).
[14] T. Onaga, F. Caccioli, and T. Kobayashi, Financial fire sales as continuous-state complex contagion, Phys. Rev. Research **5**, 043123 (2023).
[15] M. Granovetter, Threshold models of collective behavior, Am. J. Sociol. **83**, 1420 (1978).
[16] D. J. Watts, A simple model of global cascades on random networks, Proc. Natl. Acad. Sci. USA **99**, 5766 (2002).
[17] J. P. Gleeson, Binary-state dynamics on complex networks: Pair approximation and beyond, Phys. Rev. X **3**, 021004 (2013).
[18] J. P. Sethna, K. Dahmen, S. Kartha, J. A. Krumhansl, B. W. Roberts, and J. D. Shore, Hysteresis and hierarchies: Dynamics of disorder-driven first-order phase transformations, Phys. Rev. Lett. **70**, 3347 (1993).
[19] K. Dahmen and J. P. Sethna, Hysteresis, avalanches, and disorder-induced critical scaling: A renormalization-group approach, Phys. Rev. B **53**, 14872 (1996).
[20] J.-P. Bouchaud, Crises and collective socio-economic phenomena: Simple models and challenges, J. Stat. Phys. **151**, 567 (2013).
[21] H. E. Daniels, The statistical theory of the strength of bundles of threads. I, Proc. R. Soc. London, Ser. A **183**, 405 (1945).
[22] S. Pradhan, A. Hansen, and B. K. Chakrabarti, Failure processes in elastic fiber bundles, Rev. Mod. Phys. **82**, 499 (2010).
[23] J. Lorenz, S. Battiston, and F. Schweitzer, Systemic risk in a unifying framework for cascading processes on networks, Eur. Phys. J. B 71, 441 (2009).
[24] R. Burkholz, H. J. Herrmann, and F. Schweitzer, Explicit size distributions of failure cascades redefine systemic risk on finite networks, Sci. Rep. 8, 6878 (2018).
[25] S. Biswas and B. K. Chakrabarti, A fiber bundle model of systemic risk in financial networks, in Systemic Risk and Complex Networks in Modern Financial Systems, edited by V. Pacelli (Springer, Cham, 2025), pp. 179–189.
[26] F. Bucci, M. Benzaquen, F. Lillo, and J.-P. Bouchaud, Crossover from linear to square-root market impact, Phys. Rev. Lett. **122**, 108302 (2019).
[27] B. Tóth, Y. Lempérière, C. Deremble, J. de Lataillade, J. Kockelkoren, and J.-P. Bouchaud, Anomalous price impact and the critical nature of liquidity in financial markets, Phys. Rev. X **1**, 021006 (2011).
[28] J. Donier, J. Bonart, I. Mastromatteo, and J.-P. Bouchaud, A fully consistent, minimal model for non-linear market impact, Quant. Finance **15**, 1109 (2015).
[29] A. S. Kyle, Continuous auctions and insider trading, Econometrica **53**, 1315 (1985).
[30] F. Caravelli, L. Sindoni, F. Caccioli, and C. Ududec, Optimal growth trajectories with finite carrying capacity, Phys. Rev. E 94, 022315 (2016).
[31] A. Kavya, Stop-loss density and liquidation cascades: A threshold theory and its measurement, SSRN 7293038 (2026), doi:10.2139/ssrn.7293038.
[32] R. M. Garcia Seuma, Where does the criticality live? Early-warning signals are event-heterogeneous across seven crypto-perpetual liquidation cascades, arXiv:2607.27070 (2026).
[33] R. M. Garcia Seuma, Measuring the engine of a liquidation cascade: Subcritical branching inside a first-order transition, arXiv:2608.03616 (2026).
[34] R. L. Smith, The asymptotic distribution of the strength of a series-parallel system with equal load-sharing, Ann. Probab. **10**, 137 (1982).
[35] H. E. Daniels and T. H. R. Skyrme, The maximum of a random walk whose mean path has a maximum, Adv. Appl. Probab. **17**, 85 (1985).
[36] H. E. Daniels, The maximum of a Gaussian process whose mean path has a maximum, with an application to the strength of bundles of fibres, Adv. Appl. Probab. **21**, 315 (1989).
[37] S. Janson, G. Louchard, and A. Martin-Löf, The maximum of Brownian motion with parabolic drift, Electron. J. Probab. **15**, 1893 (2010).

[38] P. C. Hemmer and A. Hansen, The distribution of simultaneous fiber failures in fiber bundles, J. Appl. Mech. **59**, 909 (1992).
[39] S. Pradhan, A. Hansen, and P. C. Hemmer, Crossover behavior in burst avalanches: Signature of imminent failure, Phys. Rev. Lett. **95**, 125501 (2005).
[40] L. Eisenberg and T. H. Noe, Systemic risk in financial systems, Manage. Sci. **47**, 236 (2001).
[41] P. Groeneboom and J. A. Wellner, Computing Chernoff's distribution, J. Comput. Graph. Stat. **10**, 388 (2001).
[42] X. Gabaix, Power laws in economics and finance, Annu. Rev. Econ. **1**, 255 (2009).
[43] X. Gabaix, The granular origins of aggregate fluctuations, Econometrica **79**, 733 (2011).
[44] F. Caccioli, M. Shrestha, C. Moore, and J. D. Farmer, Stability analysis of financial contagion due to overlapping portfolios, J. Bank. Finance **46**, 233 (2014).
[45] A. Fosset, J.-P. Bouchaud, and M. Benzaquen, Endogenous liquidity crises, J. Stat. Mech. (2020) 063401.