

The Geometry of Bitcoin

ABSTRACT

Self-reinforcing feedback appears across disciplines, and where it is superlinear it produces not unbounded growth but a finite-time singularity. The classical treatment assumes a fixed feedback exponent: for $dy/dx = k(1 + y^m)$, a singularity occurs if and only if $m > 1$, at the distance $x^* = (1/k)(\pi/m)/\sin(\pi/m)$. **This paper asks what happens when that exponent is not fixed but fluctuates across the critical value.** The question arose from an attempt to determine the exponent of a specific empirical curve, Bitcoin's historical price lows, where a flattening power law, a superlinear trajectory, and a noise-driven one proved difficult to distinguish over the observed range.

The main result is that the mean exponent does not decide the outcome. If the exponent process has a non-vanishing probability of sustained excursions above the critical value, together with a renewal condition, the singularity occurs almost surely **whatever the mean may be, including means below the critical value**, with the time to the event bounded in expectation and its tail geometric. The proof rests on two elementary facts: the state is non-decreasing, so the level reached is never given back, and the duration of supercritical feedback still required to complete a singularity shrinks as that level rises. Numerical results are given for mean exponents down to 0.80. **A complementary result sharpens where the fluctuation must act:** a fluctuation entering the rate rather than the exponent cannot produce a singularity at any amplitude, however persistent.

The second half of the paper concerns identification rather than dynamics. Explicit, checkable membership conditions are stated, including the requirement that the fluctuation reach the exponent and that an irreversible state variable be identified. Simulations of nested instances of the same equation, with code included, illustrate a difficulty that applies to every candidate system: **an observed quantity may follow none of the levels present**, so that an exponent fitted to a record belongs to no part of the system producing it. Two physical domains are examined against these conditions. In tokamak edge instabilities the currently established equation reduces to the same form but with a fixed quadratic exponent and a linear damping, placing them in a distinct class of metastable escape rather than in the one described here; the status of that fixed exponent, which is derived rather than measured, is discussed. In accelerating creep before slope failure, the governing relation used in that field coincides with the equation above, and its exponent is measured above the critical value rather than assumed.

A circuit realizing the equation is described, though not yet built. An oscillator runs between a fixed threshold and a rising one, so that its frequency is inversely proportional to the gap between them, and a second capacitor closes that gap as it accumulates. **The exponent such an arrangement produces is not a design choice.** It follows from how the gap closes and how the observable reads it, through a relation derived here that accounts for every case considered: charging once per cycle gives one value, charging at a constant rate gives another, and the same relation reproduces the exponents of two mechanical arrangements, one of which coincides with the isothermal and adiabatic compressions of elementary kinetic theory. **Every passive arrangement examined proves to be supercritical**, and the reason is visible in the relation itself; reaching the critical value requires a limiting case, and going below it requires the charging law to be imposed rather than found. The circuit is a constructed instance and cannot show that the class is non-empty in nature, but it fixes the exponent by geometry rather than by assumption, and it permits the estimator used throughout to be calibrated against a system whose exponent is known.

No prediction is offered for any system, and the conjectures are labelled as such, in particular whether the fluctuations can be generated by nesting rather than imposed, and whether any real system satisfies the conditions. The framework is put forward as a descriptive instrument for recognizing intermittent superlinearity, together with the observations that would exclude it.

The Geometry of Bitcoin:

Superlinear Feedback and the Finite-Time Singularity Problem

1. THE QUESTION AND ITS SCOPE

Standard growth descriptions of Bitcoin's price history, particularly those built on permanent diminishing returns, contain no superlinear, self-reinforcing term, and they are typically recalibrated as new market eras arrive. This paper asks a narrower question: **whether a self-reinforcing feedback loop can be detected at all**, using Bitcoin's absolute, fixed supply cap of 21 million coins as a reference point. To this end, it compares a tangent-based model, which contains a superlinear feedback term, with the best possible fit of the standard **Power Law** on the same historical data. The Power Law is the mathematical formulation of the diminishing-returns principle, in which price is assumed to grow as a fixed power of time, $P(x) = A \cdot x^\eta$. From this comparison, the paper then estimates an effective feedback exponent for the observed stretch (Section 3.3.1). The geometric divergence between the two models is documented statistically in Section 6, not merely asserted here.

A note on terminology: the title refers to the broader family this model belongs to. **Any system governed by superlinear feedback produces a finite-time singularity** (Section 3.3), and **the tangent function is the member of that family that admits a closed-form solution**. The word "feedback" is deliberate: the Power Law also grows superlinearly in time, but the feedback it implies is *sublinear* (exponent $1 - 1/n \approx 0.82$), whereas the tangent-based model adds a *superlinear* feedback term on top of it (exponent $1 + 1/n \approx 1.18$; Section 3.3). Throughout this paper, **"tangent-based model" refers to this specific closed-form instance**, while the general thesis concerns the superlinear-feedback family as a whole. Section 2.1 shows that this structure does not even require fixed exponents: exponents that fluctuate around the boundary value, or somewhat below it, can produce it as well.

It is worth being precise, however, about where the emphasis of this paper ultimately falls. Bitcoin's fixed supply is not, by itself, advanced here as a sufficient explanation for the geometric divergence documented in Section 6. As Sections 3.1 and 3.2 show, neither Bitcoin's own supply-side contraction nor ordinary demand growth, examined on their own terms, is capable of hosting the kind of self-reinforcing feedback loop that divergence requires — each produces only a mild, feedback-free, logarithmic structure. **The process this paper ultimately points to as the more likely host of that feedback loop is therefore external to Bitcoin's own supply dynamics**. Scarcity's role is best understood as a **necessary precondition** rather than the complete mechanism: it is what allows an asset to register the full effect of that external process undamped, in the absence of any elastic supply response that would otherwise absorb it. This distinction is developed across Sections 3.1 and 3.2 and revisited in Section 8.

Scope of this paper: This is a descriptive, exploratory measurement framework built on Bitcoin's historical price data. It is not a price-prediction tool, not financial advice, and not a claim of physical or economic law. **It does not forecast a price, a date, or any market outcome**, including at the mathematical asymptote described in Section 5. Its parameters are empirically calibrated on a specific historical stretch, are not independently derived, and are subject to the limitations discussed in Section 7. The singularity date implied by these parameters is a property of the calibrated curve, not a prediction. **Ultimately, as Section 8 discusses, this paper's most durable contribution may lie less in the specific price curve than in proposing this framework as a candidate instrument for observing the structural dynamics of the external system referenced in Sections 3.1–3.2. This broader framing is offered as a direction, not as a conclusively established claim.**

2. THE MATHEMATICAL MODEL AND EARLY-STAGE EQUIVALENCE

$$f(x) = A \left(\frac{\tan(B \cdot x)}{B} \right)^n$$

Where:
 x = Time (e.g., Days since Genesis)
 $f(x)$ = Valuation (Price)
 A, B, n = Structural system parameters

This form originated as $\sqrt{\tan(x)}$, the case $A = B = 1, n = 1/2$; the specific parameters used in this paper are the result of calibration, not of the original construction.

A critical structural property of this tangent-based model is its relationship to conventional descriptions of diminishing returns, such as standard power law frameworks. In the network's early history, the two models are not merely similar—they are mathematically equivalent. This underlying behavior can be demonstrated using the Taylor series expansion for the tangent function.

For small values of x , the inner function expands as:

$$\frac{\tan(B \cdot x)}{B} \approx x + \frac{B^2 \cdot x^3}{3}$$

Assuming the parameter B is a sufficiently small positive constant, the squared term B^2 becomes microscopic. During the initial phases of the network's history (when x is relatively small), the cubic term remains virtually negligible. Consequently, the inner function effectively collapses to x , simplifying the overarching equation to:

$$f(x) \approx A \cdot x^n$$

This mathematical reality clarifies why standard power law formulations fit early historical nodes so effectively: **the tangent-based model mathematically contains the power law as its early-stage linear approximation**. Curve-fitting algorithms inherently struggle to isolate the B parameter using only early-stage data because the geometric curvature—the phase transition driving toward the asymptotic boundary—has not yet manifested measurably within the system's noise.

2.1 A Noise-Driven Formulation Without Fitted Exponents

The function above can also be written as the solution of a simple growth equation. Let $y(x)$ be a dimensionless state variable satisfying

$$\frac{dy}{dx} = k(1 + y^m), \quad y(0) = 0, \quad k > 0$$

and let it be mapped to the observed valuation by

$$f(x) = A \left(\frac{y(x)}{B} \right)^n$$

For $m = 2$ and $k = B$, the solution is $y(x) = \tan(B \cdot x)$, and the mapping reproduces the function above exactly. **For a constant exponent, m determines the character of the feedback.** A finite-time asymptote exists exactly when the integral $\int dy / (1 + y^m)$ converges, which holds if and only if $m > 1$. For $m < 1$, the feedback is **sublinear**: for large y , the state grows as a power of time, $y \propto x^{1/(1-m)}$, so f follows a flattening, power-law-type path with no asymptote. For $m > 1$, the feedback is **superlinear**, and y diverges at a finite point. The value $m = 1$ is the boundary between the two: the solution $y(x) = e^{kx} - 1$ grows exponentially but remains finite for every finite x .

The early-stage equivalence holds for every value of m . While $y \leq 1$, the feedback term satisfies $0 \leq y^m \leq 1$ for any $m \geq 0$, so the growth rate stays between k and $2k$. The state therefore grows at a nearly constant rate, and f keeps a power-law-type shape in the early phase whatever the character of the feedback. The feedback exponent only becomes visible once the state has grown large enough for the term y^m to dominate.

Replacing the fixed exponent with noise. Instead of a fitted value, let the feedback exponent fluctuate around a mean value at or below the boundary, while the mapping exponent is held fixed:

$$m(x) = \bar{m} + \varepsilon_m(x), \quad \bar{m} \leq 1 \quad n = 1$$

Here ε_m is a zero-mean noise term, bounded so that the exponent stays within $[0, 2]$. The noise is described by two properties: its **amplitude**, meaning how far it deviates from zero, and its **persistence**, meaning how long a deviation keeps the same sign. **The mapping exponent carries no noise.** With $n = 1$ the observed quantity is simply proportional to the state, so the whole of the stochastic behaviour enters through the feedback exponent m . The equation is solved step by step, with the exponent held constant within each step. With $\bar{m} = 1$, the average feedback sits exactly at the boundary; with $\bar{m} < 1$, **the average feedback is sublinear**.

Without noise, there is no asymptote. With the exponents held at their mean values, the equation is the constant-exponent case above. For $\bar{m} = 1$, the solution grows exponentially; for $\bar{m} < 1$, it grows only as a power law. In both cases it remains finite at every point, and in the early phase it reduces to a power-law-type shape.

With noise, the trajectory alternates between the two kinds of feedback. Whenever $m(x) < 1$, the path locally takes the flattening, power-law-type form. Whenever $m(x) > 1$, it locally takes the accelerating, tangent-type form. **Even when the average feedback is sublinear, the noise supplies superlinear stretches.** Because $dy/dx \geq k > 0$ at every point, **the state of this single equation never decreases**: the gains made during superlinear stretches are kept, while sublinear stretches only slow the growth down. **This applies to the latent state of one instance, not to an observed price series**, whose declines arise between nested instances of the equation rather than within one (Section 7.6.6). As the state grows, a superlinear stretch also needs less and less time to complete an asymptote. For a constant exponent $m > 1$ starting from a level y , the remaining time is at most

$$\frac{y^{1-m}}{k(m-1)}$$

which shrinks toward zero as y grows. Sufficiently persistent noise therefore eventually produces a superlinear stretch that is long enough, relative to the current level, to reach the asymptote, **whether the mean exponent lies at the boundary or below it**. This does not contradict the constant-exponent criterion. That criterion concerns a fixed exponent; with a fluctuating exponent, what decides the outcome is not the mean but the excursions above the boundary.

Numerical simulations of this construction support this. With a mean feedback exponent of exactly 1, **every run with noise reached a finite-time asymptote**. The median time to the asymptote was about 70 420 days with fast, weak noise, and about 8 850 days with slower, stronger noise. With a sublinear mean exponent of 0.96, **all 200 runs still reached an asymptote**, with a median of about 11 290 days and a 5th–95th percentile range of about 5 390 to 22 260 days. With mean exponents of 0.90 and 0.80, every tested run also reached an asymptote. A lower mean exponent delayed the asymptote, while more persistent noise brought it forward: at a mean of 0.80, the median time ranged from about 12 180 days with the most persistent noise tested to about 171 270 days with the least persistent.

Both pictures can therefore appear within a given observation window. The same equation can produce a trajectory whose asymptote lies beyond the window. Within the window, such a trajectory shows only the flattening, logarithmic-looking, power-law-type shape. With a mean exponent of 0.96 and only fast, weak noise, for example, just 7 of 20 simulated runs reached an asymptote within 400 000 days. The same equation can equally produce a trajectory that reaches its asymptote within the window, driven by the superlinear stretches that the noise supplies. Which of the two is observed depends mainly on the amplitude and persistence of the noise, and only secondarily on the mean exponent.

This is the central point of the formulation. The structure described by the function above, with an early resemblance to a power law and a possible finite-time asymptote, does not require specific fitted exponents, nor even an average feedback that is superlinear. It can arise from exponents whose average sits at the boundary value 1, or somewhat below it, once a sufficient degree of persistent noise is present. The constants A , B and n are therefore better read as a description of one observed trajectory than as defining constants of the underlying process. This is a property of the mathematical construction. No claim is made here about what the noise terms represent.

*** Note on the mapping exponent:** The value $n = 1$ is used here primarily to demonstrate structural similarity and to simplify the mathematical framework of the second part of this work. For Bitcoin specifically, a value greater than 1 would more accurately reflect reality. **The choice of n does not decide whether an asymptote exists.** The mapping is monotone, so it can neither create nor remove a singularity, and the criterion $m > 1$ is unchanged. What n alters is how the feedback appears in the observed valuation: the effective exponent seen there is $E = 1 + (m - 1)/n$, so a value $n > 1$ makes the observed feedback look closer to the boundary than the feedback acting on the state itself.

*** Note on the roles of the parameters:** The emphasis on m in this section is deliberate but incomplete. **m decides whether an asymptote can be reached, but it does not act alone.** The rate k is the internal clock of the system: it rescales the horizontal axis without changing the shape of the trajectory, and therefore sets *when* an asymptote is reached. It also sets the scale against which the noise is measured. The persistence of the noise is meaningful only relative to this clock, because the superlinear stretch required to complete an asymptote, $y^{1-m}/[k(m-1)]$, is inversely proportional to k . **The same noise can therefore be sufficient at one rate and insufficient at another**, and the times in days quoted above hold for the rate used in these simulations, not in general. A and n belong to the observation and affect neither the existence nor the location of the asymptote. In short, **m decides whether, k decides when, and the mapping decides how it appears. The exact relations between the parameters, and the precise conditions under which fluctuations of m force an asymptote, are set out in the second part of this work (Sections 2 and 3), which adopts the normalization $B = k$, so that the observed quantity reduces to $A \cdot x$ in the early phase.**

3. THE THEORETICAL CORE: SYSTEM DYNAMICS AND STRUCTURAL REFLEXIVITY

To move beyond pure curve-fitting, this section proposes one possible structural motivation for the valuation function from a small set of economic assumptions. **The approach taken here is an attempt, not a demonstration.** Whether a self-reinforcing feedback loop actually exists cannot be observed directly, so this paper tries to detect it indirectly. **A tangent-based model, which contains a superlinear feedback term, is compared with a baseline without such a term, the Power Law (Section 6.3), and the effective feedback exponent is then estimated from the fitted model (Section 3.3.1).** A closer fit by the tangent-based model is consistent with the presence of such a loop, but does not by itself prove it. Because the same qualitative structure can also arise from fluctuating exponents (Section 2.1), the estimated exponent is best read as an effective value, not as a fixed constant of the process. The argument proceeds in three steps. **Section 3.1 first establishes a negative result: a contraction of Bitcoin's liquid supply cannot, on its own, produce the geometry documented in Section 6. Section 3.2 extends this exclusion to ordinary demand growth. Section 3.3 then introduces the minimal additional ingredient that can — a superlinear feedback term — and shows that it maps exactly onto the empirical function of Section 2.**

The structural question this section addresses is narrower than it may first appear. **The distinction documented in Section 6 — between a feedback-free structure and the superlinear, feedback-driven structure of the tangent-based model — is not a claim about whether price goes up or down.** It is a claim about which side of the price equation, if either, contains a self-reinforcing loop at all: the perpetually expanding external capital supply, or the strictly fixed, inelastic supply of the Bitcoin network, set by protocol at 21 million units as a matter of definition rather than observation. **Zero elasticity removes the dampening that an elastic production response would otherwise apply, but removing a dampening force is not the same as introducing a self-reinforcing one.** A fixed or contracting supply therefore does not, by itself, force any particular price outcome, and the available data show no mechanical link between the two (see the note below). This is a statement about the mathematical structure of the supply side, not a forecast of how Bitcoin's liquid supply will actually evolve, and it holds whatever those dynamics turn out to be. **At least one additional process, such as the reflexivity dynamic of Section 3.3, is needed to connect this structural condition to accelerating price behavior.** If a feedback loop is present in the data at all, it most plausibly operates on the demand side: not because rising valuation is assumed in advance, but because Sections 3.1 and 3.2 exclude the alternatives.

*** This absence of a mechanical link is visible in the available exchange-reserve data for the roughly 2.5-year period for which it could be obtained.** Before March 2024, reserve and price rose together. From March 2024 onward, however, reserves entered a steady declining trend that continued largely unchanged through both the subsequent 2024 bull market and the bear-market decline that followed it — the same downward reserve trend persisted regardless of price direction. This pattern is consistent with the simplest, structure-free depletion case, $L(x) = L_0 - cx$, used as the baseline assumption in Section 3.1, rather than with a supply response that reacts to price. Under the strict reading of the differential equation in Sections 3.1–3.2, a supply trend of this kind does not, by itself, determine the sign of dP/dx without also knowing $F(x)$ — consistent with, not contrary to,

the argument developed there. This is offered as a qualitative, visually-observed illustration from a specific historical window, not a statistical test of the broader claim.

3.1 A Structural Lower Bound: Why Supply Contraction Alone Is Insufficient

To test whether Bitcoin's fixed-supply property is, by itself, capable of hosting a feedback loop of the kind this paper's model requires, consider a simple market-impact framework, in which the rate of price change is proportional to external demand and inversely proportional to available liquid supply:

$$\frac{dP}{dx} = k \cdot \frac{F(x)}{L(x)}$$

where $F(x)$ represents external capital inflow and $L(x)$ represents available liquid supply. **As the simplest possible case — one requiring no additional structure or feedback of its own — consider a supply that depletes linearly** toward a finite exhaustion point x^* :

$$L(x) = L_0 - c \cdot x, \quad x^* = L_0/c$$

Assuming, for simplicity, a constant external demand $F(x) = F_0$, the differential equation can be solved directly by separation of variables:

$$P(x) = P_0 + \frac{k F_0}{c} \cdot \ln \left(\frac{L_0}{L_0 - cx} \right)$$

As $x \rightarrow x^*$, the denominator $L_0 - cx \rightarrow 0$, and $P(x)$ does indeed diverge — but only **logarithmically**. **This is a structurally feedback-free divergence**: each tenfold reduction in the remaining liquid supply adds only a fixed increment, $(k F_0/c) \cdot \ln 10$, to price, regardless of how large price itself has already grown — there is no term through which price's own level feeds back into its future rate of change. Even at 99.99999% of the way to the exhaustion point x^* , the logarithmic factor has reached only about 16. **A linear, constant-rate contraction of supply is therefore mathematically incapable, on its own, of introducing the self-reinforcing structure this paper's model requires** — independently of whether Bitcoin's real liquid supply actually behaves this way, this specific candidate mechanism does not contain a feedback loop.

This conclusion is robust to the exact shape of $L(x)$ near x^* : in general, if liquid supply approaches zero as $L(x) \propto (x^* - x)^r$, integration shows that $P(x)$ diverges as a genuine power law only when $r > 1$ — that is, only if the depletion itself *accelerates* as the exhaustion point is approached, which is itself a feedback-like property of the supply function, not a consequence of price. The case $r = 1$, the simplest, structure-free depletion case considered above, sits exactly at the boundary between a power-law blow-up and no blow-up at all, and produces only the weak logarithmic case derived above.

The same framework also quantifies how demanding a supply-side explanation would have to be. Near its asymptote $x^* = \pi/(2B)$, the tangent satisfies $\tan(B \cdot x) \approx 1/(B(x^* - x))$, so the empirical function of Section 2 diverges as:

$$f(x) \propto (x^* - x)^{-n}$$

Under constant external demand, reproducing this divergence from supply contraction alone would require $dP/dx \propto (x^* - x)^{-(n+1)}$, that is, a liquid supply vanishing as:

$$L(x) \propto (x^* - x)^{n+1}$$

With the published $n = 5.4666$, this corresponds to a depletion exponent of $r \approx 6.5$, compared with $r = 1$ for the simplest, structure-free case considered above. The gap is qualitative rather than marginal: reproducing the model's structure from supply alone would require the depletion function itself to already contain a sharply accelerating, self-reinforcing shape — which is precisely the kind of structural feedback this section set out to test for, not something a plain, unadorned depletion process supplies for free.

A useful structural inference follows from this: a supply side that depletes without any additional, independently-motivated accelerating structure of its own cannot host the feedback loop this model's geometry requires. If such a loop is present in the observed data at all, it must therefore originate predominantly in the external capital-inflow term, $F(x)$, rather than in Bitcoin's own supply-side mechanics. This paper does not take a position on the nature or origin of that external term — doing so falls outside its scope, which is limited to Bitcoin's own price geometry — and treats it here strictly as a structural placeholder required to close the mathematical picture.

The reflexivity term introduced in Section 3.3 is the simplest shape this placeholder can take. If the state variable is written in dimensionless form y , and inflow grows with it as $F = F_0 \cdot (1 + y^2)$ while liquid supply remains approximately constant at L_0 over the relevant interval, the

market-impact equation above reduces exactly to:

$$\frac{dy}{dx} = \frac{k F_0}{L_0} \cdot (1 + y^2)$$

This constrains only the *shape* of the external term, not its nature. **The same market-impact equation therefore contains both cases side by side:** a feedback-free demand term against a linearly depleting supply yields the logarithmic, non-reinforcing structure derived above, while a feedback-carrying demand term against a finite supply yields the self-reinforcing structure developed in Section 3.3 — the same equation can host either, depending entirely on which side, if any, contains the loop.

3.2 A Second Structural Exclusion: Nominal Demand Growth Does Not Change the Order of the Singularity

A natural extension of the argument in Section 3.1 is to test whether a steadily *growing*, rather than strictly constant, external demand term could itself introduce the missing feedback structure, without requiring the reflexivity term proposed in Section 3.3. This can be tested within the same market-impact framework used previously.

In place of the constant demand assumed in Section 3.1, suppose demand grows exponentially at some constant rate i — a steadily growing, but still structurally feedback-free, term:

$$F(x) = F_0 \cdot e^{ix}$$

Combined with the same linear supply contraction established in Section 3.1, $L(x) = L_0 - cx$, the governing equation becomes:

$$\frac{dP}{dx} = k \cdot \frac{F_0 \cdot e^{ix}}{L_0 - cx}$$

As $x \rightarrow x^* = L_0/c$, the exponential term e^{ix} approaches a finite, well-behaved value — it is the supply term, $1/(x^* - x)$, that governs the order of the singularity. Numerical integration confirms this directly: for any rate i , the increase in $P(x)$ across each successive order-of-magnitude approach to x^* stabilizes at a constant value — the defining signature of a logarithmic, feedback-free divergence — with i affecting only the overall scale of the curve, not its underlying structure. Even at $i = 0.20$, an extreme, arbitrarily-chosen growth rate, the divergence remains logarithmic in order.

This allows a specific structural possibility to be excluded: **a linearly contracting supply combined with any merely exponential growth in demand, however rapid, cannot by itself introduce anything beyond the weak, feedback-free divergence already derived in Section 3.1.** Neither Bitcoin's supply contraction alone, nor its combination with a steadily growing but structurally passive demand term, is capable of hosting the self-reinforcing loop this paper's model requires. This reinforces the point stated in Section 1: whatever process is ultimately responsible for that structure is external to Bitcoin's own supply dynamics, and is not reducible to ordinary, non-reinforcing demand growth either.

3.3 The Reflexivity Term and Its Exact Link to the Empirical Function

We can model this feedback directly using differential calculus. The rate of change of the system's state with respect to time (dy/dx) is driven by two distinct forces: a constant baseline representing continuous external inflow, and a reflexivity term proportional to y^2 , **representing price-dependent inflow that compounds as valuation rises**. Here y is a **dimensionless state variable**; its exact relationship to price in USD is given below.

Mathematically, this relationship is expressed by the differential equation:

$$\frac{dy}{dx} = k \cdot (1 + y^2)$$

Where k is a proportionality constant reflecting the structural pressures of the system.

In calculus, the fundamental analytical solution to this specific non-linear differential equation is the tangent function:

$$y(x) = \tan(k \cdot x + C)$$

*** Note on Structural Derivation vs. Empirical Parameters:** The relationship between this solution and the empirical function $f(x) = A(\tan(Bx)/B)^n$ of Section 2 is **exact rather than merely qualitative**. The Genesis condition $y(0) = 0$ fixes $C = 0$, and identifying $k = B$ gives $y(x) = \tan(B \cdot x)$. The empirical function is then a monotone power transform of this solution:

$$f(x) = A \cdot \left(\frac{y(x)}{B} \right)^n$$

The reflexivity equation therefore governs the dimensionless state variable y , while A and n parameterize the mapping from y to price in USD; **dividing by B normalizes the initial slope, so that $f(x) \approx A \cdot x^n$ in the early phase** (Section 2). Expressed directly in price terms, f satisfies:

$$\frac{df}{dx} = n \cdot A^{1/n} \cdot f^{1-1/n} + n \cdot B^2 \cdot A^{-1/n} \cdot f^{1+1/n}$$

The first term alone yields exactly the Power Law $A \cdot x^n$: the Power Law is precisely the trajectory this system would follow in the absence of feedback. The second term is a superlinear feedback term with an **effective exponent of $1 + 1/n \approx 1.18$** at the published $n = 5.4666$; because this exponent exceeds 1, it produces a finite-time singularity at $x^* = \pi/(2B) \approx \text{day } 20,974$. This date is a property of the calibrated curve of Section 6.2, not a forecast; as Section 7.1 shows, it is highly sensitive to the calibration.

Why a superlinear term is needed at all is established in Sections 3.1 and 3.2, not here: reproducing the fitted curve from supply contraction alone would require a depletion exponent of $r \approx 6.5$, against $r = 1$ for the simplest, structure-free case. This section shows only that a minimal feedback term of this kind is **sufficient** to generate the observed geometry — the tangent-based curve tracking marginally flatter than the Power Law in the network's early history, then measurably steeper in its later history, as confirmed by the MACE comparison in Section 6.3 and directly visible in the Pine Script visualization of Section 6.5 — and makes its mapping onto the empirical function explicit. **It does not derive the numerical values of A , B and n , or the choice of the quadratic form over other superlinear closed forms, from first principles.**

It is worth noting that the emergence of a finite-time singularity is not narrowly dependent on the specific quadratic form of the reflexivity term. Any superlinear feedback assumption ($y' = k(1+y^m)$ for $m > 1$) produces the same qualitative outcome — a mathematical collapse of the time axis. The quadratic case ($m=2$) is adopted here specifically because it yields a closed-form, empirically tractable solution — the tangent function — rather than because market structure is assumed to dictate this exact exponent.

This choice is not arbitrary in an even stronger sense, however: it was not selected in isolation and then merely hoped to fit the data. The tangent-based functional form was independently supported in Section 6 through direct empirical curve-fitting against historical price nodes, where it demonstrated a measurably tighter fit (lower Mean Absolute Centered Error) than the best possible Power Law calibrated on the same data. The quadratic form should not, however, be read as a claim that price feeds back on itself with exponent 2. **Through the transform $f = A(y/B)^n$, the effective feedback exponent in price terms is $1 + 1/n \approx 1.18$;** the value 2 in the latent equation is the choice that yields a closed-form solution, while the empirically fitted n maps to this effective exponent — a figure whose precise value is only as reliable as the specific historical data used to derive it, as Section 7.1 illustrates. This paper does not, and cannot, assess whether this effective exponent will continue to hold going forward; only that its current value rests on a fitted, not independently derived, quantity. Section 6's fit comparison supports this functional form empirically, but that is a statistical finding, not a derivation of the exponent from market structure.

*** Note on Market Reflexivity:** In traditional finance, the concept of "reflexivity" was famously pioneered by George Soros to describe cognitive and behavioral feedback loops—where investor perception alters market fundamentals, which in turn validate the perception. However, the reflexivity modeled in this framework is strictly *structural and mathematical*, not psychological.

The two accounts are not unrelated, however. **The measurable divergence this paper documents between the tangent-based model and the Power Law (Section 6.3) is coherent with, rather than independent of, Soros's account:** a market whose participants act on reflexive perception — where the expectation of further gains draws in demand that helps validate the expectation — would be expected to leave behind exactly this kind of structural signature, a departure from the diminishing-returns baseline in the direction of self-reinforcing acceleration. This paper does not claim to have measured Soros's mechanism directly, or to have established it as the specific source of the effect documented here; it notes only that the structural direction of the statistical result and the qualitative prediction of reflexivity theory, arrived at independently, point the same way.

3.3.1 Exact Derivation of the Effective Feedback Exponent

To ensure complete structural transparency, the explicit transformation from the dimensionless state variable y to the price domain $f(x)$ is derived below. The base equations are the dimensionless reflexivity condition and the price mapping function:

$$\frac{dy}{dx} = B \cdot (1 + y^2) \quad \text{and} \quad f = A \cdot \left(\frac{y}{B} \right)^n$$

Applying the chain rule to differentiate f with respect to x gives:

$$\frac{df}{dx} = n \cdot A \cdot \left(\frac{y}{B} \right)^{n-1} \cdot \frac{1}{B} \cdot \frac{dy}{dx}$$

Substituting the reflexivity condition $B \cdot (1 + y^2)$ for dy/dx eliminates the denominator B :

$$\frac{df}{dx} = n \cdot A \cdot \left(\frac{y}{B} \right)^{n-1} \cdot (1 + y^2)$$

Distributing this across the terms in the parenthesis separates the baseline from the feedback mechanism:

$$\frac{df}{dx} = n \cdot A \cdot \left(\frac{y}{B} \right)^{n-1} + n \cdot A \cdot \left(\frac{y}{B} \right)^{n-1} \cdot y^2$$

To express the derivative purely in terms of price f , the initial mapping function is inverted to isolate the state variables:

$$\frac{y}{B} = \left(\frac{f}{A} \right)^{1/n} \quad \text{and} \quad y = B \cdot \left(\frac{f}{A} \right)^{1/n}$$

Substituting these inverted relations back into the first distributed term (the Power Law baseline) yields:

$$n \cdot A \cdot \left(\frac{f}{A} \right)^{(n-1)/n} = n \cdot A^{1/n} \cdot f^{1-1/n}$$

Substituting into the second term (the superlinear reflexivity) yields:

$$n \cdot A \cdot \left(\frac{f}{A} \right)^{1-1/n} \cdot B^2 \cdot \left(\frac{f}{A} \right)^{2/n} = n \cdot B^2 \cdot A^{-1/n} \cdot f^{1+1/n}$$

Recombining these components produces the final exact differential equation evaluated in price terms:

$$\frac{df}{dx} = n \cdot A^{1/n} \cdot f^{1-1/n} + n \cdot B^2 \cdot A^{-1/n} \cdot f^{1+1/n}$$

Using the empirically calibrated parameter $n = 5.4666$, the exponent of the feedback term evaluates to $1 + 1 / 5.4666 \approx 1.1829$. Because this exponent strictly exceeds 1, it shows that, for this calibration, the superlinear character of the feedback carries over from the dimensionless state space into the observed price domain.

3.4 A Structural Observation: The Biological Converse

From a purely mathematical perspective, it is worth noting that the differential equation proposed in this framework shares a precise structural relationship—an exact sign inversion—with classical biological models of population growth.

In biology, the standard logistic growth model introduces a negative feedback loop: as a population (y) expands, it consumes the finite resources of its environment, limited by a physical carrying capacity (K). Mathematically, this is expressed as:

$$\frac{dy}{dx} = r \cdot y \cdot \left(1 - \frac{y}{K} \right)$$

This negative feedback forces the growth rate to eventually approach zero, resulting in a horizontal asymptote (a flattening S-curve) where the population stabilizes below the carrying capacity K .

To construct a rigorous, rather than merely rhetorical, counterpart to this model, we can take the exact sign inversion of the logistic equation—flipping the restraining term into an accelerating one:

$$\frac{dy}{dx} = r \cdot y \cdot \left(1 + \frac{y}{K} \right)$$

Structurally, this is not merely an analogous or loosely inspired formulation—it is the **literal converse of the logistic equation**, produced by inverting the sign of the single term separating unconstrained growth from self-limiting growth. Where the logistic model's carrying capacity K acts as a dampening ceiling, this inverted structure transforms K into an accelerating pressure: as y grows, the very term that once suppressed further growth now compounds it. Solving this equation confirms that, like the reflexivity-based model introduced in Section 3.3, this exact converse of the logistic equation likewise produces a finite-time singularity rather than a stabilizing plateau.

This structural relationship is offered as an illustration, not as support for any particular mechanism. In the logistic model, the term containing K is what makes growth self-limiting; in its converse, the same term makes growth self-reinforcing. **The external capital-inflow term treated as a structural placeholder in Section 3.1 would play the role of this reinforcing term**, not named or identified further here. Rather than approaching a horizontal limit where growth flattens, the inversion drives the geometric curvature upward into a vertical asymptote. While the exact exponents and constants of Section 3.3's reflexivity term remain a modeling simplification (Sections 7.1 and 7.5), this sign-inverted counterpart to a well-established biological model shows that the qualitative behavior, a finite-time singularity, is not an artifact unique to this framework. **It follows from the presence of a self-reinforcing term, not merely from the absence of a dampening one**, which is the same distinction drawn in Section 3.

4. STRUCTURAL VALIDATION VIA TEMPORAL TRANSLATION (PHASE SHIFTING)

To rigorously test whether the derived asymptotic geometry represents a genuine systemic property rather than a mathematically overfitted baseline, we can apply a simple phase shift—a horizontal translation along the time axis. By shifting the original baseline function forward in time by a constant offset, we can examine how this translated geometry interacts with the market's upper cyclical extremes.

$$f_{shift}(x) = f(x + \Delta t)$$

Where Δt represents a constant positive time offset.

The evidentiary weight of this validation method relies on geometric alignment and statistical probability. If an arbitrarily chosen, translated mathematical function happens to intersect or bound a single cyclical peak, it proves absolutely nothing. Intersecting two distinct peaks could still be reasonably dismissed as a geometric coincidence.

However, the statistical significance scales non-linearly with each additional point of contact. If the exact same rigid curvature—shifted by a single, constant temporal offset—consistently aligns with or bounds the upper cyclical extremes across multiple, fundamentally distinct market eras, the probability of random coincidence drops drastically. If such a curve were to accurately map tens of thousands of data points or multiple macro-cycle limits, it would substantially shift the balance of probability away from coincidence and toward a structural mathematical rule.

While no model can claim absolute deterministic certainty, this temporal translation serves as a substantive structural test. When the overarching geometry naturally fits both the macro-bottoms (the baseline) and the macro-tops (the shifted boundary) **without requiring any alteration to the function's internal curvature**, it lends meaningful empirical support to the idea that tangent geometry plays a role in shaping the asset's macro-dynamics.

5. THE MATHEMATICAL ASYMPTOTE, MARGINAL PRICING, AND PHASE TRANSITION

By its fundamental nature, the tangent function possesses a structural boundary where it reaches a vertical asymptote. In a generalized, parameter-free form, this mathematical limit is defined as follows:

$$\lim_{x \rightarrow \pi / 2B} f(x) = \infty$$

Note: this is not a price forecast. A literal infinite price cannot exist in any real market. This limit describes a mathematical boundary of the function's domain, not a target valuation — see below for what it does describe.

The mathematical asymptote is read here instead as **a possible structural boundary**: a point at which, if the feedback loop described in Section 3 genuinely exists, the assumptions underlying continuous, liquid price discovery, most notably the presence of a willing marginal seller at any given price, may no longer hold in their current form.

What would happen at or beyond that point lies outside the scope of this paper. This paper does not examine what would follow if such a feedback loop genuinely exists and meets a shrinking pool of available liquid supply. The asymptote should therefore not be read as a forecast of any valuation, date, or economic outcome. The only claim advanced here is narrower and conditional: **if such a loop exists, the mechanism of marginal pricing may run into difficulties as available liquid supply approaches zero**, and the vertical asymptote of the tangent function is intended to represent this possible instability geometrically.

This macroscopic dynamic can be conceptually viewed through the lens of *econophysics*, drawing a structural analogy to statistical mechanics. As explored in the foundational work of researchers like R. N. Mantegna, H. E. Stanley, and Didier Sornette (who specifically modeled finite-time singularities in complex market systems), daily price fluctuations are highly chaotic—much like the unpredictable, microscopic collisions of gas molecules in a closed system. Yet, within this framework of extreme micro-volatility, the macroscopic trajectory of a system does not need to be strictly deterministic to carry a definite structural tendency. The noise-driven formulation of Section 2.1 illustrates this: fluctuating exponents can still produce an asymptotic structure, even though the timing of the asymptote remains highly variable.

Finally, it is essential to address the domain of the equation. Strictly mathematically, the tangent curve is periodic, repeating infinitely across the coordinate system. However, this economic model defines the Genesis Block as the absolute origin (zero time, zero value). Therefore, only the primary, initial branch of the tangent function holds any physical or chronological significance. Subsequent repeating branches of the function exist in mathematical theory but carry no valid meaning for the asset's real-world lifecycle.

6. HISTORICAL APPLICATION, VISUAL COMPARISON, AND PARAMETERIZATION

The parameterization below describes one specific historical stretch, not the process as a whole. It is calibrated on seven tangency nodes between Day 745 and Day 6388, and its validity does not extend automatically beyond that stretch. As Section 2.1 shows, the same qualitative structure can arise from fluctuating exponents without any fitted values. **Fixed parameters are used here out of necessity:** a quantitative estimate of the effective feedback exponent (Section 3.3.1) requires a concrete curve, fitted to a concrete stretch of data. The resulting exponent is therefore best read as **an effective value describing that stretch as a whole**, not as a constant of the underlying dynamic.

To transition from the theoretical framework to observed market action, the mathematical model must be mapped directly onto live market data. However, establishing this fit requires an objective methodology for defining historical "support nodes."

A note before proceeding: the specific numerical parameterization derived below is a statistical curve fit, not a claim of physical law, and it can be broken. Section 7.3 examines this directly, treating *this particular parameterized curve* as a measurement structure with a limited calibration range, distinct from the *underlying structural dynamic* it attempts to describe (Section 3). An extreme enough market event could end the validity of the parameters below, and of the effective exponent derived from them, without deciding anything about the broader thesis in either direction. Readers are encouraged to keep this distinction in mind throughout the following sections.

6.1 Methodology of Tangency and Constraint Satisfaction

A common pitfall in financial modeling is arbitrary curve-fitting—selecting random local lows simply because they visually align with a desired trajectory. To avoid this, our historical nodes are defined by a strict mathematical condition: **tangency**.

A point qualifies as a valid structural support node if it satisfies two simultaneous criteria on a logarithmic scale:

1. **Value Equivalence:** The actual price and the model's value are virtually identical.
2. **Slope Equivalence:** The rate of change (momentum or derivative) of the actual price matches the slope of the model's curve.

Visually, this means the mathematical curve does not awkwardly slice through the price action; instead, it smoothly "kisses" the lower boundary of the price data before separating. Because of this dual requirement, true mathematical touch points do not always land on the absolute lowest daily wick of a bear market; they often manifest as lateral contact points where the market's momentum perfectly aligns with the function's structural geometry.

The seven tangency nodes used to calibrate this model were themselves subjectively identified through visual inspection — this is not disputed. What is not subjective, however, is the outcome once those seven points are fixed: the resulting curve is heavily over-constrained, since it must simultaneously avoid ever exceeding price across multiple, independent multi-year market eras. This is a **constraint-satisfaction condition, not a free optimization** — arbitrary alternative candidate points (for example, the 2018 or 2022 bear-market bottoms) could not have been substituted without violating the non-crossing condition elsewhere in the historical record.

This can be tested directly. The 2018-12-15 bear market low (~\$3,122) and the 2022-11-21 post-FTX low (~\$15,479) were never used to calibrate any parameter of this model. Evaluating the currently published curve at these two dates shows it remains below both realized prices (by 37.6% and 11.0%, respectively) — the curve was never at risk of being crossed by data it had no knowledge of. This does not constitute proof of the model's validity, but it is a genuine, independently checkable consistency test: the seven-node calibration did not merely fit seven arbitrary points — **it produced a curve that also respects historical price action outside its own calibration set.**

6.2 Calibrated Parameters and Baseline Functions

Applying this tangency methodology across Bitcoin's macroeconomic history yields the following integrated system parameters for the tangent-based model. These constants describe the calibrated historical stretch closely, serving as a descriptive geometric fit rather than a guaranteed deterministic law.

$$f(x) = A \left(\frac{\tan(B \cdot x)}{B} \right)^n$$

Tangent Parameters:
x = Days since Genesis Block (2009-01-03)
f(x) = Price in USD
A = 5.8535 × 10⁻¹⁷
B = 7.4893 × 10⁻⁵
n = 5.4666

The Standard Power Law Model (Optimized Baseline):

To provide a rigorous, non-arbitrary comparison, the Power Law parameters provided below were independently derived via Ordinary Least Squares (OLS) log-log linear regression on the exact same seven tangency nodes used to calibrate the tangent-based model.

$$P(x) = A_{PL} \cdot x^{n_{PL}}$$

Optimal Power Law Parameters:
A_{PL} = 2.0616 × 10⁻¹⁷
n_{PL} = 5.6218

6.3 Statistical Divergence and Historical Comparison: Tangent-based model vs. Power Law

The comparative advantage of the tangent function model becomes visible when visualized alongside the optimized Power Law using the same historical dataset. Even when the Power Law is mathematically forced into its best possible fit (OLS regression), it struggles to capture the structural curvature of the asset over time. It systematically overprices the middle eras while missing the geometric acceleration of recent macro-bottoms.

Identified Tangency Nodes & Model Comparison:

Days (x)	Actual (USD)	Tangent f(x)	Tangent (Δ ln)	Power Law P(x)	Power Law (Δ ln)
745	0.313	0.296	-0.057	0.289	-0.080
1049	2.05	1.931	-0.060	1.978	-0.036
1242	4.98	4.882	-0.020	5.112	0.026
1324	7.05	6.940	-0.016	7.323	0.038
2425	198	197.97	-0.000	219.91	0.105
4087	3850	3848.87	-0.000	4136.64	0.072
6388	57735	57716.98	-0.000	50938.45	-0.125

Statistical Summary (Centered Residuals)

For explicit methodological transparency, the Mean Absolute Centered Error (MACE) referenced in the summary is calculated by first defining the individual logarithmic residuals (r_i), determining their arithmetic mean (μ) to isolate systemic bias, and then computing the average of the absolute deviations from this mean:

$$MACE = \frac{1}{N} \sum |r_i - \mu|$$

Where:
N = Total number of tangency nodes
r_i = ln(Model) - ln(Actual) for each given node
μ = The arithmetic mean of all r_i values

By design, the OLS Power Law regression centers its Mean Error at ~0.000. For a direct structural comparison, we isolate the shape accuracy of both curves by mean-centering their residuals (shifting the baseline to a true zero-error mean). The resulting **Mean Absolute Centered Error for the Power Law is ~0.069 (6.9%)**. In contrast, the tangent-based model achieves a **Mean Absolute Centered Error of approximately 0.021**

(2.1%). This suggests that the inherent geometry of the tangent function maps the observed macro-bottoms more than three times more accurately than the best mathematically possible linear log-log fit.

This asymmetry can be demonstrated directly. If the Power Law is instead forced to pass exactly through the two most recent nodes (Day 4087 and Day 6388) — the same nodes the tangent-based model tracks almost perfectly without being specifically forced to — its fit to the earlier historical nodes collapses, with log-residuals of roughly -0.9 to -0.5 (the curve running roughly 40–60% below the realized price) at Days 745 through 1324. The tangent-based model faces no equivalent trade-off: it achieves near-zero error at the recent nodes while remaining within 5-6% at the early ones, simultaneously. **This is a genuine structural asymmetry between the two functional families, not merely an artifact of the tangent model's additional free parameter** — a two-parameter Power Law cannot satisfy both epochs at once, regardless of how its parameters are chosen, while the tangent-based family can. This illustrates that **the two curves' early-stage resemblance is not evidence of equivalence, and the structural difference between them becomes visible precisely where the two curves are forced to account for the network's more recent history**. A visual comparison of the two curves against realized price action is available via the TradingView Pine Script referenced in Section 6.5.

*** Note on Error Measurement and R²:** Traditional metrics like the Coefficient of Determination (R²) are intentionally omitted here. In hyper-exponential datasets spanning multiple orders of magnitude, standard R² heavily overweights absolute errors at late-stage high valuations, allowing models to artificially inflate accuracy by overfitting recent data while ignoring massive percentage errors in early stages. Instead, accuracy is evaluated strictly using individual logarithmic residuals ($\ln(\text{Model}) - \ln(\text{Actual})$), providing a symmetric measure that weights all eras equally.

*** Note on Data Provenance:** To ensure consistency, historical pricing data from Day 1242 onward is sourced directly from the Bitstamp BTC/USD historical aggregate, representing one of the earliest and most continuous exchange records. Data points prior to Day 1242 rely on earlier, fragmented market aggregates (e.g., Mt. Gox or historical synthetics) and inherently contain a slightly higher degree of statistical noise due to early-stage liquidity constraints.

6.4 Temporal Translation (Phase Shift) and Cyclical Nodes

As theoretically outlined in Section 4, applying a constant temporal offset to the baseline geometry allows us to test the structural validity of the model against the asset's macro-cycle tops. Using a phase shift of $\Delta t = 2000$ days, the identically curved geometry interacts with three major cyclical extremes.

While isolating only a small number of contact points does not constitute absolute mathematical proof of a universal rule, their precise geometric alignment with the unmodified curvature of the baseline function shifts the balance of probability away from mere coincidence.

Days (x)	Actual Top (USD)	Shifted Model $f(x + 2000)$	Log Residual ($\Delta \ln$)
3270	19666	17420	-0.121
4452	61782	61525	-0.004
4484	64895	63513	-0.022

*** Note on Phase Shift Accuracy:** Applying the same Mean Absolute Centered Error (MACE) methodology to these three cyclical tops yields a shape variance of approximately **0.048 (4.8%)**. While a three-point sample size is inherently too small to draw definitive statistical conclusions, this minimal geometric error further demonstrates the structural alignment of the tangent function even when translated blindly across multi-year market extremes.

6.5 Independent Verification and Data Visualization

For independent verification against live market data, readers can apply the following Pine Script to overlay the baseline and temporal shift functions directly onto a logarithmic Bitcoin chart on TradingView (simply add the script to your chart as an indicator, and it will overlay directly on the price data):

```
//@version=6
indicator(title="Tangent-Based Model & Power Law Models", shorttitle="BTCtgX", overlay=true)

A = input.float(5.8535e-17, title="A", group="Tangent-Based Model Parameters")
B = input.float(7.4893e-5, title="B", group="Tangent-Based Model Parameters")
n = input.float(5.4666, title="n", group="Tangent-Based Model Parameters")

APwL = input.float(5.8535e-17, title="A", group="Power Law Model Parameters (Equivalent)")
nPwL = input.float(5.4666, title="n", group="Power Law Model Parameters (Equivalent)")

APL = input.float(2.0616e-17, title="A", group="Optimized Power Law Parameters (OLS)")
nPL = input.float(5.6218, title="n", group="Optimized Power Law Parameters (OLS)")

showHigh = input.bool(true, title="Show Shifted Curve (High)", group="Display Settings")
showLow = input.bool(true, title="Show Base Curve (Low)", group="Display Settings")
```

```

showPowerLaw = input.bool(true, title="Show Equivalent Power Law Curve", group="Display Settings")
showOptPowerLaw = input.bool(true, title="Show Optimized Power Law Curve", group="Display Settings")

timeShift = input.int(2000, title="Time Shift (Days)", minval=0, group="Time Settings")

genesis_time = timestamp(2009, 1, 3, 0, 0, 0)

daysElapsed = (time - genesis_time) / 86400000

var float modelPrice = na
var float modelPriceShift = na
var float powerLawPrice = na
var float optPowerLawPrice = na

if daysElapsed > 0
    angle = B * daysElapsed
    baseVal = math.tan(angle) / B
    modelPrice := A * math.pow(baseVal, n)

    daysElapsedShift = daysElapsed + timeShift
    angleShift = B * daysElapsedShift
    baseValShift = math.tan(angleShift) / B
    modelPriceShift := A * math.pow(baseValShift, n)

    powerLawPrice := APwL * math.pow(daysElapsed, nPwL)

    optPowerLawPrice := APL * math.pow(daysElapsed, nPL)

isBTC = syminfo.ticker == 'BTCUSD' or syminfo.ticker == 'BTCUSDC' or syminfo.ticker == 'BTCUSDT'

plot(isBTC and showLow ? modelPrice : na, color=color.green, linewidth=1, title="Baseline (Tangent)")
plot(isBTC and showHigh ? modelPriceShift : na, color=color.red, linewidth=1, title="Shifted (Tangent)")
plot(isBTC and showPowerLaw ? powerLawPrice : na, color=color.gray, linewidth=1, title="Power Law (Equivalent)")
plot(isBTC and showOptPowerLaw ? optPowerLawPrice : na, color=color.purple, linewidth=1, title="Power Law (Optimized)")

```



Figure 1.

Green: Tangent-Based Curve (Baseline)
Red: Tangent-Based Curve (Shifted)
Gray (in the Pine Script): Equivalent Power Law Curve
Purple (in the Pine Script): Optimized Power Law Curve

7. LIMITATIONS

A frequent objection to frameworks of this kind is that they are difficult to falsify. For a single, fixed asymptotic curve, however, this concern is somewhat overstated: **there is no indefinite middle ground between validation and falsification**. If a given calibrated curve is genuinely asymptotic, continued observation will, sooner or later, either keep confirming that curve or decisively break it — a fixed asymptotic curve cannot remain ambiguous indefinitely.

Mathematically, this lack of indefinite ambiguity can be formalized by evaluating the absolute divergence between the asymptotic model $f(x)$ and a continuous, non-asymptotic alternative, such as the standard Power Law $P(x)$. If $f(x)$ possesses a vertical asymptote at a finite time x^* , while $P(x)$ remains finite, the distance between them mathematically must diverge to infinity:

$$\lim_{x \rightarrow x^*} |f(x) - P(x)| = \infty$$

By the formal definition of a limit, this guarantees that for any arbitrarily large acceptable error margin E , there exists a critical time threshold $x_E < x^*$ beyond which the divergence strictly exceeds E :

$$|f(x) - P(x)| > E \quad \text{for all } x > x_E$$

In simpler terms: the geometric separation between the finite Power Law curve and the singular tangent curve accelerates so violently as $x \rightarrow x^*$ that they cannot mathematically remain within a plausible margin of error of one another. The passage of time forcefully breaks the statistical overlap, forcing a definitive resolution.

It is important to be clear about what this eventual separation does, and does not, establish. It guarantees only that these two specific curves cannot remain indistinguishable forever — it is a statement about the long-run behavior of two fixed mathematical objects, not a validation or falsification test for this paper's parameterized model. **Continued tracking of this paper's specific curve, even if it happened to hold precisely, would not by itself confirm that curve's validity**, since the space of curves capable of producing such tracking is considerably larger, and easier to satisfy after the fact, than a single confirmed prediction would suggest. The remainder of this section develops this distinction — between *this paper's specific parameterization* and the *general family of curves* it instantiates — in stages: an illustration of parameter sensitivity (7.1), the relativity of disruption across nested time-horizons (7.2), an extreme disruption viewed from the perspective of the curve as a measurement structure (7.3), the domain within which the framework applies (7.4), and the identifiability of the latent exponent (7.5). Section 7.6 then turns to a simulation direction in which the exponent itself fluctuates.

7.1 A Worked Example of Parameter Instability

To illustrate this sensitivity with a concrete case: an independent refitting of this model using seven refined historical price points (rather than the original six) produced a materially different parameter set — $A = 5.8535 \times 10^{-17}$, $B = 7.4893 \times 10^{-5}$, $n = 5.4666$, compared to the originally published $A = 10^{-16}$, $B = 8.3 \times 10^{-5}$, $n = 5.395$. Despite B alone shifting by roughly 10%, the implied singularity date moved from Day 18,925 (~2060) to Day 20,974 (~2066) — a shift of several years arising from what appears, at the level of individual parameters, to be a minor refinement.

This is a direct empirical demonstration of a broader point: **the model's fit to historical nodes is comparatively stable** under reparameterization — both parameter sets track the seven calibration nodes closely — **while the extrapolated asymptote is considerably more sensitive**. The geometric shape of the curve across the observed historical range remains broadly consistent between refits; the precise forward-projected date does not share that stability. **Consequently, the specific parameters presented in this paper (A, B, n) should not be treated as fixed, immutable constants** — they are the current best descriptive calibration, subject to revision as further data emerges. The same caution applies to the reflexivity exponent itself: **the choice of y^2 in Section 3.3 is a modeling simplification, not a fixed physical constant**, and is examined further in Section 7.5.

7.2 Nested Time-Prefixes, Black Swans and Antifragility

Every observation takes place within a given **time-prefix**: a window whose length determines which events appear ordinary and which appear extreme. An event that is rare, has a large impact, and could not have been anticipated from within that window is what Taleb calls a **Black Swan** (Taleb, *The Black Swan*, 2007). As Taleb emphasizes, this designation depends on the observer's position and knowledge. **The same event can be a complete surprise within one frame and an expected occurrence within another.**

A fractal-like structure. The framework of this paper assumes that time-prefixes are nested inside one another, and that similar dynamics recur across these scales. This view is closely aligned with Benoit Mandelbrot's work on self-similarity and fractal geometry (Mandelbrot, *The Fractal Geometry of Nature*, 1982). Studying river flows and precipitation records, Mandelbrot and Wallis described two properties of such systems: extreme events can be very extreme (the "Noah Effect"), and runs of unusually high or low values can be very long (the "Joseph Effect") (Mandelbrot & Wallis, *Water Resour. Res.* 4, 909, 1968). The simulation of Section 7.6 shows a qualitative counterpart to the second property: **the more persistent the fluctuations, the sooner the finite-time singularity arrived**. The persistence of its noise is simple and finite, however, and does not reproduce the long-range dependence of Mandelbrot's self-similar models.

Nested functions as a possible source of noise. In this fractal-like picture (Part II, Section 4.3, Simulation), the noise of a slower base function would be supplied, at least in part, by the smaller, faster functions nested inside it (Section 7.6.6). Each nested function accelerates, reaches its own asymptote and ceases to exist, while the observed level returns to the base beneath it. Seen from within the short time-prefix of a nested function, its end is an extreme, terminal event. Seen from the time-prefix of the base function, the same event is only one of many excursions that together make up its noise. **Events of this kind, which do not break the larger time-prefix, strengthen its structure rather than weaken it.** In the construction, each excursion leaves the underlying state permanently higher, and the accumulated excursions drive the base function toward its own asymptote (Sections 7.6.5 and 7.6.6).

Consistency with Taleb's framework. This is antifragility in Taleb's sense: a system that gains from variability, shocks and disorder, rather than merely withstanding them (Taleb, *Antifragile*, 2012). Two further elements of that framework are preserved here. First, Taleb notes that the antifragility of a larger system often rests on the fragility of its smaller parts. In the nested picture, it would be precisely the termination of the nested functions that supplies the gains of the base function. Second, antifragility holds only up to a certain intensity of shocks. **The construction therefore does not eliminate Black Swans; it locates them relative to scale.** An event large enough to end the base function itself would be a Black Swan at that level as well, and beyond that point the structure it supported would cease to exist (Section 7.6.7). Nor does the construction make such events predictable: even in the simulation, the timing of the singularity varied widely from run to run.

Scope. This section describes a conceptual structure. The nesting has not been simulated explicitly across multiple levels (Section 7.6.8); the simulation replaces the nested functions with a simplified noise model. The parallels above are offered as structural correspondences, not as evidence that any particular real-world system is organized in this way.

7.3 The Parameterized Curve as a Measurement Structure: What an Extreme Event Does to It

The parameterized curve of Section 6 is best treated as a **measurement structure**. It is calibrated on a set of tangency nodes within a specific historical stretch, between Day 745 and Day 6388, and its reading, the effective feedback exponent of Section 3.3.1, refers to that stretch. Like any calibrated instrument, it has a **calibration range**, and its reading is only as well determined as its calibration points allow. An extreme scenario makes this concrete, using the tangency rules of Section 6.1: suppose price collapsed to \$0.001 at day 7,000.

Applying the non-crossing and tangency conditions strictly, this scenario produces a deterministic, not arbitrarily chosen, outcome. Testing shows that no coherent tangent curve can simultaneously satisfy exact tangency at all seven previously established historical nodes *and* this new point; any attempt produces errors of several hundred percent and a degenerate, non-physical implied singularity. **Under the model's own rules, this forces the exclusion of the prior nodes as valid tangency points once this new data exists.** What survives as valid anchoring points is only the Genesis origin and this single new node. This holds not only against the seven curated nodes discussed in this paper, but, in principle, against the full historical daily price record, which a disruption of this magnitude would similarly fail to remain consistent with.

The current calibration would not survive such an event. The published parameters ($A = 5.8535 \times 10^{-17}$, $B = 7.4893 \times 10^{-5}$, $n = 5.4666$), and the effective exponent derived from them, would remain a description of the stretch on which they were calibrated. They could not, however, be carried across the event. In measurement terms, the event lies outside the calibration range.

After the event, the instrument has no determined reading. With only the Genesis origin and one new point constraining the system, the exponent n and the curvature parameter B are both free. Because $f(0) = A(\tan(0)/B)^n = 0$ automatically, for *any* choice of A , B and n , the Genesis condition imposes no real constraint at all, and the entire remaining freedom is available to satisfy the single new data point. Given any desired future singularity day x^* and any chosen exponent n , the required parameters can be solved directly:

$$B = \frac{\pi}{2x^*} \quad A = \frac{y_{target}}{(\tan(B \cdot x_{target})/B)^n}$$

For any reader wishing to verify this independently: substituting $x_{target} = 7000$, $y_{target} = 0.001$, and any exponent n (for example, 3, 5, or 10) into this formula, for any chosen future singularity day x^* beyond day 7,000, always yields a valid parameter pair reproducing the required point exactly, with f at Genesis equal to zero by construction. This freedom runs in *both* directions: x^* can be pushed later, but it can equally be pulled earlier than the current calibration implies. Since n is free as well, **the effective feedback exponent is left entirely undetermined.**

The conclusion is therefore a statement about measurement, not about resilience. An extreme disruption ends the validity of the current calibration and leaves the instrument without a determined reading, because a single new node cannot fix both B and n . A new measurement of the effective feedback exponent becomes possible only once enough new tangency nodes have formed after the event to constrain the parameters again. That measurement would then describe the new stretch, not the old one. **A difference between the two readings would itself be informative:** it would indicate that the effective feedback changed across the event. This is consistent with Section 2.1, where fitted constants describe an observed trajectory rather than fixed constants of the process. Conversely, the freedom shown above means that curve-fitting alone cannot exclude the qualitative structure. **This is a limitation of the instrument, not evidence in favor of the structure** (Sections 7.4 and 7.5).

7.4 Domain of the Model: The Boundary at Zero

The governing equation introduced in Section 3.3, $dy/dx = k(1+y^m)$, has $1+y^m > 0$ for any non-negative y and any m , with $k > 0$. This means dy/dx is strictly positive everywhere in the model's domain, for every value of m . **On its own, the equation is structurally incapable of producing a decrease, let alone a decline toward zero.** Consequently, $f(x) > 0$ holds by construction at every point after Genesis. A price of exactly zero is therefore **outside the domain of this model**. The model has nothing to say about it, in the same sense that it has nothing to say about a discrete external event such as an exchange delisting or a network shutdown. For the same reason, a price of zero is not a falsification test in either direction: a genuine test requires a condition the model can meaningfully fail, and a monotonically increasing function offers none in this direction.

Within this domain, the model's reach is wide. As shown in Section 7.3, the Genesis condition $f(0) = 0$ is satisfied automatically for any choice of parameters, so the remaining freedom (A , B , and the exponent, whether fixed at n or generalized to any $m > 1$ in the underlying term) can always match a single new touch point exactly. This does not depend on the price level of that point, and it applies again to every subsequent touch point, however often a collapse and recovery repeats. **No observed price path that stays strictly above zero can therefore exclude every member of this functional family.** This is a statement about the extent of the domain, not support for the model. It also means that price data alone cannot rule the family out, nor establish that $m > 1$ (Section 7.5).

This picture applies to a single function, and it becomes more nuanced for nested functions. In the nested construction of Sections 7.6.6 and 7.6.6.1, the underlying state still never decreases. The observed moving average, however, can fall back onto the function beneath it, and once the outermost base function reaches its own asymptote, it can fall to zero. In that construction, zero is reached not through a gradual decline of the state, but through the end of the structure that supported it.

7.5 The Latent Exponent Is Likely Unidentifiable from Price Data Alone

A further limitation concerns the choice of $m=2$ itself in the latent equation $dy/dx=k(1+y^m)$. Repeating the derivation of Section 3.3.1 for general m shows that, near the singularity, the effective feedback exponent in price terms is $E = 1 + (m-1)/n$ — which reduces to $1+1/n$ only at $m=2$. Any other $m > 1$ paired with a correspondingly rescaled $n = (m-1)/(E-1)$ reproduces the *same* effective exponent, and therefore plausibly the same near-singularity behavior.

This offers a limited defense of the model: the specific value $m=2$ is likely not a substantive, falsifiable commitment, since the empirically observable quantity is the effective exponent, not m directly — choosing $m=2$ for closed-form convenience does not obviously cost anything in fit quality. **But this cuts both ways:** it also means m is likely unidentifiable from price data alone, since any $m > 1$ can plausibly be rescaled via n to match the same observed curve. This equivalence has only been established asymptotically, near the singularity; whether it extends to the full historical trajectory, including the early, near-power-law regime, has not been tested and remains an open question.

7.6 A Possible Simulation Direction: The Exponent as a Stochastic, Regime-Switching Process

This subsection documents, in purely mathematical terms, one possible direction for future simulation work: a generalization of the reflexivity equation in Section 3.3 to a time-varying, stochastically-driven feedback exponent, rather than the fixed value used in the paper's baseline model. No claim of any kind is made here about what the underlying random variable represents, or that this construction describes Bitcoin specifically; the signal is treated strictly as an abstract binary variable, offered as a starting point for further exploration.

7.6.1 The driving signal

Let $S = (S_1, S_2, \dots, S_N)$ be a sequence of independent, identically distributed discrete random variables, each taking one of two values:

$$S_i = \begin{cases} 2 & \text{with probability } p \\ 0 & \text{with probability } 1-p \end{cases}$$

for $i = 1, \dots, N$, with a single asymmetry parameter $p \in (0,1)$. By construction:

$$E[S_i] = 2p \quad \text{Var}[S_i] = 4p(1-p)$$

A single driving signal is sufficient. Only the feedback exponent m enters the differential equation of Section 7.6.3; the mapping exponent n does not. The construction therefore uses one noisy signal, which drives m , and holds n fixed at 1.

7.6.2 The moving average and the driving exponent

Fix a window length W , a positive integer. Define the centered moving average of the sequence S with window W as:

$$M_W[S](i) = \frac{1}{W} \sum_{|j-i| \leq W/2} S_j$$

The two exponents of the model are then set as:

$$m(i) = M_W[S](i) \quad (\text{stochastic, slow-varying})$$

$$n \equiv 1 \quad (\text{deterministic constant})$$

By the law of large numbers, as $W \rightarrow \infty$, $M_W[S](i) \rightarrow 2p$ almost surely. The exponent $m(i)$ therefore fluctuates around the long-run mean $2p$, with $\text{Var}[M_W[S]] \approx 4p(1-p)/W$ in the i.i.d. case; with the holding times introduced in Section 7.6.4, $\text{Var}[M_W[S]] \approx 4p(1-p) \cdot 2H/W$. **The amplitude of the fluctuations is thus governed by the ratio H/W** , that is, by the persistence of the generator relative to the length of the window, and not by either quantity on its own.

7.6.3 The governing equation

Let $y(x)$ solve the non-autonomous ordinary differential equation with piecewise-constant, stochastic coefficient $m(x)$:

$$\frac{dy}{dx} = k(1 + y^{m(x)}), \quad y(0) = y_0 > 0 \text{ small, } k > 0 \text{ constant}$$

and define the observable quantity:

$$f(x) = A \frac{y(x)}{B} \quad (n = 1)$$

for constants $A, B > 0$. **The mapping exponent is held fixed at $n = 1$** , so the observable is a linear rescaling of the state. This is not merely a simplification of convenience: n does not appear in the differential equation at all, and a monotone rescaling of y cannot move the point at which y diverges. **The location of the singularity x^* is therefore determined by $m(x)$ and the rate k , and not by the mapping**; since k is held fixed, every stochastic result reported below originates in that single noisy exponent. Since $m(x)$ is piecewise constant (constant over each unit interval of x , by construction of the moving average on integer-indexed data), the equation for y is solved by numerically integrating forward across each interval with the locally-constant exponent, then updating m at the next interval; this requires no assumption beyond local (sub-interval) numerical stability of the chosen integration scheme.

7.6.4 Observed Result: A Noise-Driven Finite-Time Singularity

The construction in Sections 7.6.1–7.6.3 was simulated with one generalization of the driving signal, anticipated in Section 7.6.2: **the generator holds its current state for a random number of steps before redrawing**, rather than drawing a new, independent value at every step. Holding times are geometrically distributed with mean H . At each switching epoch a new value is drawn with the same probabilities as in Section 7.6.1 (2 with probability p , 0 otherwise). The long-run mean $E[S_i] = 2p$ is therefore unchanged. What changes is how long the moving average can remain on one side of the critical boundary $E[S] = 1$. The i.i.d. construction of Section 7.6.1 is recovered as the special case $H = 1$.

The central test tilts the generator slightly *below* the critical value:

Quantity	Driving signal $m(x)$
Asymmetry parameter p	0.48 (long-run mean $0.96 < 1$)
Window length W (days)	1461
Mean holding time H (days)	40
Standard deviation of the moving average	≈ 0.23
Mapping exponent n	fixed at 1, noise-free

Remaining constants: $k = 10B \approx 7.4893 \times 10^{-4}$, $y_0 = 10^{-6}$, with A and B taken from Section 6.2. With this value of k , the reference configuration (generator tilted upward to $p = 0.6$) produces a finite-time singularity at a median of approximately day 5100, the same order as the singularity reported in earlier runs of this system (day 5778). The equation for y was integrated in logarithmic form, $u = \ln y$. Once the feedback term dominated the constant term, each unit interval was solved in closed form with its locally constant exponent, including the exact location of the blow-up within that interval. All results below are sample statistics over independent random seeds.

Configuration (p_m)	Runs	Runs reaching a singularity	Median x^* (days)	5th–95th percentile of x^*
Reference: 0.60	200	100%	≈ 5100	—

Configuration (p_m)	Runs	Runs reaching a singularity	Median x^* (days)	5th–95th percentile of x^*
Sub-critical: 0.48	200	100%	$\approx 11\,290$	$\approx 5390 - 22\,260$

Every run of the sub-critical configuration ended in a finite-time singularity, even though the exponent's long-run mean (0.96) lies below the critical boundary. The downward tilt did not remove the singularity. It only delayed it, roughly doubling the median time to blow-up relative to the upward-tilted reference.

The role of noise can be isolated directly. If the fluctuations are removed and the exponent is held at its mean, $m(x) \equiv 2p = 0.96$, the equation has no finite-time singularity. For large x it then grows only as a power law:

$$y(x) \approx [(1 - m) \cdot k \cdot x]^{1/(1-m)}, \quad \text{finite for every finite } x$$

The deterministic, noise-free system is therefore stable, while the same system with the same mean exponent but with fluctuations is not. **The singularity is produced by the fluctuations themselves, not by the average tilt.** The following table varies only the speed of the generator, and hence the amplitude and persistence of the fluctuations in the moving average. Each cell gives the median day of the singularity over 30 runs, over a horizon of 1 000 000 days, and every run in every cell reached a singularity.

Mean holding time H (days)	Std. dev. of $m(x)$	$p_m = 0.48$	$p_m = 0.45$	$p_m = 0.40$
10	≈ 0.11	26 591	48 992	171 267
20	≈ 0.16	17 777	25 521	71 959
40	≈ 0.23	11 428	14 608	24 134
80	≈ 0.32	7 952	9 254	12 183

Two regularities are visible:

- For a fixed tilt, a slower generator produces larger and more persistent fluctuations, and the singularity arrives earlier.
- For a fixed generator speed, a stronger downward tilt delays the singularity but, across the tested range down to $p = 0.40$, does not prevent it.

The i.i.d. construction ($H = 1$) is the low-noise limit of the same system. There the moving-average fluctuations are an order of magnitude smaller (std. dev. of $m(x) \approx 0.03$), and at $p_m = 0.48$ only 7 of 20 runs reached a singularity within 400 000 days. Over the full horizon all 20 eventually did, but very late (median $\approx 448\,000$ days). The ordering is therefore monotone:

$$\text{no noise (no singularity)} \rightarrow \text{weak, fast noise (rare and very late singularity)} \rightarrow \text{strong, slow noise (early singularity)}$$

This replaces the earlier reading of the parameter $p^* = 1/2$ as a sharp boundary between a singular and a non-singular regime. In the presence of noise, p^* does not decide *whether* a singularity occurs. It shifts *when* it occurs, together with the amplitude and persistence of the fluctuations. Below p^* , the power-law-like growth described above persists only as a transient phase, whose length depends on how rarely the moving average makes a sufficiently long excursion above the critical boundary.

7.6.5 Mechanism: Why Fluctuations Force the Singularity, and the Link to Antifragility

The mechanism follows directly from the governing equation. During any interval in which the exponent stays constant at a value $m > 1$, the time the trajectory needs to reach infinity, starting from its current level y , is at most:

$$\Delta x^* \leq \frac{y^{1-m}}{k(m-1)}$$

(the $+1$ term in the equation can only shorten this time). Because $1 - m < 0$, this required duration **shrinks toward zero as y grows**. At the same time, as shown in Section 7.6.7, $dy/dx \geq k > 0$ at every point, so y never decreases, whatever regime the exponent is in. The consequences follow in sequence:

- Periods with $m(x) < 1$ slow the growth but never reverse it.
- Periods with $m(x) > 1$ that are too short to cause a blow-up still lift y to a permanently higher level.
- From that higher level, an even shorter excursion above the boundary suffices next time.

The trajectory therefore climbs in a stepwise fashion: long, nearly flat stretches are interrupted by sudden jumps coinciding with excursions of $m(x)$ above 1. Eventually one excursion is long enough relative to the current level, and the system reaches its finite-time singularity. As long as the

moving average has a nonzero probability of staying above 1 for a sustained interval, which holds for any finite window length, this outcome is only a matter of time. This is why the noise-free system (Section 7.6.4) is the only fully stable case.

The same conclusion can be stated as a convexity property. At a fixed level $y > 1$, the growth term $y^m = e^{m \cdot \ln y}$ is a convex function of the exponent:

$$\frac{\partial^2}{\partial m^2} y^m = (\ln y)^2 \cdot y^m > 0$$

Upward fluctuations of m therefore add more to the growth rate than equally large downward fluctuations take away (Jensen's inequality). The time-averaged growth rate of the fluctuating system consequently exceeds that of the noise-free system with the same mean exponent. **Because the curvature grows with $(\ln y)^2$, this asymmetry strengthens as the system itself grows:** the larger the system becomes, the more it gains from each fluctuation.

This is precisely the formal signature of antifragility in Taleb's sense: a convex response to variability, under which the system benefits from volatility, disorder, and shocks rather than merely withstanding them. The simulation thus provides a quantitative counterpart to the conceptual description of antifragility given in Section 7.2. **In this system, noise, that is, shocks, is not an obstacle to the upward dynamic but its driver.** Two features make this explicit:

- The same mean exponent produces a stable, power-law path without noise, yet a finite-time singularity with it.
- Stronger and more persistent fluctuations bring the singularity earlier.

At the same time, the zero boundary of Section 7.6.7 guarantees that downward fluctuations can never push the underlying state back. The system is thus exposed asymmetrically: it keeps the gains from its upward excursions and cannot be eroded by its downward ones. As with the rest of Section 7.6, this is a statement about the mathematical construction, not a claim that the abstract signals represent any specific real-world quantity.

The speed of the generator matters as much as its tilt. The results in Section 7.6.4 show that the holding time, meaning how long the generator remains in a given state, is at least as important as the average asymmetry in determining when the singularity occurs. In a nested, multi-level treatment, the holding time would plausibly grow longer at the slower, outer levels of the hierarchy. By the mechanism above, this would make the slowest levels the most consequential ones, since their excursions are the most persistent. **A qualification from the second part of this work:** in this construction, a slower generator also enlarges the fluctuations of the moving average (Section 7.6.2), so persistence and amplitude change together. When the two are separated, amplitude appears to be the dominant factor, and persistence matters mainly by making sufficiently long excursions possible at all (Part II, Section 3).

The present construction contains a single fluctuating level, and the asymmetry of its generator was fixed exogenously. In a fuller treatment, the asymmetry parameter and holding time of each level would plausibly be weighted functions of the states of the other levels, with every level, at every timescale, exerting some influence on every other. Specifying these weights is left for further work; a conceptual construction is outlined in Section 7.6.6.

7.6.6 Nesting of Functions: Supporting Structures as a Possible Source of Noise

The random generator of Section 7.6.1 is a simplification. This section describes, at a conceptual level, one possible reading of what it stands in for: **a hierarchy of nested functions of the same qualitative kind, in which the smaller, faster structures would supply the fluctuations of the larger, slower ones. Where the fluctuations actually come from is not settled here.** What the simulations show at the level of a single function is a high-frequency, structureless fluctuation whose source cannot be localized, and which behaves, as far as the equation is concerned, much like a random generator. **Terminating nested functions are a plausible contributor to it, not an established one,** and none of the results of Section 7.6.4 depend on the answer.

Construction. Consider a large, long-horizon base function, such as the function $f(x)$ of this paper, calibrated on the tangency nodes of Section 6. Select a point x_0 on this function and treat it as a new origin. From this origin a second function g starts, with local time $\xi = x - x_0$. This function has the same superlinear, asymptotic geometry, but a much faster rise: its own asymptote lies at a local time ξ_g^* far shorter than the remaining distance to the base function's asymptote,

$$\xi_g^* \ll x^* - x_0$$

A faster moving average is taken over this structure. Let $L(x)$ denote the level of the innermost function that exists at x : $L(x) = g(x - x_0)$ while g is defined, and $L(x) = f(x)$ otherwise. The moving average is taken on a logarithmic scale, over a window W_M that is short compared with the lifetime of the base function:

$$\ln M(x) = \frac{1}{W_M} \int_{x-W_M}^x \ln L(s) ds$$

The moving average follows g upward, but **it does not reach the asymptote of g** . The logarithm of a function diverging at a finite point remains integrable, so $M(x)$ stays finite as g diverges. The moving average therefore detaches from g as g approaches its asymptote. Beyond the asymptote, at $\xi > \xi_g^*$, the smaller function no longer exists. The window of the moving average fills up with values of the function beneath, and $M(x)$ falls back onto the base function $f(x)$. The base function therefore acts as a **supporting structure** for the moving average: the nested excursions rise above it and return to it, but do not break through it. **This support lasts only as long as the base function itself exists.** The base function has the same asymptotic geometry as the structures nested in it, so it also terminates at its own asymptote x^* . What happens once the outermost base function itself ceases to exist is taken up in Section 7.6.7. Writing each such excursion as a non-negative log-deviation D_j from the base, starting at its own origin x_j , the moving average can be written schematically as:

$$\ln M(x) \approx \ln f(x) + \sum_j D_j(x - x_j), \quad D_j \geq 0$$

Here D_j is the windowed average of $\ln(g_j/f)$ over the part of the window in which g_j exists. The approximation reflects that $\ln f$ itself changes slowly across the short window. The profile D_j is zero before its origin. It then rises along the accelerating path of its own function g_j , detaches as g_j reaches its asymptote, and decays back to zero once g_j has ceased to exist.

The construction is recursive: each g_j can itself serve as the base function for still smaller and faster nested structures, which in turn would supply its fluctuations. Read this way, **the sum of the nested excursions below a given level would be that level's noise**, and the generator $m(x)$ of Section 7.6.1 a statistical stand-in for that aggregate. Its holding time (Section 7.6.4) corresponds to how long a nested structure keeps accelerating before the moving average detaches from it.

The validity of the parameters comes to the fore again. In this picture, the parameters of the base function (Section 6.2) are not fixed constants of a single process. They are effective averages over everything nested beneath it. The parameters of the nested structures themselves (origin, speed, lifetime, amplitude) are unstable and vary from one excursion to the next. This is the same fragility of parameters emphasized throughout Section 7, now appearing as a structural feature rather than as a limitation of the fit.

Instability of the parameters alone is sufficient for an asymptote. The single-level simulation of Section 7.6.4 supports this directly. Setting the generator exactly to the critical value, $p_m = 0.5$, gives the mean exponent $E[m] = 1$. All other settings were as in Section 7.6.4.

Exponent with mean $m = 1$	Runs	Runs reaching a singularity	Median x^* (days)	5th–95th percentile of x^*
No noise ($m \equiv 1$)	—	none (pure exponential growth)	—	—
Fast, weak noise (i.i.d., $H = 1$)	20	100%	$\approx 70\,420$	—
Slower, stronger noise ($H = 40$)	200	100%	$\approx 8\,850$	$\approx 4\,120 - 18\,910$

Without fluctuations, the system with mean exponent 1 has no finite-time singularity: the equation reduces to $dy/dx = k(1 + y)$, whose solution grows exponentially and remains finite for every finite x . With a fluctuating exponent of exactly the same mean, every run reached an asymptote, and the more persistent the fluctuations, the sooner it arrived. **A function whose parameters are unstable, but whose average feedback exponent is exactly $m = 1$, can therefore still terminate in an asymptote.** In the nested picture, the fluctuating structures beneath it are what make this happen.

Detachment and accumulation. In the nested picture, the moving average never completes the asymptotes of the nested levels: each nested function reaches its own asymptote and ceases to exist, while the moving average detaches and falls back onto the level beneath (Section 7.6.6.1). In the terms of Section 7.6.5, this is consistent with how the asymptote of the level that contains them is nevertheless reached. Excursions of the exponent above the critical boundary that are too short to cause a blow-up do not vanish without trace: they leave the underlying state y permanently higher. The moving average returns to its supporting base, but the excursion's contribution to the underlying state is not erased (Section 7.6.7). Asymptotes that the moving average does not complete at the nested levels thus accumulate into the completed asymptote of the level that contains them. This is a further expression of the antifragile structure described in Section 7.6.5: on this reading, the shocks arriving from the nested levels are what drive the larger structure forward.

As stated in Section 7.6.8, this nesting has not been simulated explicitly across multiple levels here. The random generator remains a simplified substitute for it. The construction above is offered as the interpretation that the generator is meant to approximate, and as a concrete starting point for a fully nested simulation.

7.6.6.1 A simple numerical illustration

The following script, run in JupyterLab on a Fedora Workstation system using only NumPy, pandas and Matplotlib, reproduces the two effects described in this section on a single level. It is not the simulation of Section 7.6.4: **the parameters differ, and the exponent here is drawn independently on each step rather than taken from a moving average of a driving signal**, so the fluctuations carry no persistence at all. The exponent alternates between 0.8 and 1.2, with the higher value drawn with probability 0.45, **giving a mean exponent of 0.98, below the critical value of 1**. The state nonetheless reaches a finite-time singularity, which illustrates the central point of Section 7.6: the outcome is decided by the excursions above the boundary, not by the average. It also shows why persistence is not indispensable here. As the state grows, the superlinear stretch required to complete the divergence becomes shorter, until even the brief runs produced by independent draws are long enough. The second half of the script forms the observed quantity and its logarithmic moving average, and **plots the moment at which the function ceases to exist and the moving average detaches from it**, falling back toward the base level drawn beneath it.

```
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt

# --- 1. PARAMETERS ---
A = 10e-17
B = 10e-5
k = 10 * B
y0 = 1e-6
windowM = 1000      # Moving average window size in days

# Noise configuration
p = 0.45
numDays = 20000

# --- 2. NOISE GENERATION ---
m = np.random.choice([1.2, 0.8], size=numDays, p=[p, 1-p])

# --- 3. SOLVING THE DIFFERENTIAL EQUATION ---
y = np.zeros(numDays)
y[0] = y0

blowupDay = numDays # Store the day when the blow-up occurs

for i in range(1, numDays):
    dy = k * (1 + y[i-1]**m[i-1])
    y[i] = y[i-1] + dy

    # Safety break: singularity (blow-up)
    if y[i] > 1e30:
        blowupDay = i
        # The latent state drops back to the base level (y0) to compute the moving average decay.
        y[i:] = y0
        print(f"Singularity (blow-up) reached on day {i}! The function ceased to exist.")
        break

# --- 4. OBSERVABLE QUANTITY AND MOVING AVERAGE ---
f = A * (y / B)
baseLevel = A * (y0 / B) # The starting height (supporting structure)

# Calculating the logarithmic moving average on the latent state
logF = np.log(f)
logM = pd.Series(logF).rolling(window=windowM, min_periods=1).mean()
M = np.exp(logM)

# For plotting, the original function genuinely disappears after the singularity
fVisible = f.copy()
fVisible[blowupDay:] = np.nan

# --- 5. PLOTTING ---
lastVisibleDay = min(blowupDay + windowM + 500, numDays)

plt.figure(figsize=(40, 20))

# Plotting the horizontal base level
plt.axhline(baseLevel, color='gray', linestyle=':', linewidth=2, label='Base level')

# Plotting f(x) (which disappears) and the moving average
```

```
plt.plot(fVisible[:lastVisibleDay], color='lime', alpha=0.5,linewidth=1, label='Stochastic blow-up trajectory')
plt.plot(M[:lastVisibleDay], color='orange', linewidth=2, label=f'Logarithmic Moving Average')

plt.yscale('log')
plt.title('Stochastic blow-up trajectory and detachment of the logarithmic moving average')
plt.xlabel('x')
plt.ylabel('Value (Log scale)')
plt.legend(loc='upper left')
plt.grid(True, alpha=0.3)

plt.tight_layout()
plt.show()
```

Figure 2.

Three points should be kept in mind when reading the output. First, the script shows the supported case, not the terminal one: after the singularity the function is removed from the plot, and the moving average decays over the length of its window toward the base level that stands beneath it. **The case in which no such level remains, and the moving average falls to zero, is the separate matter treated in Sections 7.6.7 and 7.4** Second, the integration is a fixed-step explicit scheme with a step of one day, so **the exact day of the singularity is scheme-dependent and only the qualitative behavior is meaningful.** Third, the random draws are not seeded, so individual runs differ; fixing a seed makes a particular run reproducible.

7.6.6.2 Implications for shorter horizons and for the parameterized model

This simulation suggests that, because of the instabilities described above, the approach of Sections 6 and 7.3, which reads the structure through a curve calibrated within a given time-prefix, is likely to be **considerably less accurate over shorter "trends" than over longer cycles.** At shorter horizons, the nested excursions and their unstable parameters make up a larger share of what is observed, while over longer cycles they tend to average out into the slower supporting structure. The simulation also raises a further possibility, stated here tentatively. The structural pattern examined earlier in this paper was observed through a model with **specific, fitted parameters.** The simulation indicates that a feedback exponent averaging $m = 1$ (or somewhat below 1; Section 2.1), with the mapping exponent held at $n = 1$ and combined with **a sufficient degree of noise,** may produce **structural properties similar to those of the parameterized model,** including termination in a finite-time asymptote. If this holds, the particular parameter values of Section 6.2 may be better read as one realization of a more general noise-driven structure than as its defining constants. This remains an interpretation of a deliberately simplified simulation (Section 7.6.8), not an established result.

7.6.7 A Note on the Zero Boundary

For a single function, this stochastic extension does not alter the zero-boundary property already established for the fixed-exponent case (Section 7.4): since $1 + y^{m(x)} \geq 1 > 0$ for every finite value $m(x)$ can take — including the (rare, but not impossible) case where the moving average of the driving signal, and with it the exponent, equals exactly zero — dy/dx remains strictly positive at every point along the trajectory, regardless of which regime $m(x)$ is currently in. **A single function therefore still cannot decrease on its own, let alone reach zero, under this generalization,** for the same structural reason as in the fixed-exponent case, not because of any special property of the moving-average construction itself. Within a single function, a price of zero remains outside the domain of the model.

This no longer holds for the nested system as a whole. In the nested construction of Section 7.6.6, the observed quantity is not the state of one function but the logarithmic moving average $M(x)$, which follows the innermost function that currently exists. When a nested function completes its singularity, it ceases to exist, and the moving average detaches from it and falls back onto the supporting function beneath. **The observed level can therefore decline, even though no single function ever gives back any of its own growth.** Such declines are bounded from below only as long as a supporting function exists. The outermost base function has no support beneath it: **if it, too, ends in a singularity and the moving average detaches from it, the moving average falls to zero and the process as a whole stops.** Zero is therefore reachable in the nested system, but not through a gradual decline of any state variable; it is reached through the end of the last structure that supported the observed level. **The zero boundary of this section is a property of each function taken individually, not of the observed dynamics of the nested system.**

7.6.8 A Necessary Caveat on Simplification

While this simulation points in a direction that appears credible, it is only a very simplified model. It contains a single fluctuating level, a simplification made purely because of resource constraints, and it therefore does not handle the full set of nested timescales that a complete treatment of this idea would require.

7.6.9 Possible research or simulation directions

Possible directions include fractal geometry, biology, psychology, economic cycles, technological development, the hierarchical nesting of solar cycles, phantom dark energy cosmology (the "Big Rip" future singularity), and other systems exhibiting similar properties. A structurally related phenomenon appears in magnetic confinement fusion: in nonlinear ideal ballooning theory, the explosive onset of **tokamak edge instabilities** (edge localized modes), a key stability concern for next-generation devices such as ITER, is attributed to a finite-time singularity in the governing model (Wilson & Cowley, *Phys. Rev. Lett.* 92, 175006, 2004). There, a finite perturbation can trigger the transition even where the configuration is stable to infinitesimal ones. The parallel offered here is purely structural: superlinear feedback combined with fluctuations near a critical boundary. It is not a claim of shared physical mechanism.

This direction is noted at greater length than the others for a specific reason. Its connection to Bitcoin is indirect. It is included because the simulation of this section, together with the nested construction of Section 7.6.6, is not specific to any asset: it describes a general class of systems with superlinear feedback, fluctuating parameters and nested timescales. The discussion below is offered as an illustration of how far such a construction might reach, not as a contribution to plasma physics.

Where a downward path appears. The construction contains two levels that behave differently. The underlying state y has no downward path: by Section 7.6.7, $dy/dx > 0$ at every point. The observed level, however, does have one. In the nested picture of Section 7.6.6.1, the logarithmic moving average $M(x)$ follows a nested function g_j upward and detaches from it as g_j approaches its own asymptote. Once g_j has ceased to exist, **the moving average falls back onto the function beneath it**. Each excursion D_j therefore rises, peaks and decays, and the sequence of such excursions forms **a recurring build-up-and-crash cycle at the observed level**. The lower-level singularity is, in this sense, **regularized as seen from above**: the nested function diverges, but the moving average stays finite, because the logarithm of a function diverging at a finite point is integrable.

The lower bound of the moving average is zero, not the outermost function. The supporting role of each function holds only as long as that function itself exists. The outermost base function, the outermost level on which all others rest, is of the same qualitative kind as the ones nested within it: it also terminates in a finite-time asymptote. When it does, there is no further structure beneath it. As soon as the window of the moving average reaches beyond this point, **the moving average falls to zero: the process can stop**. This does not contradict Section 7.6.7. The underlying state never decreases; the path to zero does not lead through a decline of y , but through its completed singularity, beyond which the function no longer exists. The accumulation described in Section 7.6.6 therefore has a definite end point. Uncompleted asymptotes at the nested levels accumulate into the completed asymptote of the outermost level, and that completion is the termination of the whole structure.

Read this way, the construction contains **two kinds of downward movement**. The first is a partial crash, after which the moving average returns to a still-existing supporting function and the cycle continues. The second is a terminal crash, after which it falls to zero. The plasma case shows a structurally similar distinction. An edge localized mode only partially relaxes the edge profile, which then builds up again. **A disruption**, the abrupt loss of confinement that terminates the discharge, corresponds more closely to the terminal case.

The differences are equally important. In the plasma, the regularization of the partial crashes is a physical process: saturation, relaxation, and the loss of energy and particles from the edge. In the construction, it is a property of the averaging operator. The nested function does complete its asymptote; only the observed average fails to follow it. In addition, **a partial crash in the construction does not erase what the excursion contributed to the underlying state** (Sections 7.6.6 and 7.6.7), and it is this accumulation that leads toward the terminal crash. It is an open question whether repeated edge instabilities leave any comparable cumulative trace in the plasma that brings a terminal event closer, or whether the parallel is confined to the separate shapes of the two kinds of crash. It is equally open whether noise persistence plays a role in the timing of such events comparable to its role in Section 7.6.4.

A possible implication, should the structural parallel prove to be more than an analogy. If a genuine connection exists between the dynamics of tokamak plasmas and the system dynamics described in this section, the results of Section 7.6.4, read through the nested picture, point toward a possible stabilizing strategy along two lines. In the nested picture, the holding time of the generator corresponds to how long a nested structure keeps accelerating before the moving average detaches from it. In the simulation, **a faster generator, meaning a shorter holding time, produced smaller and less persistent fluctuations, and the singularity arrived considerably later**. In nested terms, the excursions are cut short before they can build up to a large amplitude, and the completion of the outermost level, and with it the terminal crash, is postponed.

The first physical counterpart would be to **accelerate the plasma flow, in particular its sheared component**, so that fluctuations decorrelate more quickly. A related principle is already established in plasma physics on independent grounds: sheared flow is known to suppress edge turbulence by speeding up turbulent decorrelation (Biglari, Diamond & Terry, *Phys. Fluids B* 2, 1, 1990). The second counterpart concerns the partial crashes themselves: **triggering them more often than they would occur on their own, so that each has less time to build up**. This, too, has an established experimental counterpart. In ELM pacing by pellet injection, raising the ELM frequency above its intrinsic value was found to make the individual events milder (Lang et al., *Nucl. Fusion* 44, 665, 2004).

Within the present construction, however, **shorter excursions only delayed the singularity and made it rarer; they did not remove it entirely**. Full stability was reached only in the noise-free limit. The nested picture also indicates why: an excursion that is cut short still leaves its contribution in the underlying state, and these contributions continue to accumulate toward the completion of the outermost level. Should the parallel hold, such measures would therefore be expected to reduce the size of individual partial crashes, and to postpone a terminal one, more readily than to exclude it. Whether this correspondence reflects a shared mechanism, or only a structural resemblance, is left open here.

8. CONCLUSIONS

This paper is, in the end, best understood not as a price-prediction tool, but as an attempt to build a measurement structure, which makes it both narrower and broader than a curve fit. It is narrower in the sense repeatedly emphasized throughout: **the specific numerical parameters presented here describe one calibrated historical stretch, can be invalidated by a sufficiently extreme market event, and carry no claim to physical or economic necessity.** As Section 2.1 shows, even the qualitative structure they describe does not depend on these particular values.

It is broader in a different sense. The recurring structure of this paper's argument — first excluding Bitcoin's own supply contraction as a sufficient driver (Section 3.1), then separately excluding ordinary demand growth on the same grounds (Section 3.2) — points toward a more general framework than a single asset's price curve. What this exercise amounts to is the use of an asset with a hard, independently verifiable supply constraint as a fixed reference point, against which the behavior of everything on the other side of the price ratio can be examined. Read this way, **the contribution of this paper is not only a descriptive curve fit to one asset's history, but a candidate measurement instrument:** a fixed point against which the internal dynamics of an external, otherwise difficult-to-observe system might be studied indirectly, through the shape of the divergence it produces.

A direct corollary follows from this framing, worth stating explicitly. **If the source of the feedback loop this paper's mathematics requires is, by the exclusions in Sections 3.1 and 3.2, external to Bitcoin's own supply dynamics, then no amount of further analysis confined to Bitcoin's own price and supply data could ever supply a positive derivation of that mechanism** — only further exclusions of additional internal candidates, of the same kind already carried out. Confirming what the external process actually is would require studying that external system directly, on its own terms, which lies outside this paper's stated scope. **This is not a gap this paper could close by being more thorough within its own boundaries; it is a structural consequence of where the argument itself locates the missing term.**

This corollary is not merely theoretical. Exploratory attempts, in the course of preparing this paper, to detect a signature of this feedback mechanism directly within Bitcoin's own price statistics — momentum-style dependence between trend-relative price level and subsequent returns, and scale-invariance in the distribution of price fluctuations across different time horizons — did not yield supporting evidence, and in some cases pointed toward the opposite, ordinary market behavior. **This is consistent with, though it does not by itself prove, the structural conclusion above: if the mechanism is genuinely external, examining the instrument — or its wiring, Bitcoin's own supply and demand dynamics — should not be expected to illuminate it, and the negative results obtained here are compatible with exactly that.** This should not be read as a claim that no analysis of Bitcoin's own data could ever bear on the question; only that the specific attempts made here did not.

This paper does not establish this broader framing conclusively, and **it should be read as the direction the preceding sections point toward rather than a claim defended in full here.** It is offered as the most honest summary of what this body of work, taken as a whole, appears to be reaching for.

9. PLAIN-LANGUAGE SUMMARY

This paper compares a mathematical curve based on the tangent function, which contains a self-reinforcing feedback term, with a simpler alternative, the Power Law, on Bitcoin's historical price lows, in an attempt to detect whether such feedback is present. The tangent-based curve fits these lows more closely than the best possible Power Law fit on the same data, and the paper uses this fit to estimate an effective feedback exponent for that historical stretch. A simplified simulation also shows that the same overall shape can emerge from fluctuating, noise-driven exponents, without any finely tuned values. The curve has a mathematical feature, a point where its value becomes undefined, which is interpreted as a theoretical marker of possible extreme illiquidity, not as a price prediction. This is exploratory, descriptive work, not a forecasting tool, and it carries real uncertainties, discussed openly in Section 7. **Read together with the exclusions in Sections 3.1 and 3.2, this points toward a broader, more tentative possibility: that the model's greatest value may lie less in the price curve itself than in its potential use as an instrument for observing the structural pressures of the external system it points to** — a direction this paper proposes rather than conclusively establishes.

REFERENCES

- Santostasi, G. & Perrenod, S. (2026). *A mechanistic derivation of the Bitcoin price power law: Network adoption dynamics and generalised Metcalfe scaling*. Nonlinear Science (Elsevier), 8, 100172. DOI: 10.1016/j.nls.2026.100172.
- Nakamoto, S. (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. <https://bitcoin.org/bitcoin.pdf>
- Soros, G. (1987). *The Alchemy of Finance*. Simon & Schuster, New York.
- Soros, G. (2013). *Fallibility, reflexivity, and the human uncertainty principle*. Journal of Economic Methodology, 20(4), 309-329. DOI: 10.1080/1350178X.2013.859415.
- Johansen, A. & Sornette, D. (2001). *Finite-time singularity in the dynamics of the world population, economic and financial indices*. Physica A, 294(3-4), 465-502. DOI: 10.1016/S0378-4371(01)00105-4.
- Sornette, D. (2003). *Why Stock Markets Crash*. Princeton University Press, Princeton, NJ.

- Mandelbrot, B. B. & Wallis, J. R. (1968). *Noah, Joseph, and operational hydrology*. Water Resources Research, 4(5), 909–918. DOI: 10.1029/WR004i005p00909.
- Mandelbrot, B. B. (1982). *The Fractal Geometry of Nature*. W. H. Freeman, San Francisco.
- Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable*. Random House, New York.
- Taleb, N. N. (2012). *Antifragile: Things That Gain from Disorder*. Random House, New York.

The Geometry of Bitcoin:

Intermittent Superlinearity — A Blow-Up Theorem for Fluctuating Exponents

1. SCOPE AND CLAIMS

The earlier part of this work arrived, by way of a single asset's price geometry, at a construction that turned out not to depend on that asset at all. A state variable grows under a feedback term whose exponent is not fixed but fluctuates, and the system reaches a finite-time singularity even when the average feedback is not superlinear. That observation began as a numerical result. **It is established as a theorem in Section 3 of this paper**, and the remaining question, which is genuinely open, is whether anything in the physical world realizes the structure.

The object of study. Consider a non-negative state variable $y(x)$ governed by

$$\frac{dy}{dx} = k(1 + y^{m(x)}), \quad k > 0$$

where **the feedback exponent $m(x)$ is not a constant but a fluctuating quantity whose values straddle the critical boundary $m = 1$** . For a constant exponent the behavior is classical: a finite-time singularity exists if and only if $m > 1$. Section 2 sets out that case exactly, and Section 3 treats the fluctuating one.

A single instance, not the whole system. Everything stated in Sections 2 and 3 concerns **one instance of this equation**. For such an instance the growth rate never falls below $k > 0$, so **its state cannot decrease**: subcritical stretches slow it, but nothing reverses it (Lemma 1). **This monotonicity is a property of the single instance, not of the full dynamics**. In the construction this paper is ultimately concerned with, instances of the same equation are **nested within one another**: faster instances ride on slower ones, and the observed quantity follows the innermost instance that currently exists. When a nested instance completes its singularity, it ceases to exist, and the observable **detaches from it and falls back onto the supporting instance beneath**. Seen at the level of the whole system, **the observable can therefore decline**, even though no single instance ever gives back any of its own growth. Such declines are bounded from below for as long as a supporting instance exists beneath them. **The outermost instance has no support beneath it**. If it, too, runs to its own singularity, there is no lower level for the observable to return to, and **the nested process as a whole terminates**. In the application that motivated this work, this corresponds to a breakdown of the logarithmic moving average itself, rather than a return to it (Part I, Section 7.6.6). None of this conflicts with the results below. The theorem of Section 3 requires only the monotone state of a single instance, and the declines described here occur at a different level, between instances rather than within one. **They are taken up once the nested construction is developed, after Section 4.**

Three claims of different status. They are kept apart deliberately, because they are not equally secure. **First, a proved result** (Section 3): the fluctuating case is not a perturbation of the constant case but a distinct regime, in which intermittent excursions above the boundary force the singularity almost surely, and in which the mean exponent plays no part in the outcome. **Second, a conjecture about the origin of the fluctuations** (Section 4): that they need not be imposed from outside, but may be produced by smaller, faster instances of the same equation nested within the system. **Third, a conjecture about the physical world** (Section 4): that the class defined here is not empty, that some real systems realize this structure. The third is the reason for the paper and the least settled part of it. **A demonstration that no such system exists would be an equally valid outcome.**

What is claimed, and what is not. This paper will use examples drawn from unrelated physical domains: edge instabilities in confined plasmas, nonlinear optical processes, accelerating creep in geomechanics, and noise-contaminated control loops that occasionally run away. **No claim is made that these share a physical or causal mechanism**, and nothing here belongs to solid-state physics, plasma physics or any unified physical theory. The claim is **mathematical isomorphism**: that the governing dynamics of these systems, once reduced, share the same structure, and that their behavior near the critical boundary is therefore governed by the same geometry of phase space.

This kind of claim is ordinary in dynamical systems theory. A pendulum and an LC circuit have no physical mechanism in common, yet the same second-order equation describes both, and consequences derived from that equation apply to both. **The shared object is the equation, not the hardware**. In the same sense, the systems examined here are treated as different physical realizations of one mathematical structure, each computing the same dynamics in its own substrate.

A note on terminology. The term "universality class" is used in this paper in a deliberately restricted sense: a set of systems sharing the same blow-up scaling exponent $1/(m-1)$ and the same qualitative response to fluctuations in m . **It is not used in the strict sense of critical phenomena**, where membership is established through renormalization group arguments and a full set of critical exponents. No such derivation is offered here, and none is claimed.

Order of argument. The mathematical case is developed first, and the physical examples are taken up only afterwards. This order is deliberate: the examples are tests of a stated structure, not the evidence from which the structure is inferred. Section 2 establishes the deterministic reference case, Section 3 proves the theorem and the criterion that follows from it, and Section 4 states what remains conjectural and how a candidate system is to be tested. **Where a candidate fails the test, this is reported as such.**

2. THE DETERMINISTIC SINGLE-FUNCTION CASE

Before any fluctuation is introduced, it is worth setting out precisely what can be said about **a single instance of this structure with a constant exponent**. Everything in this subsection is exact, elementary, and independent of the conjectures of Section 4. It serves as the reference case against which the fluctuating construction of Section 3 is compared.

The system. The structure consists of a **latent equation** governing an internal, dimensionless state variable, together with an **observation map** carrying that state into whatever quantity a given system actually exposes to measurement. Throughout this paper the observation map is taken to be linear, so that the observed quantity is the state itself expressed in the system's own units:

$$\frac{dy}{dx} = k(1 + y^m), \quad y(0) = 0 \qquad f(x) = A \cdot y(x) / k$$

with $k > 0$, $A > 0$ and $m > 0$ constant. Here k carries the dimension of inverse x , A carries the unit of the observed quantity, whatever that happens to be, and m is dimensionless. Dividing by k is a normalization choice, made so that the observed quantity reduces to $A \cdot x$ in the early phase. A more general, monotone power map is discussed at the end of this section; **it changes none of the results below, and is set aside because the feedback mechanism itself does not depend on it.**

Proposition 0 (the rate is a clock, not a shape). Substituting $\xi = kx$ gives $d y/d\xi = 1 + y^m$, in which k no longer appears. The solution of the general system is therefore $y(x) = \tilde{y}(kx)$, where $\tilde{y}$ is the single universal solution of the reduced equation. Consequently **k is a pure rescaling of the horizontal axis and A a pure rescaling of the vertical axis**: neither alters the shape of the trajectory. **With the observation map fixed, the family has exactly one shape parameter, m .**

Proposition 1 (existence and location of the singularity). Separating variables gives the exact distance to the divergence, measured from the origin:

$$x^* = \frac{1}{k} \int_0^\infty \frac{du}{1 + u^m} = \frac{1}{k} \cdot \frac{\pi / m}{\sin(\pi / m)}$$

The integral converges if and only if $m > 1$, which is the **blow-up criterion**: for $m \leq 1$ the trajectory is defined for all x and no singularity exists. Writing $G(m) = (\pi/m)/\sin(\pi/m)$, the result separates cleanly into $x^* = G(m)/k$: the exponent fixes the *dimensionless* distance to the singularity, and the rate fixes the clock. The closed form used in the earlier part of this work is the special case $m = 2$, where $G(2) = \pi/2$ and $x^* = \pi/(2k)$.

$G(m)$ decreases monotonically from infinity to one as m runs from 1 to infinity, so **the singularity can be pushed arbitrarily far out by approaching the critical boundary from above, but never closer than $1/k$, however strong the feedback**:

m	1.2	1.5	2	3	5	10
$k \cdot x^* = G(m)$	5.236	2.418	1.571	1.209	1.069	1.017

Proposition 2 (local scaling law). Near the singularity the constant term becomes negligible and $dy/dx \approx k y^m$, which integrates to

$$y(x) \approx [(m-1) k (x^* - x)]^{-1/(m-1)}$$

The blow-up exponent $1/(m-1)$ depends on m alone; k enters only through the prefactor. This exponent is the quantity by which two systems can be compared without reference to their units or substrates, and it is the candidate invariant on which the classification proposed in this paper rests.

Proposition 3 (remaining distance from a given level). The same integral taken from an arbitrary level y rather than from the origin gives what is still required to reach the singularity:

$$T(y) = \frac{1}{k} \int_y^\infty \frac{du}{1+u^m} \approx \frac{y^{1-m}}{k(m-1)} \quad (y \gg 1)$$

Because $1 - m < 0$, this quantity **shrinks toward zero as the state grows**, and it is inversely proportional to k . This is the result from the deterministic case that carries the most weight for what follows: **the higher the rate k and the higher the level already reached, the shorter the interval of superlinear feedback needed to complete a singularity**. In the constant-exponent case treated here this is merely a restatement of Proposition 1. Its significance appears only once m is allowed to vary, since it then sets **how long an excursion above the critical boundary must last in order to be sufficient**, and that requirement weakens as the state grows.

Proposition 4 (the subcritical and critical cases). For completeness, since a fluctuating exponent spends part of its time at or below the boundary: for $m < 1$ the same separation of variables gives $y(x) \approx [(1-m)kx]^{1/(1-m)}$ at large x , **a power law with no singularity**, while $m = 1$ gives exactly $y(x) = e^{kx} - 1$, which is unbounded but finite at every point. **The three regimes are therefore power-law growth, exponential growth, and finite-time divergence**, separated by the single value $m = 1$.

The early phase. While $y \leq 1$, the feedback term satisfies $0 \leq y^m \leq 1$ for every $m \geq 0$, so the growth rate stays between k and $2k$ and the state grows at a nearly constant rate. **In this phase the trajectory carries almost no information about m :** systems that will diverge and systems that never will are, at this stage, indistinguishable from one another and from a feedback-free linear growth. The exponent becomes visible only once the state is large enough for the term y^m to dominate.

Roles of the parameters. The above fixes each one exactly. **k is the internal clock:** it sets how fast the system runs and nothing else, scaling the whole trajectory horizontally, so that x^* is inversely proportional to it. **m is the feedback strength:** it determines whether a singularity exists at all, its dimensionless distance $G(m)$, and the blow-up exponent $1/(m-1)$. Together, **k and m fix the singularity: m says how far in dimensionless terms, k says how long that is.** **A is the unit of measurement:** it converts the dimensionless state into the quantity a particular system exposes — a currency, a current, an intensity — and affects neither the shape of the trajectory nor the location of the singularity.

Remark: a non-linear observation map. Some systems expose not the state but a power of it, $f = A(y/k)^n$ with $n > 0$ constant. Since this map is monotone, **it can neither create nor remove a singularity**, and x^* is unchanged; it only reshapes the approach. Eliminating the state gives a closed equation in the observed quantity,

$$\frac{df}{dx} = n k A^{1/n} f^{1-1/n} + n k^{m+1} A^{(1-m)/n} f^E, \quad E = 1 + \frac{m-1}{n}$$

in which the first term alone is a plain power law and the second is the feedback term, with E the effective feedback exponent as it appears in the observed quantity. **$E > 1$ precisely when $m > 1$** , so the blow-up criterion survives the transformation unchanged, as it must. Two points are worth keeping from this. First, **only E is accessible from an observed trajectory, and E constrains just the combination $(m-1)/n$:** any pair with the same ratio produces the same observable, so the feedback exponent cannot be recovered from a record of the output unless the observation map is known independently. Second, with the linear map adopted here, $n = 1$ and **$E = m$:** the observed feedback exponent and the exponent in the equation coincide, which is why the general map adds nothing to the argument that follows.

2.1 A worked instance, and the origin of this construction

The simplest non-trivial case is $m = 2$, $k = 1$, $A = 1$ and $n = 1/2$, for which the latent state is $y = \tan(x)$ and the observed quantity is

$$f(x) = \sqrt{\tan(x)}$$

Differentiating directly gives $df/dx = \frac{1}{2} \tan(x)^{-1/2} \sec^2(x) = \frac{1}{2}(f^{-1} + f^3)$, which is exactly what the general expression above returns: the first exponent is $1 - 1/n = -1$, the second is $E = 1 + (m-1)/n = 3$, and both coefficients are $n k = \frac{1}{2}$. **Two points are worth drawing from this.** First, $\sqrt{\tan(x)}$ is *not* itself a member of the family $dy/dx = k(1 + y^m)$ for any m , since it produces two terms of different exponent; it is that family seen through a power map. **The same dynamic therefore presents as a different equation depending on the power through which it is read**, which is the concrete form of the identifiability limitation noted above: here a latent quadratic feedback appears in the observed quantity as a cubic one. Second, this particular function is where the present line of work began, before the latent equation was identified; it is recorded here so that the relationship between the two is explicit.

3. INTERMITTENT SUPERLINEARITY: A BLOW-UP THEOREM

This section states and proves the central mathematical claim of the paper. The structure of the argument is a race between two quantities established in Section 2: **the duration of superlinear feedback still required to complete a singularity, which shrinks as the state grows, and**

the probability of an excursion of that duration, which does not. Once the first falls below a fixed window length, the second stops shrinking altogether, and the outcome follows.

Setting. Let y satisfy

$$\frac{dy}{dx} = k(1 + y^{m(x)}), \quad y(0) = y_0 \geq 0, \quad k > 0$$

where the exponent path $m(x)$ is a stochastic process satisfying:

(S1) Boundedness and regularity. $0 \leq m(x) \leq M < \infty$, with right-continuous sample paths. **(S2) A superlinear tail.** There exist $\mu > 1$ and $\Delta > 0$ such that the probability of a sustained excursion, $q = P(m(x) \geq \mu \text{ throughout a window of length } \Delta)$, is strictly positive. **(S3) Renewal.** Sufficiently separated windows are independent, or asymptotically so; a regenerative or finite-state Markov exponent process satisfies this. **This is an assumption about the exponent process, not a consequence of the construction**, and it is the condition the nested case of Section 4 is least likely to satisfy: where the fluctuations of one level are produced by the levels within it, intermediate levels transmit influence between windows that are far apart, so independence cannot be assumed at any separation. Section 5 adds a second reason to be careful with it, since a bounded system has no unlimited supply of fresh trials. **Section 4 leaves open a question that bears directly on this point:** whether a completed instance, which is brief and enormous when measured on a slower level, can itself act as a fluctuation there. Were that so, the renewal a level cannot generate within itself would be supplied from below. **That is a question and not a result**, and nothing in the theorem depends on it; the theorem applies where (S3) holds, and says nothing where it does not.

Assumption (S2) is weak: if the paths are right-continuous and the exponent exceeds μ with positive probability at all, then it stays above μ for some positive duration, and some window length Δ therefore has $q > 0$. **Note what (S2) does not mention: the mean of m , which plays no role anywhere in what follows.**

Lemma 1 (ratchet). Since $dy/dx \geq k > 0$, the state is strictly increasing and $y(x) \geq y_0 + kx$. **Any level Y is therefore reached no later than $x = Y/k$, and is never given back.** Subcritical stretches slow the growth; they cannot reverse it. **This holds within a single instance; declines of the observable in a nested system arise between instances, not within one (Section 1).**

Lemma 2 (excursion criterion). Suppose $y(a) = Y \geq 1$ and $m(x) \geq \mu > 1$ throughout $[a, a+\tau]$. For $y \geq 1$ and $m \geq \mu$ we have $1 + y^m \geq y^\mu$, so comparison with $dz/dx = k z^\mu$, $z(a) = Y$, gives $y \geq z$ on the interval. Since z diverges at $a + Y^{1-\mu}/[k(\mu-1)]$, **the singularity is completed within the excursion whenever**

$$\tau > \tau^*(Y) = \frac{Y^{1-\mu}}{k(\mu-1)}$$

Lemma 3 (threshold level). $\tau^*(Y) \leq \Delta$ as soon as

$$Y \geq Y_\Delta = [k(\mu-1)\Delta]^{-1/(\mu-1)}$$

By Lemma 1 this level is reached by $x_\Delta = \max(1, Y_\Delta)/k$ at the latest. **Beyond that point, every window of length Δ in which the exponent stays above μ is sufficient on its own.**

Theorem (intermittent superlinearity). Under (S1):

(a) No blow-up without a superlinear tail. If $m(x) \leq 1$ for almost every x , then $1 + y^m \leq 2(1 + y)$ for all $y \geq 0$, hence $y(x) \leq (1 + y_0)e^{2kx} - 1$. **The trajectory is finite at every finite x , however long it is observed.**

(b) Almost-sure blow-up with one. If (S2) and (S3) also hold, then y reaches a finite-time singularity with probability 1. Partition $[x_\Delta, \infty)$ into consecutive windows of length Δ . By Lemma 3 each is sufficient, and by (S2) each succeeds with probability at least $q > 0$; by (S3) the attempts are (asymptotically) independent, so the probability of surviving the first N windows is at most $(1-q)^N \rightarrow 0$. **Moreover**

$$E[x^*] \leq x_\Delta + \Delta/q \quad P(x^* > x_\Delta + N\Delta) \leq (1-q)^N$$

What the theorem says, in words. A system of this kind does not need superlinear feedback on average, nor even most of the time. **It needs only that superlinear feedback be possible, recurrent, and capable of lasting a fixed positive interval.** The mean exponent, which determines everything in the constant case, drops out entirely: it affects how quickly the state climbs, and therefore the timing, but not the outcome. The separating condition is not the average but **the upper tail of the exponent's distribution together with the recurrence of its excursions**, and part

(a) shows this is sharp: remove the tail above 1 and the singularity disappears. Stated in one line, this is **the tail criterion: what decides the outcome is the upper tail of the feedback exponent, not its average.**

Why the requirement collapses. Writing $u = \ln y$ for $y \geq 1$ puts the mechanism in its clearest form, $du/dx \approx k e^{(m(x)-1)u}$. **An excursion of a given size above the boundary acts on the logarithm of the state with a strength that grows exponentially in that logarithm.** An excursion that is negligible early is decisive later, not because the excursion has changed, but because the state has.

Three predictions, and a numerical check. The theorem makes claims that can be tested directly against simulation. First, blow-up should occur for exponent processes whose mean is below the boundary. Second, the distribution of blow-up times should have an exponential right tail, from the geometric bound above. Third, **the timing should be controlled by how far above the boundary the exponent reaches, more than by anything else.** In a two-state exponent process with equal switching rates and $k = 0.05$, integrated numerically over 25 runs per setting, the median blow-up time behaves as follows at a fixed mean of exactly 1:

Upper state μ	1.1	1.2	1.4	1.8
Median x^*	286	159	80	44

Every run in every setting reached a singularity, including settings with a mean exponent of 0.76. **A separate check varied the persistence alone, holding the marginal distribution fixed, and found only a weak effect:** across switching rates spanning a factor of eighty, the median moved by a few percent. This is worth stating plainly, because it revises a reading suggested by the earlier part of this work, where more persistent noise appeared to bring the singularity forward. **In that construction, slower fluctuations also had larger amplitude, and the theorem indicates that the amplitude, not the persistence, was doing the work.** Persistence matters through (S2), by making q positive at a usable window length; beyond that threshold its influence is secondary.

Proposition 5 (prefactor noise is insufficient). The theorem is specific about *where* the fluctuation has to act. Suppose instead that the exponent is fixed at or below the boundary, $m \leq 1$, and that the noise enters the rate, so that $dy/dx = k(x)(1 + y^m)$ with $k(x) \leq k_{\max} < \infty$. Comparison with the constant-rate equation at $k_{\max}$ gives $y(x) \leq (1 + y_0)e^{2k_{\max}x} - 1$ by part (a) of the theorem. **No finite-time singularity occurs, however large, however persistent and however heavy-tailed the fluctuation in the rate.**

This is elementary, but it is the sharpest practical tool the theorem provides, and it applies directly to candidate systems. **A fluctuation that scales the growth rate, shifts the curve, or modulates the gain cannot produce this behavior; only a fluctuation that reaches the exponent of the feedback term can.** In examining any physical system, therefore, the first question is not how strong its noise is, but which of the two it perturbs. Many systems whose time series show sharp, intermittent build-up-and-collapse events are excluded at this step, and the exclusion is decidable from the governing equations rather than from the appearance of the data.

What the theorem does not establish. Assumption (S3) is not decorative. If the exponent is drawn once and then frozen, a realization below the boundary never blows up, and no amount of waiting helps; the result requires that the system keep drawing. The bound on $E[x^*]$ is also crude, since it credits the system with nothing during subcritical stretches and only counts whole windows; in the numerical checks above, blow-up arrives far earlier than the bound allows. **Most importantly, the theorem assumes the fluctuations as given.** It says nothing about where they come from, and in particular it does not yet cover the nested case, where the excursions are produced by smaller instances of the same equation. That case introduces a dependency the theorem does not handle: **finer levels run on faster clocks, so their excursions are not only smaller but shorter**, and whether their durations shrink faster than $\tau^*(Y)$ does is precisely the question Conjecture 2 of Section 4 turns on. That comparison, and the conditions a candidate system must satisfy, are taken up in Section 4.

4. OPEN CONJECTURES AND MEMBERSHIP CONDITIONS

The theorem of Section 3 is a statement about an equation. **Whether any real system is governed by such an equation is a different question,** and so is the question of where the fluctuations in one would come from. This section begins with what can be said with reasonable confidence: how a singularity can appear in a physical system at all, what a candidate system has to satisfy, and what can be resolved by measurement. **The conjectures follow afterwards, and are marked as such.**

4.1 Singularities in a physical system

In the physical world, no quantity actually becomes infinite. A singularity is therefore a property of an equation, not of a measurement. In a system described by such an equation, the divergence may be completed mathematically, but **the observed quantity has to leave it at some point:** what is recorded ends, changes character, or falls back to whatever remains beneath it. In the simulations of this work the same thing appears as the point at which the function stops being defined and its value becomes NaN. **The detachment is not an anomaly of the model but the form in which a singularity can appear in a physical system at all.**

What cuts the divergence off need not act at the level where measurement takes place. The level at which an equation applies and the level at which data is recorded may differ, so a divergence may be completed at the former and appear only as a bounded event at the latter. A system may also reach values that are effectively meaningless from the point of view of the observer, whether because no measurement remains possible, or because the quantity loses the interpretation it had over the range in which it was defined. **In such cases the distinction between a completed divergence and an extremely large finite value is not decidable by measurement**, and nothing in this work depends on deciding it.

Whether anything in nature diverges without limit is a separate question. **The expansion of the universe and its long-term future are the obvious place where a genuine infinity might be at stake**, and that question lies outside the scope of this work.

One form of regularization recurs often enough to be worth stating in advance. A whole family of realizations, including every one constructed in Section 6, is driven by a **closing gap**: two surfaces, two thresholds, or two levels of any kind, between which something travels, and whose separation shrinks. The observed quantity in such a system is a frequency, which is a speed divided by that separation, and it diverges as the separation goes to zero. **The separation cannot go to zero.** Beneath it there is always a scale set by the thing moving within it, and below that scale the gap is not merely small but undefined. In the circuit of Section 6 this is the resolution of the comparators, below which it is no longer determined which threshold is higher; between two surfaces it is the size of the body bouncing between them; in a gas it is the diameter of the particle. **The singularity is therefore not prevented by something acting against it, but by the disappearance of the quantity that was to diverge.**

This gives the regularization of a physical realization a definite location rather than a vague one. Where a candidate system is driven by a closing gap, the question "what truncates the divergence?" has a first answer that can be stated before any mechanism is investigated: the gap closes to the scale of its own contents and no further. The event that terminates the process occurs there, and the observed quantity reaches whatever value corresponds to that scale, which is finite and in principle calculable. **Any further truncation, by saturation, dissipation or loss, acts earlier than this one and is specific to the system; this one is not. The remark applies only where a closing gap is present**, which is a wide class but not every candidate: a system in which the diverging quantity is a stress, a density or an accumulated damage has no gap and must be examined on its own terms. Where there is one, however, the location of the cutoff is known in advance of any investigation of the mechanism.

What happens at that point is a redistribution rather than an ending. When the gap reaches the scale of the body within it, the body no longer traverses the gap, and the motion that was being counted ceases. The energy associated with it does not. It passes into the internal degrees of freedom of the body itself, which in a mechanical case means deformation and internal oscillation, and those have their own characteristic frequencies, set by the body's size and by the speed at which disturbances travel through it. **The observed quantity collapses while the energy does not**, and an observer counting the original events records a termination where the system has undergone a transfer.

Two things follow, and the second is a limitation. The first is that this offers one reading of why systems of this kind change character at close range rather than simply stopping: the interaction that dominates at separations large compared with the bodies is not the one that dominates when they are comparable, and the change of regime is a change in what the degrees of freedom are, not only in the strength of a force. **The second is that nothing here says the redistributed energy enters another instance of the same construction.** That would require a closing gap at the new scale and a non-decreasing quantity driving it, and neither follows from the transfer. **The observation is that the energy reappears elsewhere; where it goes afterwards is outside what this construction can say**, and the open question of Section 7 concerns exactly that step.

4.2 Membership conditions

For a system to fall within the scope of this paper, five conditions must hold, and each is checkable rather than rhetorical. **A system that fails any of them does not belong to the class, however suggestive its phenomenology.** Establishing membership requires reducing the system's own governing equations to the form of Section 1, not observing a visual resemblance in its time series.

4.2.1 Condition

A state variable that does not decrease over the regime considered, **at the level of a single instance of the equation.** This condition does not extend to the observable of a nested system. The observable may decline when it detaches from a completed instance and falls back onto the supporting instance beneath, and **the process as a whole may stop** if the outermost instance also ends in its singularity and the observable detaches from it as well (Part I, Sections 7.6.6 and 7.6.7). **Nor does a non-decreasing state imply an observable that grows without bound:** in physical realizations, the approach to a singularity presumably ends in an abrupt transition and a detachment rather than in a literal infinity, although systems not bound by a physical budget, such as a market price, may prove exceptions.

4.2.2 Condition

A feedback term whose effective exponent fluctuates across the critical boundary, with a non-vanishing probability of sustained excursions above it. The exponent must also vary on a timescale short compared with the one over which the state grows appreciably, while its individual excursions still persist long enough to matter. **This is a ratio and not a rate**, and the condition should not be read as requiring the fluctuation to be fast in absolute terms. **The construction has no timescale of its own:** by Proposition 0 the rate parameter is a pure rescaling of the horizontal axis and leaves the behaviour otherwise unchanged. A system whose exponent moves over decades while its state builds over centuries satisfies this

condition exactly as well as one whose exponent moves in microseconds while its state builds in milliseconds. **A candidate should therefore not be excluded for being slow, nor admitted for being fast**, and the comparison is always between two quantities belonging to the same system.

4.2.3 Condition

Physical systems of this kind are finite, so the dynamics need only hold over stretches of their history rather than throughout it. The early spread of a virus or of a bacterial population is an example: such a phase can show superlinear feedback and the blow-up character described here, while the system as a whole remains bounded by the resources available to it. Because of that bound, **the observed quantity cannot pass a certain level**; it leaves the instance it was following and falls back to a lower one, which may be zero. What is required is therefore not that a system behave this way throughout its history, but that at least one such stretch be identifiable. **The same system may also behave differently at different scales**: feedback found at one scale need not be present at another, and where it is present, its exponent need not be the same.

4.2.4 Condition

The observed trajectory cannot be smooth. Because the dynamics are driven by instances that end in singularities, the record has to contain abrupt rises and the falls that follow them. A smooth curve of the kind a slow moving average would produce indicates a different dynamic, not this one. **For the same reason, an exact fit is not to be expected**: the exponent is driven by noise, and in the simulations of this work repeated runs under identical parameters differ substantially in the timing and the size of individual events. **What can be compared with data is the character and the scaling of the trajectory, not the position of any particular excursion.**

4.2.5 Condition

The fluctuation must reach the exponent of the feedback term, not merely its rate.

4.3 Resolvability

The conditions of the preceding subsections are easier to state than to picture. This subsection therefore sets out a simulation of the construction, not as evidence for anything, but as **an illustration of the distance between an idealized nested system and the quantity an observer of it would actually record**. The code follows; what it does, and what it does not, is described first.

What is simulated. Four instances of the equation of Section 1 run simultaneously, with rates differing by a factor of ten at each step, so that each level is ten times faster than the one supporting it. **All four share a single exponent process**, drawn at each step from two values straddling the critical boundary. When any level exceeds its cutoff, that level and every level above it are reset to the value of the level beneath, which is the construction's own account of an instance completing its asymptote and the structure above it falling back onto its support. The lowest level resets the entire system. **The observable is not any of these levels**: it is a moving average taken over the logarithm of the fastest one, and it is drawn separately.

Environment. The simulation was run on Fedora Workstation under JupyterLab, in Python 3, using NumPy for the state arrays and the exponent draws, pandas for the rolling average, and Matplotlib for the figure. It requires no other dependency and runs in a few seconds.

Seven points about what this does and does not represent.

1. It shows the unbounded case. The system has no ceiling: each level runs to its cutoff and the cascade continues indefinitely. Real systems are bounded in the sense of Section 4.2.3, and nothing here addresses the terminal behaviour of a finite structure. **What it does illustrate, and what it is included for, is the measurement problem**, which is present whether or not the system is bounded.

2. The observable is a logarithmic moving average, and it is not faithful. This is the point of the figure. Even here, where the underlying structure is exactly known and perfectly nested, **the average tracks none of the four levels**: it runs above the supporting structure, below the excursions, and it climbs steeply in the approach to a large event because the excursions beneath it are growing, not because the support is. An exponent fitted to it would belong to no level in the system. **A single moving average is offered here as an illustrative device and not as a good readout**; a faithful reading would presumably require several, related to one another as the levels are, which is a conjecture rather than a result.

3. The nesting is idealized. Each level is an exact instance of the same equation, differing only in rate, and each resets exactly onto the level beneath. **A real system organized this cleanly is close to inconceivable**, and the figure should be read as a diagram of a relationship rather than as a model of anything.

4. The noise is imposed, where in reality it would be generated. All four levels are driven here by one common exponent sequence. In the construction this paper proposes, the fluctuation at a given level would come from the faster instances running within it, or from the system's fluctuations generally, and would therefore differ from level to level and be correlated with what the faster levels are doing. **The rates and the other parameters would likewise vary rather than remaining fixed.** The cutoffs are of the same character: they are numerical safety limits chosen to keep the arithmetic finite, standing in for physical mechanisms that would truncate a real divergence, and the simulation marks those steps as missing values so that the plotted line breaks instead of drawing a connection that did not occur.

5. It is not deterministic. The exponent is drawn at random at every step, so **no two runs are identical**. The fluctuation is fast relative to the growth, so the differences between runs are modest in appearance, but the timing of individual events is not reproducible and should not be read as if it were. **The mean exponent in the configuration shown is slightly below the critical boundary**, and the cascade nevertheless proceeds, which is the theorem of Section 3 in visible form.

6. Every parameter is exposed, and the illustration is meant to be varied. The seed of the generator, the two exponent values and the probability between them, the four rates, the cutoffs, the window of the moving average, the plotted range, and the number of levels can all be changed; so can the structure itself, by making the cutoffs dynamic, by giving each level its own exponent process, or by weighting the coupling between levels. **Reading the figure is less instructive than changing one parameter and watching what moves**, which is the reason the code is given in full rather than described.

7. One variation is worth recording. In runs where the fluctuation driving a level was taken from the faster levels within it, rather than from a common external sequence, **the approach to a large event was steeper by orders of magnitude**: a long, almost flat stretch followed by a rise that arrived abruptly. This is an observation about a simulation and not a result, and the coupling that produced it is the subject of Conjecture 1; it is noted here because it is the behaviour a reader varying the code is most likely to encounter, and because it differs markedly from the figure shown.

```
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt

# --- 1. PARAMETERS ---
A = 10e-17
B = 10e-5

k1 = 0.1 * B      # Level 1 (slowest)
k2 = 1 * B        # Level 2 (slower)
k3 = 10 * B       # Level 3 (slow)
k4 = 100 * B      # Level 4 (fast)

y_init = 1e-6     # Starting value for all levels

windowM4 = 700    # Moving average window for L4
plotWindow = 150000 # Adjustable observation window on the x-axis

# Noise configuration
p = 0.45
numSteps = max(50000, plotWindow + 2000) # Ensure simulation covers the plot window safely

# --- 2. NOISE GENERATION ---
m = np.random.choice([1.2, 0.8], size=numSteps, p=[p, 1-p])

# --- 3. SOLVING THE DIFFERENTIAL EQUATIONS ---
y1 = np.zeros(numSteps)
y2 = np.zeros(numSteps)
y3 = np.zeros(numSteps)
y4 = np.zeros(numSteps)

y1[0] = y_init
y2[0] = y_init
y3[0] = y_init
y4[0] = y_init

y1Visible = np.zeros(numSteps)
y2Visible = np.zeros(numSteps)
y3Visible = np.zeros(numSteps)
y4Visible = np.zeros(numSteps)

y1Visible[0] = y_init
y2Visible[0] = y_init
y3Visible[0] = y_init
y4Visible[0] = y_init

for i in range(1, numSteps):
    # Calculate increments based on previous step
    dy1 = k1 * (1 + y1[i-1]**m[i-1])
    dy2 = k2 * (1 + y2[i-1]**m[i-1])
```

```

dy3 = k3 * (1 + y3[i-1]**m[i-1])
dy4 = k4 * (1 + y4[i-1]**m[i-1])

y1_temp = y1[i-1] + dy1
y2_temp = y2[i-1] + dy2
y3_temp = y3[i-1] + dy3
y4_temp = y4[i-1] + dy4

# --- L1 Function (Red) Singularity Check ---
if y1_temp > 1e50:
    # Absolute reset of the entire system to y_init
    y1Visible[i] = np.nan
    y2Visible[i] = np.nan
    y3Visible[i] = np.nan
    y4Visible[i] = np.nan

    y1[i] = y_init
    y2[i] = y_init
    y3[i] = y_init
    y4[i] = y_init

# --- L2 Function (Magenta) Singularity Check ---
elif y2_temp > 1e40:
    y1[i] = y1_temp
    y1Visible[i] = y1[i]

    y2Visible[i] = np.nan
    y3Visible[i] = np.nan
    y4Visible[i] = np.nan

    # Reset everything above to the L1 level
    y2[i] = y1[i]
    y3[i] = y1[i]
    y4[i] = y1[i]

# --- L3 Function (Cyan) Singularity Check ---
elif y3_temp > 1e30:
    y1[i] = y1_temp
    y1Visible[i] = y1[i]

    y2[i] = y2_temp
    y2Visible[i] = y2[i]

    y3Visible[i] = np.nan
    y4Visible[i] = np.nan

    # Reset L3 and L4 to the L2 level
    y3[i] = y2[i]
    y4[i] = y2[i]

# --- L4 Function (Green) Singularity Check ---
elif y4_temp > 1e20:
    y1[i] = y1_temp
    y1Visible[i] = y1[i]

    y2[i] = y2_temp
    y2Visible[i] = y2[i]

    y3[i] = y3_temp
    y3Visible[i] = y3[i]

    y4Visible[i] = np.nan

    # Reset L4 to the L3 level
    y4[i] = y3[i]

# --- 5. Normal Operation (No Singularities) ---
else:
    y1[i] = y1_temp
    y1Visible[i] = y1[i]

    y2[i] = y2_temp

```

```

y2Visible[i] = y2[i]

y3[i] = y3_temp
y3Visible[i] = y3[i]

y4[i] = y4_temp
y4Visible[i] = y4[i]

# --- 4. OBSERVABLE QUANTITIES AND MOVING AVERAGES ---
f1Visible = A * (y1Visible / B)
f2Visible = A * (y2Visible / B)
f3Visible = A * (y3Visible / B)

f4 = A * (y4 / B)
f4Visible = A * (y4Visible / B)

absolute_baseLevel = A * (y_init / B)

# Moving average for L4 (Orange)
logF4 = np.log(f4)
logM4 = pd.Series(logF4).rolling(window=windowM4, min_periods=1).mean()
M4 = np.exp(logM4)

# --- 5. PLOTTING ---
plt.figure(figsize=(40, 20))

plt.axhline(absolute_baseLevel, color='gray', linestyle=':', linewidth=2, label='Absolute origin level (y_init)')

# Plot the four layers from top to bottom
plt.plot(M4[:plotWindow], color='orange', linewidth=1, label=f'Logarithmic Moving Average')
plt.plot(f4Visible[:plotWindow], color='lime', alpha=0.4, linewidth=1.5, label='Level 4: Fast (Green)')
plt.plot(f3Visible[:plotWindow], color='cyan', alpha=0.6, linewidth=2, label='Level 3: Slow (Cyan)')
plt.plot(f2Visible[:plotWindow], color='magenta', alpha=0.8, linewidth=2.5, label='Level 2: Slower (Magenta)')
plt.plot(f1Visible[:plotWindow], color='red', alpha=1.0, linewidth=3, label='Level 1: Slowest (Red)')

plt.yscale('log')
plt.title('Quadruple Cascading Fractal Cycles')
plt.xlabel('x')
plt.ylabel('Value (Log scale)')
plt.xlim(0, plotWindow)
plt.legend(loc='upper left')
plt.grid(True, alpha=0.3)

plt.tight_layout()
plt.show()

```

Figure 3.

Four observations follow from the figure, and they are the reason it is included.

The observed minima are not monotone, even in an unbounded system. The simulation has no ceiling, and its levels rise without limit; one might therefore expect the moving average to be pushed steadily upward as well. It is not. **In several places a later minimum of the average lies below an earlier one**, within the same stretch of the same supporting level. The reason is that the average is a reading of the fastest level and not of the support beneath it, so where it falls to depends on where the excursions happen to be when the structure above it resets, and not only on how far the support has risen. **This is worth stating explicitly because the opposite is a natural assumption:** that a system whose underlying state never decreases must produce a record whose troughs never decrease either. The figure shows that it need not, and it shows this in the simplest possible case, with the structure exactly known.

External intervention is a further difficulty, and a distinct one. Everything above concerns a system left alone. Where a system is added to or drawn from by an external agency, the record combines the dynamic with the intervention, and the two are not separable without knowing the intervention. **This is not a limitation of the framework so much as a limit on what any record of an intervened system can be asked to show**, and it applies with equal force to a deliberate control action, to a change of operating conditions, and to an unrecorded external input.

The quantity that moves the exponent need not be the quantity that is observed. What has been discussed so far concerns the relation between a record and the level it reads. There is a second gap, and it lies earlier: the exponent of the feedback is moved by some physical quantity, and there is no reason for that quantity to be the one an instrument happens to report. **Dynamic quantities are the natural candidates**, and each field has its

own: an accumulated mechanical stress, an evolving current profile, the entanglement of a field configuration. Whichever it is must be identified within the field concerned, and this work cannot do it. **The practical consequence is that a membership test may require two quantities rather than one:** the observed level, and the quantity suspected of moving the exponent, examined together rather than separately. A single channel may show the events clearly and carry no information at all about what produced them. It is entirely possible for a discrete event to register strongly in one quantity, such as a jump in power, while what it leaves behind, such as a persistent structural change or a residual stress, registers in another and only loosely with the first.

A genuinely unbounded system is a different case, and the origin becomes a fixed point. Mathematically an unbounded system may exist, and there the situation is simpler in one respect: a single instance reaches the singularity, and there is no averaging over anything, because there is no faster structure whose excursions have to be summarized. **Whether nested structure of the kind shown here can exist in such a system, and whether it would then admit an observable quantity of some other form, not a logarithmic moving average, is left open.** What can be said is that in that case the origin serves as a fixed reference point, since the state departs from it and never returns, and the whole trajectory can be read against it. **In a bounded system that reference is not available in the same way,** and the question of what replaces it is taken up separately.

4.3.1 Sampling a Part of the System: Anticorrelation Between the Part and the Whole

The preceding subsection concerned aggregation in *time*: over what interval a quantity is averaged. **A separate difficulty arises from aggregation in space,** that is, from measuring part of a system rather than all of it. The two are independent, and the second is the more counterintuitive, because it allows a correctly measured quantity to fall while the system it belongs to is rising.

The condition is about the whole. Where the accumulating quantity is extensive, so that the state of the system is the sum of the states of its parts, the non-decreasing requirement of Section 4.2 applies to that sum. **It does not apply to any individual part,** and nothing in the construction says it should. A sum can increase while some of its terms decrease, and if the parts exchange with one another, some of them must.

The elementary case, where the effect can be computed exactly. Take N units sharing a conserved total, $\sum_i x_i = X$, exchanging among themselves. Since the total is fixed, the variance of the sum vanishes:

$$\text{Var}(\sum_i x_i) = N \sigma^2 + N(N-1) \text{Cov}(x_i, x_j) = 0 \quad \rightarrow \quad \text{Cov}(x_i, x_j) = - \frac{\sigma^2}{N-1}$$

The correlation between any two parts is negative, and it is negative by necessity rather than by accident. For two parts the correlation is $-1/(N-1)$, which is weak when the parts are many. **Between one part and the sum of all the others it is exactly -1,** since the conserved total makes each the complement of the other: whatever one gains, the rest loses. **It is this second comparison, and not the first, that a measurement actually performs,** since an instrument reports a region and the quantity of interest belongs to the whole. This is the standard result for a conserved sum, and it is the reason a sampled portion is not merely a noisy copy of the whole: **it is an anticorrelated one.** In a system of a few hundred units the effect is weak for any single pair and considerable between a part and its complement.

A gas makes the effect concrete. The speed of one molecule in a container of N molecules varies erratically under collisions, while its long-time average approaches the ensemble average, and the time required for that convergence grows as the sample shrinks. Within a short window, one molecule may be observed slowing while the gas as a whole is being heated; over a long enough window it cannot. **What is measured here is not an error but a true statement about a part,** and the only mistake available is to read it as a statement about the whole. **The example is used here for the arithmetic of sampling alone,** and no position is taken on whether a gas, or any part of one, falls within the class described in this paper; that question is independent of the point being made.

The case that matters is when the exchange is not symmetric. Where every unit is equivalent, the anticorrelation is spread evenly and a sample is unbiased on average. Systems in which a process concentrates somewhere are not like this. In rock under load, microcracking is initially distributed through the volume and then **localizes into a narrow band which becomes the eventual fracture surface,** so that activity in that band rises steeply while activity elsewhere declines. The accumulated damage of the specimen as a whole is increasing throughout. **A sensor placed outside the band therefore records a decrease during precisely the interval in which the system is closest to failing.** The same structure appears wherever a process concentrates: a measurement confined to one region of a plasma, one part of a network, or one component of a loaded structure reports the behaviour of that region and not of the system.

The consequence for membership, in both directions. A candidate system should not be excluded because a measured signal decreases, if that signal reports a part of an extensive quantity rather than the quantity itself; the decrease may be the complement of an increase elsewhere. **Nor should a candidate be accepted because a local signal rises,** since a localizing process produces exactly that in the region where it concentrates, whatever the rest of the system is doing. The condition stated in Section 4.2 applies to the whole, and establishing that a measurement reports the whole is part of the work of identifying the state variable.

What makes this harder than ordinary sampling error. A noisy sample can be improved by averaging, and the error it introduces is unbiased. **Anticorrelation is neither:** averaging a fixed region for longer does not remove it, because the region is genuinely losing what the rest is gaining, and the bias has a direction set by where the process concentrates. **The worst case is also the most interesting one,** since localization sharpens as

the system approaches its limit, so the discrepancy between part and whole is largest exactly where the measurement matters most. Where the extent of the system can be covered, covering it is the remedy; where it cannot, the relation between the sampled region and the whole has to be argued rather than assumed.

4.4 Conjectures

Conjecture 1: the class is not empty. That some real system, in some domain, satisfies the conditions below, and therefore exhibits the behavior proved in Section 3. **A single verified instance would settle the existence question completely**, and no accumulation of unverified resemblances would do so. This asymmetry governs how the rest of this paper is written: candidate systems are examined one at a time, against explicit conditions, and a candidate that fails is not evidence against the conjecture, only a candidate that failed.

Conjecture 2: the fluctuations are generated by nesting. A further element distinguishes this construction from a system merely driven by external noise. The fluctuations are not assumed to be imposed from outside: **they may be produced by smaller, faster instances of the same equation nested within the system**, each running to its own singularity, being cut off by the system's own regularization, and contributing a bounded excursion to the exponent of the slower level that contains them. In this picture, noise is not a modeling convenience added to a deterministic law, but the aggregate trace of the same dynamic operating at a finer scale. **The mathematical statement does not depend on any particular device for aggregating these contributions**; a simulation may implement the aggregation in a simplified form, but nothing in the argument rests on that choice. Whether this nesting is necessary to the result, or merely one way of generating fluctuations with the required upper tail, is open. The theorem of Section 3 does not settle it, for the reason given at the end of that section.

Conjecture 3: the levels are coupled through a shared carrier. Conjecture 2 states where the fluctuations come from, but not how they reach the level they act on. **The coupling is presumably mediated by a single dynamic quantity that the levels share**: in a plasma, a quantity such as the local density, the current or the magnetic field, where the excursions to be understood are the smallest eddies and their abrupt detachment; in a financial system, the volume and velocity of capital flows; in a circuit, a voltage or a current. **The carrier need not be a physical quantity.** It may equally be an informational one, as in the spread of an item over a network, where the carrier is the rate at which it is passed on: that rate rises while the item is being shared and falls while it is only being consumed. (Such systems are finite and saturate, so at most their early dynamics could have this structure.) In each case the same quantity carries the events of the finer levels and delivers them, aggregated, to the exponent of the coarser ones. **If this is correct, every system in the class must possess such a carrier**, a single quantity in which the activity of the finer levels is present and, at least in principle, accessible to measurement. The moving-average construction of Part I, Section 7.6, where the fluctuating exponent is driven by an average taken over the observed signal itself, is one concrete realization of this idea rather than an additional assumption.

A second element belongs to the same conjecture. **The coupling is presumably not restricted to immediately adjacent levels.** In a fully developed system, the aggregated activity of every scale would act, with some weighting, on the exponent of every level, so that the exponent governing one level is a weighted functional of the whole hierarchy rather than of the level directly beneath it. The weights, and whether they decay with scale separation, are unknown. **This does not alter the theorem of Section 3**, which uses only the statistics of the exponent process at the level considered, whatever produced it; it constrains what a realization must look like, and what has to be measured in a candidate system. It also sharpens the search: **in any candidate, this transmission medium should be named before the membership conditions are checked**, since a system in which no such quantity can be identified is unlikely to support the nesting proposed in Conjecture 2.

Conjecture 4: the boundary of a finite system acts on the detachment, not on the equation. The systems this framework might apply to differ in how far their limits lie from the range in which they are observed. Some, such as a population in a closed habitat, are plainly finite; others appear unbounded simply because their limits are far away. **The conjecture is that this difference does not reach the equation.** In both cases the exponent presumably fluctuates around the critical boundary, and in both cases the instances run to their singularities. What finiteness changes is what happens at the end of that run: **the closer a system is to its own limits, the sooner the observable must detach from the instance it was following**, because the physical budget that would be needed to continue no longer exists. A saturating growth curve would then not be evidence against the structure; it would be the detachment made visible at the level where measurement takes place (see *Regularized singularities* above).

This has a consequence that the unbounded case does not have. In a finite system the supporting structures are themselves finite in number. If the outer, slower instances run to their singularities one after another, with nothing left to replace them, **the observable eventually detaches from the last support that remains**, and the system collapses to whatever level lies beneath, which may be far below the one it had reached. On this reading, a collapse of that kind is not a separate phenomenon requiring its own explanation, and need not be preceded by any visible external cause: **it is the terminal case of the same dynamic**, reached because the hierarchy that carried the growth was exhausted rather than because the growth was interrupted. The reading also carries a prediction: **such a collapse should be preceded by the scaling behavior of Section 2 and by detachments that become more frequent at the outer levels**, and not by a gradual flattening. Whether the exponent of a finite system does in fact fluctuate around the boundary, rather than falling away from it as the limits are approached, is the part of this conjecture most open to test.

5. BOUNDED SYSTEMS AND SIMULATIONS

The construction of Section 4.3 has no ceiling. Every physical realization does, and the obvious next step is a bounded version of the same simulation. **No such version is offered here, and the reason is itself the content of this section**: boundedness is not one mechanism. It enters the

construction at several distinct places, the places are independent of one another, and which of them governs a given system is a question about that system rather than about the construction. A single bounded simulation would therefore be a model of whichever choice its author happened to make, presented as though it were general.

Four places a bound can enter. They are listed in order of increasing depth, from the reading toward the structure.

1. A bound on the observed quantity. The reading is capped while the dynamic beneath it continues. This is the mildest case and probably the most common: an instrument saturates, a thermal capacity smooths, a channel compresses. **The level is bounded, the structure is not**, and the two part company at the cap.

2. A bound on each instance. Every instance terminates at a finite value rather than diverging, as the cutoffs in Section 4.3 do. The consequential question here is not the bound itself but **what the levels above it reset to**, and that is a matter of coupling rather than of the equation.

3. A bound on the structure. The number of levels is finite. When the outermost supporting instance completes its own asymptote there is nothing left beneath the observation, and the result is the terminal collapse described in Part I rather than another cycle.

4. A bound on the coupling. In the idealized construction the observed quantity may detach from the levels beneath it without limit. **In a bounded system it cannot**: the observable remains tied to the slower structure, and how tightly is a property of the system. This is the least tractable of the four and the one least amenable to a general rule, and it is the main reason a general bounded simulation is not attempted.

One consequence is general, and it is measurable. As a supporting level approaches its own limit, the faster instances running within it complete more often. **The change is in their rate, not primarily in their size.** A reading taken over a long window averages the increasingly frequent events together and therefore grows *smoother*; a reading of the same record over a short window encounters them more often and therefore grows *busier*. The two move in opposite directions at the same time, and either one taken alone misleads.

This gives a signature with some discriminating power. Genuine stabilization reduces activity on every timescale. The approach to a bound, on this construction, reduces it on long timescales while increasing the event rate on short ones. **Comparing the same record at two window lengths therefore separates the two cases**, where a single channel at a single resolution does not. The caution attached to it is real: events occur at every level of the structure, so what a given window length samples is a mixture, and the comparison indicates a direction rather than measuring a quantity.

Two scaling laws make the preceding paragraphs quantitative. Both follow from Section 2 and Section 3 rather than from any new assumption.

Event rate. A nested instance starting from the current support level Y and running at rate k_f completes, by Proposition 3, after an interval $Y^{1-m}/[k_f(m-1)]$. The frequency with which such instances terminate is the reciprocal:

$$\nu(Y) \propto k_f(m-1) Y^{m-1} \quad (\nu \propto k_f Y \text{ for } m = 2)$$

The events become denser as the support rises, and the density is a power of the support with the feedback exponent in it. Their individual durations shrink in the same proportion.

Smoothing. If the observed quantity is a superposition of such events, a reading averaged over a window W contains $N = \nu W$ of them, and the relative fluctuation of that reading falls as $N^{-1/2}$:

$$CV(W) \propto (\nu W)^{-1/2} \propto Y^{-(m-1)/2} W^{-1/2}$$

A long reading therefore smooths as a power of the support, for the same reason and at the same time as the events beneath it grow denser. This holds under the usual superposition assumptions, independent events of comparable shape, which the construction satisfies only approximately; a second effect works in the same direction, since the logarithmic span of each excursion narrows as the support approaches its bound.

The crossover, which is the quantity that actually moves. Dividing the two, the window length at which a reading stops resolving individual events and starts averaging them is

$$W^*(Y) \approx 1/\nu \propto Y^{-(m-1)}$$

which **contracts as the support rises**. This is the precise content of the earlier observation. Nothing becomes simultaneously calmer and busier; **the resolving scale slides downward past whatever fixed window an observer happens to use**, so a long window crosses into the averaging regime while a short one remains in the resolving regime, and the same record presents both appearances at once.

Estimating where the crossover lies is a free choice, and no particular estimator is prescribed. One possibility is to compare the variability of the same record at two window lengths, $W_1 < W_2$, taken on the increments of the logarithm rather than on its level, so that the trend does not enter the variance:

$$D_{W_1, W_2}(x) = \frac{\text{Var}_{W_1}(\Delta \ln f)}{\text{Var}_{W_2}(\Delta \ln f)}$$

Computed on the raw observed quantity rather than on an already smoothed one, **this ratio rises as W^* passes between the two windows**, and it is informative only while $W_1 < W^* < W_2$. Quantities of this family are standard in other settings; the variance-time curve, and the Allan and Hadamard variances, which remove a constant offset and a linear drift respectively, serve the same purpose and are better characterized.

Equally admissible are constructions this paper has not named: two distinct measured quantities of the same system compared at one window length, where their physical smoothing differs; a direct count of events per unit interval, which estimates ν without any variance at all; or any other formulation that separates the resolving regime from the averaging one. **What the construction supplies is the pair of scaling laws and the contracting crossover between them; which statistic is used to detect that contraction is a matter for the field concerned**, and depends on which channels it can measure and on what each of them already smooths.

Two further things a bounded system does that the idealized one does not. First, **detachment of the reading is not confined to the neighbourhood of a true asymptote**; in a bounded system it may occur at any earlier stage at which the reading can no longer follow. Second, where the levels do not share a single noise source, their trajectories need not stay ordered: **one instance may flatten while a slower one beneath it completes**, so the observed quantity can enter a prolonged plateau that terminates in an event not visible in it at all. Both behaviours depend on couplings this paper does not specify.

An illustration of the first kind of bound. The code below modifies only the reading of Section 4.3: the logarithmic moving average is passed through a soft cap that dampens it increasingly as it approaches a ceiling, leaving the four levels untouched. The result is a long, nearly straight approach before the reading detaches from the structure beneath it. **This is one arbitrary choice among many, and it is not a model of anything**; it is included to make the first entry point concrete and to show how readily a bounded-looking record is produced by bounding the reading alone. **Its inadequacy is as instructive as its appearance**: while the capped average is at its smoothest, the events beneath it are at their densest, which is exactly the opposite of what the smoothness suggests.

What a domain-specific study would have to fix. Four decisions, none of which this paper can make: which of the four bounds governs; what a level resets to when an instance terminates; whether the measured channel already smooths the structure through its own physics, in the sense of Section 4.3; and how tightly the observable is tied to the slower levels. **The construction of Section 4.3 is offered as a starting point for that work rather than as a result**, and the four decisions above are the ones that turn it into a model of something in particular.

5.1 Plateaus

A system of this kind need not look like one. The simulation below produces **a level that flattens into a long plateau and nevertheless terminates in a singularity**, and it is included for a methodological reason: **a candidate system that fails to show the expected shape is not thereby excluded**, and the shapes this construction can produce are more varied than the idealized figure of Section 4.3 suggests.

What is varied. Three levels run, with rates in the ratio 7 : 15 : 100. The middle level begins part-way through, taking its starting value from the level beneath, and **its upper exponent value is compressed as it grows**: where its observed quantity is small the exponent reaches 1.3, and as the quantity rises the reach falls toward 1.04, just above the critical boundary. The lower value is unchanged throughout. This is a damping that strengthens with amplitude, which is a common enough arrangement in a real system, and it is imposed here by hand rather than derived.

The mean exponent falls below the boundary, and the singularity arrives anyway. With the probabilities used, the mean of the middle level's exponent begins at about 1.025 and ends at about 0.908, while its upper reach ends at 1.04. By the theorem of Section 3 this is sufficient, since what is required is a positive probability of sustained excursions above the boundary and not a favourable mean. **Compressing the tail changes when the singularity arrives, not whether it does**, and what the compression buys in the record is the plateau: a long, almost flat stretch in place of a visible acceleration. The flattening and the eventual termination are the same mechanism seen at two amplitudes.

Three features of the figure are worth reading carefully.

1. The plateau is quiet on its own channel and loud on the one above it. When the middle level establishes itself, the level running on top of it begins to terminate far more often, and the change is abrupt. This is the event-rate law of Section 5 at work: the fast level now rests on a support many orders of magnitude higher, and its completion interval shortens accordingly. **An observer watching the plateau sees calm; an observer watching the faster level sees the opposite**, from the same system at the same moment.

2. The order of completion is not guaranteed. In the run shown, **the base level completes before the plateau level resting on it**, and the plateau continues afterwards. Where levels do not share a noise source, nothing orders their terminations, so a slower level may end while a faster one is

still flat. For an observer reading only the flat channel, the event arrives from a structure not visible in that channel at all.

3. The plateau has an interior. Its apparent smoothness is a property of the level itself, not of what is happening within it, and the density of events above it is highest precisely where it looks flattest. **Smoothness in one channel is not evidence of quiescence in the system**, which is the same conclusion reached in Section 5 by a different route.

What this does not establish, and the risk it carries. The damping is imposed rather than derived; the three levels draw on independent noise generators rather than on one another, so the feedback this paper conjectures is absent; and the fast level rests on the structure beneath without acting back on it. **The simulation therefore reproduces a shape and not a mechanism**, and it remains an idealization even though it is not an idealized shape. The risk in a result of this kind should also be named: **a framework that accommodates both the sharp case and the flat one is in danger of accommodating anything**. What limits that danger is that the two are not interchangeable. The plateau version makes different predictions, and they are checkable: the event rate of faster structures should rise sharply as the plateau forms rather than falling with it; the termination should be delayed rather than removed; and the completion order of levels should not be reliable. **A plateau that is genuinely quiet at every scale is not this construction**, and that is the observation that would exclude it.

```
import numpy as np
import matplotlib.pyplot as plt

# --- 1. PARAMETERS ---
A = 10e-17
B = 10e-5

# Growth rates for the three layers (adjustable)
kBase = 7 * B      # Level 1 - Magenta
kTop = 15 * B      # Level 2 - Cyan
kFast = 100 * B    # Level 3 - Lime

yInit = 1e-6
numSteps = 15000

# The step index where the Cyan plateau function branches off
startStepTop = 2500

# Noise configuration: probability and amplitude boundaries
p = 0.45
baseUpper, baseLower = 1.2, 0.8
topUpperStart, topUpperEnd = 1.3, 1.04
topLower = 0.8
fastUpper, fastLower = 1.2, 0.8

# Noise reduction thresholds for the plateau (Cyan) function
fTransitionStart = 0.00001
fTransitionEnd = 10.0

# --- 2. NOISE GENERATION ---
# Independent noise generators for each layer
isUpperBase = np.random.choice([True, False], size=numSteps, p=[p, 1-p])
isUpperTop = np.random.choice([True, False], size=numSteps, p=[p, 1-p])
isUpperFast = np.random.choice([True, False], size=numSteps, p=[p, 1-p])

# --- 3. ARRAY INITIALIZATION ---
yBaseVis = np.full(numSteps, np.nan)
yTopVis = np.full(numSteps, np.nan)

yFast = np.full(numSteps, np.nan)
yFastVis = np.full(numSteps, np.nan)

yBaseVis[0] = yInit
yFast[0] = yInit
yFastVis[0] = yInit

baseActive = True
topActive = False
fastActive = True

# --- 4. SOLVING THE DIFFERENTIAL EQUATIONS ---
for i in range(1, numSteps):
```

```

# L1: Base function
if baseActive:
    currentMBase = baseUpper if isUpperBase[i-1] else baseLower

    dyBase = kBase * (1 + yBaseVis[i-1]**currentMBase)
    tempBase = yBaseVis[i-1] + dyBase

    if tempBase > 1e40:
        baseActive = False
    else:
        yBaseVis[i] = tempBase

# L2: Branching plateau function starts at a specific step
if i == startStepTop and not np.isnan(yBaseVis[i]):
    topActive = True
    yTopVis[i] = yBaseVis[i]

# L2: Plateau function calculation with gradual noise reduction
elif topActive and i > startStepTop:
    currentFTop = A * (yTopVis[i-1] / B)

    if currentFTop <= fTransitionStart:
        currentTopUpper = topUpperStart
    elif currentFTop >= fTransitionEnd:
        currentTopUpper = topUpperEnd
    else:
        logF = np.log10(currentFTop)
        logStart = np.log10(fTransitionStart)
        logEnd = np.log10(fTransitionEnd)
        alpha = (logF - logStart) / (logEnd - logStart)
        currentTopUpper = topUpperStart - alpha * (topUpperStart - topUpperEnd)

    currentMTop = currentTopUpper if isUpperTop[i-1] else topLower

    dyTop = kTop * (1 + yTopVis[i-1]**currentMTop)
    tempTop = yTopVis[i-1] + dyTop

    if tempTop > 1e30:
        topActive = False
    else:
        yTopVis[i] = tempTop

# L3: Fast nested function that falls back to the highest active layer
if fastActive:
    currentMFast = fastUpper if isUpperFast[i-1] else fastLower
    dyFast = kFast * (1 + yFast[i-1]**currentMFast)
    tempFast = yFast[i-1] + dyFast

    if tempFast > 1e30:
        yFastVis[i] = np.nan

    # Cascading fallback logic
    if topActive and not np.isnan(yTopVis[i]):
        yFast[i] = yTopVis[i]
    elif baseActive and not np.isnan(yBaseVis[i]):
        yFast[i] = yBaseVis[i]
    else:
        fastActive = False
else:
    yFast[i] = tempFast
    yFastVis[i] = tempFast

if not baseActive and not topActive and not fastActive:
    break

# --- 5. OBSERVABLE QUANTITIES ---
fBaseVis = A * (yBaseVis / B)
fTopVis = A * (yTopVis / B)
fFastVis = A * (yFastVis / B)

absoluteBaseLevel = A * (yInit / B)

```

```
# --- 6. PLOTTING ---
plt.figure(figsize=(40, 20))

plt.axhline(absoluteBaseLevel, color='gray', linestyle=':', linewidth=2, label='Base Level')

plt.plot(fBaseVis, color='magenta', alpha=0.9, linewidth=3, label='Level 1: Slow')
plt.plot(fTopVis, color='cyan', alpha=0.8, linewidth=2.5, label='Level 2: Plateau')
plt.plot(fFastVis, color='lime', alpha=0.6, linewidth=1.5, label='Level 3: Fast')

plt.yscale('log')
plt.title('Plateau')
plt.xlabel('x (Steps)')
plt.ylabel('Value (Log scale)')
plt.legend(loc='upper left')
plt.grid(True, alpha=0.3)
plt.xlim(0, numSteps)

plt.tight_layout()
plt.show()
```

Figure 4.

6. A PHYSICAL REALIZATION (WORK IN PROGRESS!)

The simulations of the preceding sections are numerical. This section describes **an electronic circuit whose governing equation is the one studied here**, and in which the feedback exponent is not a single designed value but a quantity the arrangement can be made to take across its whole range, from well above the critical value to below it. It has not yet been built, and nothing below is a measurement.

The arrangement. A capacitor is charged at a steady rate. When its voltage reaches an upper threshold it is discharged, and when it falls to a lower threshold the charging resumes, so the circuit oscillates. The upper threshold is fixed. **The lower threshold is not a fixed level but the voltage on a second, independent capacitor**, referred to here as the **accumulator**, since its voltage only rises and it is this accumulation that drives the system. Writing G for the gap between the two thresholds, the oscillation frequency of a circuit of this kind is inversely proportional to that gap,

$$f \propto 1 / G$$

so that **as the accumulator voltage rises the window narrows and the frequency increases**. A comparator pair sets the two thresholds; because two comparators acting on the same node can contend near a transition, a small microcontroller is the simpler way to arbitrate between them, and it also provides the timestamping needed to record the output.

The accumulator voltage never decreases, whichever way it is charged, which is the non-decreasing state variable required by Section 4 realized structurally rather than imposed. The gap therefore closes monotonically, and the frequency rises monotonically with it. **What remains open is the law by which it closes**, and that is what fixes the exponent: by the relation of Section 6.2, the exponent follows from how the gap approaches zero and from how the observable reads it, and nothing else about the circuit enters.

Why frequency is the right observable. It can be counted rather than measured, so its precision is set by a clock rather than by an amplitude reference, and the resulting series needs no calibration in the sense of Section 4.3. It is nevertheless an *aggregating* readout in the sense of that section, since a frequency is always taken over a gate interval, and changing that interval changes what is reported. **This makes the circuit a test of the measurement problem as well as of the dynamic**: the same run, counted over two different gate lengths, should display the window-dependence described in Section 5.

Where the singularity stops. The binding limit is the one identified in Section 4.1: **the gap cannot close below the scale of what defines it**, which here is the input offset voltage of the comparators together with their noise. Once the gap is comparable to that offset, the two thresholds are no longer distinguishable and which one is higher is no longer determined, so the quantity whose reciprocal is being measured ceases to exist rather than becoming small. **This acts well before the dynamic limits do**: the comparator propagation delay and the slew rate of the charging stage bound the frequency only when the period approaches nanoseconds, whereas the offset is reached while the period is still long. **The regularization of this circuit is therefore a resolution limit and not a speed limit**, which is convenient, since an offset can be measured in advance and trimmed, while a resolution limit that has been characterized is a known truncation rather than an unknown one. **The divergence is therefore truncated by the apparatus, which is the regularization of Section 4.1 occurring for identifiable and measurable reasons**. Leakage from the accumulator sets a second limit, since the accumulated voltage is held against the comparator's input current and the leakage of whatever charges it; where that leakage balances the charge delivered, the accumulator stops rising and the frequency stops with it. **This limit must be**

characterized before a run rather than inferred from one, since an accumulator that has reached its leakage equilibrium produces a plateau that could be mistaken for the dynamic.

What it would and would not establish. The circuit is a constructed instance, so it cannot show that the class described in this paper is non-empty in nature; it was built to be a member. **What it can do is check the predictions against a system whose exponent is known from its own construction**, which no field measurement allows: whether the terminal scaling follows Proposition 2, whether the event rate follows the law of Section 5, and how much the exponent estimator is biased by noise and by the choice of gate length. It is, in that sense, a calibration of the instrument rather than evidence about the world.

1. Charging by a coupled winding: $m = 3$. A fixed quantity of charge is delivered once per cycle, so the gap closes per event rather than per unit time. Then $G \propto (t^* - t)^{1/2}$ and the frequency diverges as $(t^* - t)^{-1/2}$.

Why the exponent is not a setting. Suppose each cycle transfers a fixed quantity of charge to the accumulator. The rate at which the accumulator rises is then proportional to the frequency, and since the frequency is inversely proportional to the gap, the gap closes according to

$$\frac{dG}{dt} = -b/G \rightarrow G(t) = \sqrt{G_0^2 - 2bt}$$

The gap reaches zero in finite time, at $t^* = G_0^2/(2b)$. Rewriting this in terms of the observed quantity $y = f \propto 1/G$ gives

$$\frac{dy}{dt} \propto y^3 \quad y \propto (t^* - t)^{-1/2}$$

which is the equation of Section 2 with $m = 3$, and the terminal scaling $1/(m-1) = 1/2$ that Proposition 2 requires for that exponent. **Nothing was tuned to produce this.** The cubic arises because a quantity accumulating in proportion to a frequency closes a window whose width sets that frequency, and the exponent is a property of that arrangement rather than of a component value.

2. Charging through a resistor toward a level above the threshold: $m = 2$. The gap obeys $dG/dt = -(\Delta + G)/RC$ with $\Delta > 0$ the amount by which the charging level exceeds the threshold. As $G \rightarrow 0$ the closing rate approaches the constant Δ/RC , so the closing becomes linear in the limit and the frequency diverges as $(t^* - t)^{-1}$. A constant-current source gives the same exponent and closes linearly throughout.

A second case, and why it matters. If instead the charge transferred per cycle scales with the period, so that the transfer per unit time is constant rather than the transfer per cycle, the gap closes linearly and the same derivation gives $m = 2$. **Two plausible coupling behaviours therefore give two different exponents, and a real winding sits between them**, since the shape of the transferred pulse depends on frequency and departs from either idealization as the oscillation speeds up. The effective exponent is thus expected to drift during a run, for a physical reason, and to lie between 2 and 3. This is the situation described in Part III, Section 1.1 for a derived exponent, with the difference that here it can be measured: the slope of the logarithm of the growth rate against the logarithm of the frequency returns it directly, and the derivation above says what range the answer should fall in.

3. Charging toward exactly the threshold: $m = 1$. At $\Delta = 0$ the gap obeys $dG/dt = -G/RC$ and decays exponentially, approaching zero without reaching it. The frequency grows exponentially and no singularity occurs. **This is the critical boundary, and in this circuit it is a single point in the parameter space** rather than a regime: it is reached by one setting of the potentiometer and left by any deviation from it.

4. Charging toward a level below the threshold: saturation, and not a subcritical exponent. At $\Delta < 0$ the gap stops at a positive value and the frequency levels off. Expressed in the observed quantity this is

$$\frac{df}{dt} = Af - Bf^2$$

the logistic equation: the same superlinear term with its sign reversed. It is easy to mistake this for a subcritical exponent, since neither produces a singularity and in neither does the frequency fall. They are nevertheless different equations: a subcritical exponent gives unbounded growth as a power of time, $y \propto x^{1/(1-m)}$, which never levels off, whereas this settles at A/B . **Setting the potentiometer close to the threshold does not approximate a subcritical exponent; it produces a very high saturation level.**

What the four cases have in common. The passive circuit reaches $m = 3$, $m = 2$, the boundary at $m = 1$, and saturation. It does not reach anything between 1 and 2, and it does not reach anything below 1. The reason is the relation of Section 6.2: with $m = 1 + 1/(pq)$ and both factors positive, **any gap that closes in finite time gives $m > 1$** . A subcritical exponent requires the gap to approach zero more slowly than exponentially, as a power of time, which means the closing rate itself decaying as a power. **No passive charging path does this:** a resistor closes the gap at a rate proportional to what remains of it, which is always exponential, and a winding or a current source closes it at a rate that does not decay at all.

5. A charging current that decays with the gap: any exponent, including below 1. The restriction above is a restriction on passive charging paths, not on the arrangement as a whole. Suppose the charging current is made to follow the gap itself,

$$I \propto G^{1+1/a} \rightarrow G \propto t^{-a} \rightarrow m = 1 - 1/a$$

so that $a = 10$ gives $m = 0.9$, $a = 2$ gives $m = 0.5$, and $a = 1$ gives $m = 0$. **The gap then approaches zero without reaching it, and the frequency grows without bound and without a singularity**, which is the subcritical regime of Section 2 and is distinguishable from saturation precisely because it does not level off. The required current is a function of the present gap rather than of time, so the arrangement is a feedback loop and not a schedule: the gap is already measured by the circuit, and the charging current is set from it.

This exponent is imposed rather than found, and the distinction matters. The exponents of cases 1 to 4 follow from the geometry of the charging path and could not have been otherwise; the exponent of case 5 is whatever the feedback law is written to be. **Nothing in case 5 contradicts the conclusion above**, which concerns what a closing gap does when left to a passive mechanism. It does mean that the arrangement as a whole spans the full range of the construction, from well above the boundary to well below it, within one piece of hardware.

Why this matters for the theorem. Cases 1 to 4 are all at or above the critical value, so an arrangement restricted to them can test the deterministic results of Section 2, the scaling of Proposition 2 and the event-rate law of Section 5, but not the theorem of Section 3, whose subject is an exponent crossing the boundary. **With case 5 available the crossing becomes constructible:** a supercritical branch and a subcritical one, switched between at random, give an exponent process whose mean can be placed below 1 while its upper tail remains above. That is the configuration the theorem concerns, and it can be made to fluctuate at a chosen amplitude and persistence. **Whether the mean or the upper tail decides the outcome is then a measurement rather than a simulation.**

A caution that applies to case 5 alone. Because the law there is imposed, two artifacts belong to the apparatus rather than to the dynamics. The charging current is adjusted at discrete intervals, so the realized law is a stepwise approximation of the intended one, and the deviation grows as the gap closes and the intervals become long relative to the remaining time. And the gap is measured rather than known, so the feedback inherits the resolution limit of Section 4.1. **Both are quantifiable on this hardware precisely because the intended exponent is known**, which makes them a calibration of the estimator rather than a limitation of it.




Figure 5.

6.1 Mechanical Analogues of the Same Equation

The circuit of the preceding section is not the simplest realization of its own dynamics. **A mechanical system of two surfaces and a single ball between them obeys the same equation**, in one dimension throughout, and setting the two side by side shows which features of the construction

are essential and which belong to the particular hardware. The correspondence is offered as an analogy between equations; no claim is made that either system is a model of the other.

The elementary case. Let a ball travel at speed v between two parallel surfaces separated by a gap G , and let its speed be unchanged by each impact, so that no energy is exchanged. **This requires the surfaces to absorb the work they do**, by acting as a heat bath or by being held at rest at each impact and advanced between impacts; it is not what slowly moving elastic surfaces do, for the reason given below. The interval between impacts is G/v , so the collision frequency is

$$f = v / G$$

which is the relation $f \propto 1/G$ of Section 6, with the gap between two surfaces in place of the gap between two thresholds. **Now let the surfaces approach at a constant rate** $u = -dG/dt$. Differentiating,

$$\frac{df}{dt} = \frac{v u}{G^2} = \frac{u}{v} f^2 \rightarrow f \propto (t^* - t)^{-1}$$

This is the equation of Section 2 with $m = 2$, and the terminal scaling $1/(m-1) = 1$ that Proposition 2 requires for that exponent. The singularity occurs at $t^* = G_0/u$, when the gap closes.

The two drives, and why the circuit gives a different exponent. The same arrangement yields a different exponent depending on what closes the gap. If the gap closes *per unit time*, as above, the result is $m = 2$. If instead it closes *per event*, by a fixed amount δ at each impact, then $dG/dt = -\delta f$ = $-\delta v/G$, so G^2 decreases linearly and

$$G \propto (t^* - t)^{1/2} \quad f \propto (t^* - t)^{-1/2} \quad \frac{df}{dt} \propto f^3$$

which is $m = 3$, the exponent derived for the circuit of Section 6. **The difference between the two is not a difference of substrate but of bookkeeping:** in the circuit the coupled winding delivers its charge once per cycle, so the closing is per event; between two approaching surfaces the closing is per unit time. **The same structure therefore produces $m = 2$ or $m = 3$ according to whether the accumulating quantity is driven by the clock or by the events themselves.** This is worth noting beyond these two examples, since the distinction is easy to overlook in a candidate system and changes the predicted exponent by a whole unit.

What accumulates, and what does not. In the elementary case the ball's momentum is constant: it neither gains nor loses speed, and the frequency rises solely because the gap narrows. **The accumulating quantity is the displacement of the surfaces, not anything belonging to the ball.** The distinction matters, because the obvious alternative is to treat the ball's speed as the state variable, and that choice does not give a singularity at all.

Connection to Fermi acceleration. If the surfaces are *not* at rest at impact, each collision changes the ball's speed. In the classical treatment of this arrangement, the Fermi–Ulam model, randomly phased wall motion produces a diffusive growth of speed, with v increasing roughly as the square root of time. **Unbounded as that is, it is not a finite-time singularity**, and the reason is exactly the distinction of Proposition 5. Taking the surfaces to approach at rate u while also moving at impact, each collision adds $2u$ to the speed, and carrying this through the same derivation gives

$$\frac{df}{dt} = 3 \frac{u}{v} f^2$$

This relation is instantaneous rather than asymptotic, because the speed is no longer constant. Note that making the surfaces slow does not recover the previous case. Each impact adds $2u$, which vanishes as $u \rightarrow 0$, but the number of impacts grows as $1/u$, and the two cancel exactly: the invariant below holds whatever the closing speed. **The quasi-static limit of elastic surfaces is therefore the adiabatic case and not the isothermal one**, and the two differ by where the work goes rather than by how fast the surfaces move. Dividing the two rates gives $dv/v = -dG/G$, so $v G$ is conserved: this is the adiabatic invariant of the one-dimensional billiard, and it holds increasingly well as the speed grows past the wall speed. With $v \propto 1/G$ the frequency scales as $f \propto 1/G^2$, and a gap closing at constant rate then gives

$$\frac{df}{dt} \propto f^{3/2} \quad f \propto (t^* - t)^{-2}$$

so that moving walls change the exponent from 2 to 3/2 rather than merely rescaling it. This is not a counterexample to Proposition 5, which concerns noise in the rate of an equation whose exponent is fixed; here the relation between the two state variables is itself altered, which is a change of structure rather than of coefficient. **The distinction is worth keeping in view, since a mechanism can act on the exponent without being a fluctuation at all.**

One feature of the Fermi–Ulam literature bears directly on assumption (S3) of Section 3. **Smooth, deterministic wall motion does not produce unbounded acceleration**, because invariant tori in the phase space bound the attainable energy; randomly phased motion does. The dividing line is not the strength of the driving but whether successive impacts are effectively independent. **That is the same condition the theorem requires, arising in a system with no connection to the present construction**, and it is some evidence that the condition is structural rather than a convenience of the proof.

Where the exponent itself can move. Neither of the above makes m fluctuate. What does is the collision law. If the restitution of the impact depends on impact speed, the loss per collision scales as some power of the speed, and that power enters the exponent directly. **For real materials this dependence is measured rather than assumed:** for viscoelastic bodies the coefficient of restitution decreases with impact speed, following a fractional power. The exponent is then a material property, in the same way that the exponent of a hyperabrupt varactor is a property of its doping profile rather than of the circuit around it. **A mechanical realization of a fluctuating exponent would therefore be built by varying the collision law, not the wall motion**, which is the mechanical statement of the distinction this paper draws throughout.

A caution from the simplest case of all. A ball bouncing on a single fixed surface under gravity, losing a fixed fraction of its speed at each impact, undergoes infinitely many collisions in finite time and comes to rest at a definite moment. **Here the collision frequency diverges while the ball's speed goes to zero**, and the accumulating quantity is the energy already lost rather than anything that grows. This is the reverse of the case treated above, and it carries a warning for the measurement problem of Section 4.3: **a diverging event rate does not establish that any state variable is diverging**, and a record of event times alone cannot distinguish a quantity running away from one running down. Which of the two is occurring follows from the governing equation, not from the frequency.

What these analogues establish. Nothing about the physical world: each is a textbook arrangement, and the correspondences are between equations. **What they do show is that the exponent of this construction is not a free parameter** but is fixed by the geometry of the accumulation in three independent realizations, an electronic one and two mechanical ones, giving 2 or 3 according to the bookkeeping rather than according to the substance.

6.2 What the Exponent Encodes: Closing Gaps and the Relation $m = 1 + 1/p$

The distinction drawn above between fixed and moving surfaces is not a peculiarity of this construction. It coincides with a distinction long established in the kinetic theory of gases, and the coincidence is worth recording, since it shows that the exponent separates two processes that were already known to differ.

Consider the same ball, of mass M , and average the force it exerts on one surface over many impacts. Each impact reverses the momentum, delivering $2Mv$, and impacts on that surface occur at a rate $v/2G$. The mean force is therefore

$$F = 2Mv \cdot \frac{v}{2G} = \frac{Mv^2}{G} \quad \rightarrow \quad FG = Mv^2$$

which is **the one-dimensional ideal gas law**, with G in place of the volume and F in place of the pressure; assigning $\frac{1}{2}kT$ to the single degree of freedom gives $FG = kT$. **This derivation is the standard one of kinetic theory and is not offered as a new result**; what follows from it is the point of this note.

The two cases of this subsection are the two classical compressions. Where the surfaces are at rest at impact, the speed is constant, so $F \propto 1/G$. That is Boyle's law, and the process is **isothermal**; it is also the case that gives $m = 2$. Where the surfaces move at impact, the invariant derived above is $vG = \text{constant}$, so that

$$F = \frac{Mv^2}{G} \propto \frac{1}{G^3} \quad \rightarrow \quad FG^3 = \text{constant}$$

and $F G^\gamma = \text{constant}$ with $\gamma = 3$ is **the adiabatic law for a one-dimensional monatomic gas**, for which $\gamma = 1 + 2/d = 3$ at $d = 1$. **The restriction to one dimension is not incidental and applies to every derivation in this section:** in three dimensions the same compression gives $\gamma = 5/3$, and the corresponding exponent is $m = 7/4$ rather than $3/2$, since the contents then gain speed more slowly as the gap narrows. That is the case giving $m = 3/2$. **The feedback exponent therefore distinguishes isothermal from adiabatic compression:** the walls that take no energy from the ball leave the temperature unchanged, and the walls that do work on it heat the gas.

Two things are worth taking from this, and neither is a claim about thermodynamics. First, **the exponent of this construction is not a free parameter even in the simplest mechanical case:** its two available values correspond to two processes that were distinguished long before this equation was written down, and the correspondence is exact rather than suggestive. Second, the quantity that accumulates differs between the two, being the displacement of the surfaces in the first case and the energy of the ball in the second. **Identifying which one accumulates is the same question that Section 4 requires a candidate system to answer**, and here it has a textbook answer, which is rarely the case elsewhere in this paper.

A rule covering every case derived so far. The exponents obtained in this section and the preceding one were each derived separately, but they follow from a single relation. Let the observed quantity diverge as a power of the remaining time,

$$y \propto (t^* - t)^{-s} \rightarrow m = 1 + 1/s$$

which follows by differentiating: $dy/dt \propto (t^* - t)^{-(s+1)} = y^{(s+1)/s}$. **The exponent is fixed by the terminal rate of divergence of whatever is observed, and by nothing else.** In the arrangements of this section that rate factors into two parts, because the observed quantity is a frequency and a frequency is a speed divided by a gap. Writing $G \propto (t^* - t)^p$ for the closing of the gap and $y \propto G^{-q}$ for the reading taken of it,

$$s = p q \rightarrow m = 1 + 1/(p q)$$

Here p belongs to the mechanism that closes the gap and q to what is being read. Where the speed is constant the frequency is simply the reciprocal of the gap, so $q = 1$ and the relation reduces to $m = 1 + 1/p$. Where the contents gain speed as the gap narrows, as they do whenever the surfaces do work at impact, $q > 1$. Applied to the cases treated here:

Arrangement	Gap closes as	p	q	m
Surfaces approaching at constant rate; accumulator charged at constant current	linearly	1	1	2
Fixed amount removed per event; accumulator charged by a winding	$\sqrt{t^* - t}$	1/2	1	3
Approaching surfaces that also move at impact, in one dimension (adiabatic case)	linearly	1	2	3/2
Inverse-square attraction, idealized and taken as a separation closing	$(t^* - t)^{2/3}$	2/3	1	5/2

Only the behaviour near the singularity enters the table. A case not listed there illustrates the point. Charging an accumulator through a resistor toward a fixed level does not close the gap linearly, but exponentially, as $dG/dt = -(\Delta + G)/RC$ where Δ is the amount by which the charging level exceeds the threshold. **As $G \rightarrow 0$ the rate approaches the constant Δ/RC ,** so the closing becomes linear in the limit and the exponent is the same as for a constant-current source, which closes it linearly throughout. The two differ in the approach and agree at the end, which is all that m records. This also makes plain why $\Delta > 0$ matters: at $\Delta = 0$ the gap closes as $e^{-t/RC}$ and never reaches zero, which is the critical case $m = 1$, and at $\Delta < 0$ it stops at a positive value and the behaviour is self-limiting.

Two consequences, and the second is a limitation. The first is that an exponent measured from data is a statement about the closing law and can be read as one: where the possible closing mechanisms of a system are known and give different values of p , the measurement discriminates between them, which is the situation described for accelerating creep in Part III. **The second is that the exponent does not identify a mechanism on its own.** Two systems with no physics in common return the same m whenever their gaps close in the same way, as the first row of the table shows for an electronic and a mechanical arrangement. **Agreement of exponents across domains is therefore not evidence of a shared mechanism,** and the isomorphism this paper claims is an isomorphism of equations for exactly that reason. The membership conditions of Section 4 ask for the governing equation and the accumulating quantity precisely because the exponent alone does not supply them.

A third point follows for measurement. The relation above assumes the observed quantity is the reciprocal of the gap. **Read a different quantity and the same system returns a different exponent:** in the inverse-square case above, taking the speed rather than the reciprocal separation as the observable gives $m = 4$ for the same motion. This is the mapping of Section 2 appearing again, and it reinforces the requirement of Section 4.3 that the relation between the record and the underlying state be settled before any exponent read from it is interpreted.

7. A NOTE ON SCALE

One property of the construction deserves comment, because it invites a stronger reading than it supports. **The equation has no intrinsic scale.** By Proposition 0, the rate sets only the clock and the mapping constant only the unit, so nothing in the mathematics fixes a magnitude, a duration, or a domain at which it is meant to apply. The same is true of the nested form: the construction of Section 5 can be extended downward to faster instances or upward to slower ones without any term acquiring a preferred size.

This is a property of the equation, and it is not evidence about the world. Scale invariance in a model means that the model does not select a scale; it does not mean that the systems the model might describe are related to one another across scales. The membership conditions of Section 4 apply at every level independently: a candidate at any scale must have its governing equation reduced to this form, its non-decreasing state variable identified, and the quantity that moves its exponent established. **Nothing is inherited from one scale to another, and a correspondence found at one says nothing about any other.**

The temptation the absence of scale creates should therefore be named and set aside. **A framework that can be written at any scale is not thereby a framework that connects scales,** and the distance between the two is the entire content of the membership conditions. This paper makes the first claim and not the second.

One question is worth leaving open, stated as a question. In the construction as it stands, an instance that completes its asymptote leaves the observation: the moving average detaches from it and falls back onto the structure beneath. **The construction says what happens to the reading. It does not say what happens to the instance.** That silence is worth noticing, because it is not the silence of something that has been ruled out.

Two facts make the question well posed rather than idle. **The first is that nothing is lost at the detachment.** The moving average remains finite through it, not because the instance beneath it stops, but because the logarithm of a function diverging at a finite point is integrable; the instance completes its asymptote in full while the reading of it stays bounded. **The detachment is a property of the observation, not an event in the thing observed.** The second is that the completion occurs at a point. A finite-time singularity is reached in finite x , so whatever the completed instance amounts to, **it is brief when measured on the clock of any slower level**, by the factor separating their rates.

The question, then, is this. Suppose one sets aside the structure the instance grew out of, and asks only what a completed instance is, viewed from a slower level. It is a quantity that arrives, is enormous on its own scale, and is over in an instant on the slower one. **On that level it has the form of a fluctuation.** And a fluctuation, in this framework, is not a disturbance but the thing that moves the exponent, and hence the thing from which a sequence of instances can begin. **Whether a completed instance can serve in that role, becoming the starting point of a new structure rather than only the end of an old one, is not answered here.**

What this paper does not ask, deliberately. How such an event would act on the slower structures already present, whether it would be absorbed, amplified or ignored, is a question of coupling and of parameters, and is left alone. So is any question of what the completed instance consists of, which the mathematics does not address. **The question retained here is narrower and structural:** whether the construction admits a second kind of transition alongside the one it describes, in which a completion is not only an ending in one place but an origin in another.

This is recorded as a question and nothing more. It is not a conjecture in the sense of Section 4, since no observation is proposed that would bear on it and no candidate system is named. **It is placed in this section, among the things the absence of scale does not license, precisely so that it is read as a question about the mathematics rather than as a claim about the world.**

REFERENCES

- Osgood, W. F. (1898). *Beweis der Existenz einer Lösung der Differentialgleichung $dy/dx = f(x, y)$ ohne Hinzunahme der Cauchy-Lipschitz'schen Bedingung*. Monatshefte für Mathematik und Physik, 9, 331–345. DOI: 10.1007/BF01707876.
 - Teschl, G. (2012). *Ordinary Differential Equations and Dynamical Systems*. American Mathematical Society, Providence, RI.
 - León, J. A. & Villa, J. (2011). *An Osgood criterion for integral equations with applications to stochastic differential equations with an additive noise*. Statistics & Probability Letters, 81(4), 470–477. DOI: 10.1016/j.spl.2010.12.001.
 - Foondun, M., Khoshnevisan, D. & Nualart, E. (2024). *Instantaneous everywhere-blowup of parabolic SPDEs*. Probability Theory and Related Fields, 190(1), 601–624. DOI: 10.1007/s00440-024-01263-7.
 - Fernández Bonder, J. & Groisman, P. (2009). *Time-space white noise eliminates global solutions in reaction-diffusion equations*. Physica D: Nonlinear Phenomena, 238(2), 209–215. DOI: 10.1016/j.physd.2008.09.005.
 - Jensen, J. L. W. V. (1906). *Sur les fonctions convexes et les inégalités entre les valeurs moyennes*. Acta Mathematica, 30, 175–193. DOI: 10.1007/BF02418571.
 - Taleb, N. N. & Douady, R. (2013). *Mathematical definition, mapping, and detection of (anti)fragility*. Quantitative Finance, 13(11), 1677–1689. DOI: 10.1080/14697688.2013.800219.
 - Norris, J. R. (1997). *Markov Chains*. Cambridge University Press.
 - Davis, M. H. A. (1984). *Piecewise-deterministic Markov processes: a general class of non-diffusion stochastic models*. Journal of the Royal Statistical Society: Series B, 46(3), 353–388. DOI: 10.1111/j.2517-6161.1984.tb01308.x.
 - Lawley, S. D., Mattingly, J. C. & Reed, M. C. (2014). *Sensitivity to switching rates in stochastically switched ODEs*. Communications in Mathematical Sciences, 12(7), 1343–1352.
 - Liberzon, D. (2003). *Switching in Systems and Control*. Birkhäuser, Boston. DOI: 10.1007/978-1-4612-0017-8.
 - Platt, N., Spiegel, E. A. & Tresser, C. (1993). *On-off intermittency: A mechanism for bursting*. Physical Review Letters, 70(3), 279–282. DOI: 10.1103/PhysRevLett.70.279.
 - Durrett, R. (2019). *Probability: Theory and Examples*. Cambridge University Press. DOI: 10.1017/9781108591034.
 - Johansen, A. & Sornette, D. (2001). *Finite-time singularity in the dynamics of the world population, economic and financial indices*. Physica A, 294(3–4), 465–502. DOI: 10.1016/S0378-4371(01)00105-4.
 - Sornette, D. (2003). *Why Stock Markets Crash*. Princeton University Press, Princeton, NJ.
 - Mandelbrot, B. B. & Wallis, J. R. (1968). *Noah, Joseph, and operational hydrology*. Water Resources Research, 4(5), 909–918. DOI: 10.1029/WR004i005p00909.
 - Mandelbrot, B. B. (1982). *The Fractal Geometry of Nature*. W. H. Freeman, San Francisco.
 - Fermi, E. (1949). *On the Origin of the Cosmic Radiation*. Physical Review, 75(8), 1169–1174. DOI: 10.1103/PhysRev.75.1169.
 - Ulam, S. M. (1961). *On some statistical properties of dynamical systems*. Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 3: Contributions to Astronomy, Meteorology, and Physics, 315–320. University of California Press.
-

The Geometry of Bitcoin:

Physical Candidates and Isomorphisms

1. POSSIBLE CONNECTIONS: TOKAMAK EDGE INSTABILITIES AND METASTABLE ESCAPE

The systems examined so far were candidates for the mechanism of Part II, Section 3: an exponent that fluctuates across the critical boundary. **Tokamak edge localized modes (ELMs), as they are currently described in the literature, are not such a system, and the purpose of this section is to say precisely why, and what the established description amounts to instead.** That literature derives a genuine finite-time singularity from first principles; but in the derivations available at the time of writing, **the nonlinearity is quadratic and fixed**, arising from the second term of a Taylor expansion about marginal stability rather than from any stochastic quantity. What varies is the linear coefficient and the size of the perturbation that arrives. On that description, ELMs fall into a third class, **metastable escape**, which this section introduces and which the remainder of the paper treats as distinct from both exponent noise and prefactor noise.

The generalized equation. To accommodate this class, the equation of Part II, Section 1 requires one additional term, a linear damping:

$$\frac{dy}{dx} = -a(x)y + ky^m$$

The constant term of the original equation is dropped, because here the origin is an equilibrium rather than a driven state. **The sign of a decides the character of the system.** For $a < 0$ the origin is unstable and every trajectory departs it. For $a > 0$ the origin is *linearly stable*, yet for $m > 1$ the superlinear term still dominates at large amplitude, so **the system is stable against infinitesimal perturbations and unstable against sufficiently large ones.** This is metastability, and the original equation of Part II, Section 1 cannot represent it: with a positive constant term, every trajectory grows from the start and no threshold exists.

The threshold and the escape time. For $m = 2$ the equation is separable and yields a closed solution. Writing the threshold amplitude as $y_c = a/k$:

$$y(x) = \frac{a}{k - (k - a/y_0)e^{ax}} \quad x^* = -\frac{1}{a} \ln(1 - y_c/y_0)$$

For $a < 0$ the blow-up time is finite for every initial amplitude. **For $a > 0$ it is finite if and only if $y_0 > y_c$;** below the threshold the trajectory decays to zero. Two features are worth noting, because both reappear in the plasma. The escape time diverges logarithmically as the initial amplitude approaches the threshold from above, so **events triggered just above threshold are slow and those triggered well above it are abrupt.** And as a is driven toward zero, the threshold y_c falls in proportion, so **an ever smaller perturbation suffices.**

The isomorphism. Nonlinear ideal ballooning theory describes the eruption of a plasma flux tube. In the drag-dominated limit, Cowley, Cowley, Henneberg & Wilson (*Proc. R. Soc. A* 471, 20140913, 2015) obtain for the displacement ξ of a field line

$$(v_D/\rho_0) \frac{d\xi}{dt} = \Gamma_0^2 \xi + C_4 \xi^2$$

where v_D is a drag coefficient, ρ_0 the mass density, Γ_0^2 the local linear drive and C_4 the coefficient of the explosive nonlinearity. **Dividing through by v_D/ρ_0 gives exactly the equation above,** under the correspondence

$$y \leftrightarrow \xi \quad a \leftrightarrow -\Gamma_0^2 \rho_0 / v_D \quad k \leftrightarrow C_4 \rho_0 / v_D \quad m = 2$$

The correspondence is exact, not approximate, and the plasma literature derives the same two consequences independently. The critical amplitude for a linearly stable field line is $\xi_c = |\Gamma_0^2|/C_4$, which is the threshold $y_c = a/k$ above; and in the linearly unstable case the eruption occurs at $t_\infty = (1/\gamma) \ln(1 + \xi_c/\xi_0)$, which is the escape time above with $a = -\gamma$. **The tokamak edge, in this limit, is computing the same equation.**

Where the singularity comes from, and why the exponent is fixed. The general theory (Cowley & Artun, "Explosive instabilities and detonation in magnetohydrodynamics", *Phys. Rep.* 283, 185, 1997; Wilson & Cowley, "Theory for Explosive Ideal Magnetohydrodynamic Instabilities in Plasmas", *Phys. Rev. Lett.* 92, 175006, 2004) expands the energy functional about the marginal state. The quadratic term of that expansion is the explosive nonlinearity, so **$m = 2$ is structural within that treatment: it follows from the expansion itself, and not from any property of the plasma that the treatment allows to fluctuate.** Whether a description in which the effective exponent is not fixed in this way would better match the observations is not a question this paper can settle, and it is left open here. In slab geometry the displacement diverges as $(t_0 - t)^{-2.1}$; numerical

solutions of the full equation give exponents near -2.5 . In toroidal geometry the inertia term becomes a fractional time derivative of order λ , set by the Mercier indices, and the divergence exponent becomes λ itself, reducing to $(t_f - t)^{-1}$ for a vanishing Mercier coefficient.

A remark that follows from this, and that applies well beyond the plasma. Balancing a time derivative of order λ against a nonlinearity of order m gives $y \propto (t_f - t)^{-\lambda/(m-1)}$. Part II, Section 2 is the case $\lambda = 1$. **A measured divergence exponent therefore does not identify m on its own:** the same value is produced by first-order dynamics with one exponent and by second-order dynamics with another. The slab value near 2.1 is consistent with $m = 2$ and $\lambda = 2$, not with a larger nonlinearity. Any attempt to read a feedback exponent from an observed blow-up profile must therefore establish the order of the governing dynamics first. This observation is the present author's; it is not drawn from the plasma literature.

The mechanism, on the established picture: a slow ramp and a sudden kick. Between events the pedestal is heated and fuelled, its pressure gradient rises toward the peeling–ballooning limit, and the linear coefficient drifts toward zero: in the notation above, **$\alpha(x)$ decreases and with it the threshold $y_c = \alpha/k$** . The edge is metastable throughout this phase. Linear instability is not required for eruption, since a finite—though not infinitesimal—perturbation will destabilize a linearly stable profile. The asymmetry involved is severe: in the model atmosphere of Cowley et al. (2015), a critical perturbation of about 1% of a field line's magnetic energy releases roughly 17% of it. In tokamak geometry, Ham et al. (*Phys. Rev. Lett.* 116, 235001, 2016) find metastable saturated states over a window of the normalized pressure gradient lying *below* the linear stability boundary.

This account also offers a reading of why pellet pacing works. Injecting a pellet is an engineered perturbation above the threshold, and it triggers an ELM before the natural ramp would have delivered one. Because the event is triggered while α is still comparatively large, the accumulated pedestal is smaller and the resulting crash is milder; raising the ELM frequency above its intrinsic value was found to reduce the severity of the individual events (Lang et al., *Nucl. Fusion* 44, 665, 2004). **In the framework above, pacing does not alter the dynamics; it chooses the moment of escape.**

Detachment: the singularity is never reached. As Part II, Section 4.1 requires of any physical realization, the divergence is truncated. Several mechanisms are proposed in the literature, and on the standard account they act in sequence. The weakly nonlinear expansion itself breaks down well before the formal singularity. The erupting tube then narrows, twists and either saturates into a finitely displaced equilibrium or is broken up by the sheared edge flow. Finally, resistive reconnection and cross-field heat transport along the erupted tube vent the stored energy. **What is observed is therefore not an infinity but an abrupt, bounded event:** a type-I ELM expels of the order of ten percent of the pedestal energy, strongly dependent on collisionality, after which the gradients recover within a few milliseconds and the cycle repeats at tens of hertz. The correspondence with the regularized-singularity picture of Part II, Section 4.1 is close, and in this respect the plasma is the best-documented example available.

Where the state condition does and does not apply. It would be easy to conclude from the ELM cycle that the non-decreasing state condition fails, since the pedestal rises monotonically only within a single inter-ELM interval and is reset at each crash. **That conclusion would apply the condition at the wrong level.** As Part II, Section 4.2.3 sets out, physical systems of this kind are bounded, the observed quantity cannot pass a certain level, and it must leave the instance it was following and fall back to a lower one, which may be zero. A reset of the observable is therefore *expected* in any bounded realization, and is not by itself evidence against membership. **The condition constrains the latent variable, not the record.**

What is genuinely missing is the identification of that variable. In the established description, the quantity that rises between crashes is the pedestal itself, and that quantity is reset; impurities, so far from accumulating, are expelled by the crashes, which is one reason ELM-free regimes require an alternative exhaust channel. No slower, irreversible quantity has been identified that would play the role of the underlying state in Part II, Section 3. **This is an absence of identification rather than a demonstrated absence**, and the distinction matters: the same would have been said of any candidate system in this paper before its governing equation was reduced. The question is taken up again at the end of this section. A related one is whether the ELM sequence has a terminal member at all, in the sense in which the nested construction ends when its outermost supporting instance completes its own asymptote, and whether such an event would be distinguishable from the other ways a discharge can end. **Neither question is addressed by the literature reviewed here.**

How the sequence is currently described. Independently of the above, the established treatment of the ELM sequence is a renewal process with reset, and the ELM literature has adopted one: under the assumption that the likelihood of an ELM rises monotonically with time since the last, the waiting times follow a Weibull-class distribution, which was confirmed on JET data (Webster & Dendy, *Phys. Rev. Lett.* 110, 155004, 2013). **This description is compatible with both readings**, since a renewal structure with an increasing hazard rate is what a ramped threshold produces and also what a bounded nested structure produces between its terminal events; it therefore does not discriminate between them.

An open question about the reset. If the state were fully reset at each crash, and nothing carried across, the cycle would be governed entirely by the ramp and the noise, and sustained suppression would be a matter of holding the threshold above the ambient perturbation level. In part this is what is observed: suppression over long pulses has been achieved with resonant magnetic perturbations and in quiescent regimes. **It is nevertheless not obvious that a fully reset system would recover its previous behavior as reliably as it does**, and the absence of material accumulation does not by itself establish the absence of accumulation. Candidates for a slower carried quantity would include the evolving current profile, the conditioning of the wall, and resonant structures that persist across events. **This is raised as a question rather than a claim:** within

the framework of this paper it is the question of whether an irreversible variable exists at a level slower than the one at which the cycle is observed, which is precisely the variable Part II, Section 4 requires a candidate system to identify.

Three classes, not two. Taking this section together with Part II, Section 3 and Proposition 5, candidate systems divide, *according to the equation currently used to describe them*, as follows. **Prefactor fluctuation:** noise in the rate, which cannot produce a singularity at any amplitude.

Exponent fluctuation: the mechanism of Part II, Section 3, in which the outcome is decided by the upper tail of the exponent and the mean is irrelevant. **Metastable escape:** a fixed superlinear exponent with a linear damping whose coefficient is slowly driven toward zero, where the outcome is decided by whether a perturbation exceeds a falling threshold. **The three are distinguishable from the governing equations, and that reduction is what assigns a system to one of them.** They are not reliably distinguishable from the shape of a record, and the plasma case shows why: a build-up followed by a collapse is produced by all three. Assigning a candidate to the wrong class on the strength of that shape is the most likely error this framework invites.

What is not claimed. The correspondence established here is between equations, and the plasma satisfies the membership conditions of Part II, Section 4 only within a cycle. **What this section claims is narrower: that the currently established governing equation of the drag-limited plasma eruption reduces to the damped form of this paper's equation with $m = 2$, and that the plasma, as currently modeled, is therefore an instance of metastable escape rather than of intermittent superlinearity.** Nothing here should be read as a claim about what the plasma does, as distinct from what the present models say it does; those models may be superseded, and this paper is itself an attempt to add something to the same body of description. **Where the established treatment is unsettled or disputed within its own field, this paper takes no position**, and the questions raised in the following subsection concern the status of a derived quantity, not the physics from which it was derived.

An observation on prior work. Cross-domain mapping is not new in this field: the same explosive instability has been applied to solar flares, magnetospheric substorms and tokamak disruptions by the same authors, and ELMs have separately been modelled as sandpile and self-organized-criticality systems (Chapman, Dendy & Hnat, *Phys. Rev. Lett.* 86, 2814, 2001). **What appears to be absent is any connection to the general blow-up literature:** the plasma papers derive and use the elementary singularity of $d\xi/dt \propto \xi^2$ without reference to the Osgood criterion, and no link has been drawn to the finite-time-singularity treatments of rupture, or to the accelerating-creep relation of Part I, Section 7, whose inverse-rate forecasting method is the same mathematics applied to a different substance.

1.1 A Caution About the Fixed Exponent

The preceding section takes the exponent of the plasma's nonlinearity as fixed at 2, because that is what the established derivations give. **It is worth being precise about what that establishes and what it does not**, since the classification of the previous section turns on it.

The exponent is derived, not measured. It arises from expanding the energy functional about the marginal state and retaining the leading nonlinear term, which is quadratic. **An expansion of this kind is local by construction:** it holds in a neighbourhood of marginality and for displacements small compared with the width of the unstable region, and the same literature notes that this ordering breaks down before the singularity is reached. What the derivation establishes is that the leading correction near marginality is quadratic. **It does not establish that the effective exponent remains 2 through the phase that produces the observed event.**

The coefficients are not constants, and they are known to vary. Both the linear drive and the coefficient of the nonlinear term depend on the equilibrium; the linear drive varies with the pressure gradient, with sheared flow, and with the turbulence that the previous section describes as shifting the threshold, and the coefficient of the nonlinear term is itself a function of the equilibrium rather than a universal number. **A quantity that depends on the equilibrium is not a constant of the dynamics.**

This matters because a fixed exponent and a varying coefficient are not separate descriptions. For the damped equation of the previous section, the local logarithmic slope of the right-hand side is

$$m_{\text{eff}}(\xi) = \frac{\Gamma_0^2 + 2C_4\xi}{\Gamma_0^2 + C_4\xi}$$

which rises from 1 at small amplitude toward 2 at large amplitude, and which moves whenever the linear drive moves. Whether one says that the exponent is fixed and the coefficient fluctuates, or that the effective exponent fluctuates, is a matter of which variable the equation is written in.

Within the quadratic truncation this slope stays between 1 and 2, so it is not in itself the mechanism of Part II, Section 3; the narrower point stands regardless. **"The exponent does not fluctuate" is a property of how the equation is written, and not an independent fact about the plasma.**

A fluctuating exponent is not excluded by the exponent that governs an eruption. Suppose the effective exponent varied over some range rather than being fixed. By Part II, Section 2, the terminal behaviour of a completing event is set by the value prevailing during the excursion that completes it, and not by the average of the process. **That excursion is necessarily supercritical**, since by the blow-up criterion no subcritical stretch can complete a singularity. The exponent that governs an eruption is therefore a value drawn from above the critical boundary whatever the

mean may be, and a value of 2 there is consistent with a fixed exponent of 2 and equally consistent with a fluctuating one whose mean lies at or below 1 and whose range includes 2. **The two descriptions differ in the mean, and the eruption carries no information about the mean.**

A note on the origin of this caution. The model from which this work began also had a fixed exponent, and it was taken as given for the same reason: it fitted the observations closely, and the functional form in which it was written made no other value natural. **It was only later established that the exponent could not be recovered from the observed trajectory at all** (Part I, Section 7.5), and that the fit was consistent with a whole family of values. The exponent was, in the end, a property of the description rather than of the thing described. **That experience is the reason for this section**, and it is offered as a methodological caution rather than as a criticism of the plasma treatment, which is derived carefully and from first principles. The narrow point is this: **an exponent that enters a model by construction is not thereby a measured feature of the system, and is worth measuring independently.**

2. POSSIBLE CONNECTIONS: ACCELERATING CREEP AND SLOPE FAILURE

The strongest correspondence found so far between the structure of this paper and an established body of physical measurement does not come from any of the domains listed in Part II, Section 1. It comes from geomechanics, and it is worth setting out precisely, because **the equation was written down independently, is in operational use, and its exponent is measured rather than assumed.**

Voight's relation. In the terminal stage of material failure, the accelerating creep of a slope, a rock mass or a loaded specimen is commonly described by

$$\ddot{\Omega} = A \cdot \dot{\Omega}^\alpha$$

where Ω is the observed displacement, the dots denote time derivatives, and A and α are empirical constants (Voight, 1988, 1989; building on Saito and on Fukuzono, 1985). **For $\alpha > 1$ this relation predicts divergence of the velocity at a finite time**, which is read as the transition from accelerating creep to complete failure. Setting $y = \dot{\Omega}/v_{ref}$ and $x = t$ turns it into $dy/dx = A' y^\alpha$, which is **the equation of Part II, Section 2 with the constant term removed and with α in the role of m** . The critical boundary is the same value, 1, and it is crossed for the same reason.

The exponent is measured, and it lies above the boundary. Voight's compilation gives α in the range 1.74–2.01 for metals, 1.90–2.11 for soils, and close to 2 for natural landslides; a more recent compilation spans roughly 1.53–2.15, with the prefactor A varying across three orders of magnitude. This is the first candidate encountered in this work where **the supercritical exponent is an outcome of measurement rather than an assumption of the model.**

The inverse-velocity method is the $\alpha = 2$ case of Proposition 3. Part II, Section 2 gives the distance still required to reach the singularity from a level y as $T(y) \approx y^{1-m}/[k(m-1)]$. At $m = 2$ this is simply proportional to $1/y$, so the remaining time is a linear function of the inverse velocity. This is exactly the standard forecasting procedure in the field: plot $1/\dot{\Omega}$ against time and extrapolate to zero. Applied retrospectively to the Vaiont slide, Voight noted that the failure date could have been identified more than ten days in advance by this route. **What is a forecasting heuristic there is, in the present framework, a corollary of the blow-up integral.**

A mathematical addition: the constant term restores the earlier stage. Voight's relation describes only the accelerating, tertiary stage; it says nothing about the steady, secondary creep that precedes it, and it cannot, since at $\dot{\Omega} = 0$ it predicts no motion at all. The equation of Part II, Section 2, written for this case as

$$\ddot{\Omega} = k [1 + (\dot{\Omega}/v_{ref})^\alpha]$$

contains both stages in one expression. While $\dot{\Omega} \ll v_{ref}$ the constant term dominates and the creep is steady; once the velocity passes the reference scale the feedback term takes over and Voight's relation is recovered as the asymptotic form. The crossover velocity v_{ref} then becomes a measurable quantity in its own right, and the blow-up time follows from the closed form of Proposition 1 rather than from a fit restricted to the final phase. Whether this single-expression form fits creep data better than the two-stage description is an open, and directly testable, question.

An observation about the measured values. Integrating the local scaling law of Part II, Section 2 shows that the velocity diverges at t_f for every $\alpha > 1$, but the *displacement* reached at that moment is finite only for $\alpha > 2$, and diverges for $\alpha < 2$. **The measured exponents cluster at almost exactly this dividing value.** No explanation is offered here, and the clustering may well be an artifact of fitting procedures that assume $\alpha = 2$ in the first place. It is recorded because it is the kind of coincidence that is either meaningless or important, and the distinction is decidable from data.

Aborted accelerations. The field also records precursory accelerations that do not complete: the 1986–87 acceleration of the La Clapière landslide is a documented case. Such episodes are normally treated as false alarms, or as failures of the forecasting method. **In the framework of Part II, Section 3 they are neither: they are excursions above the critical boundary that ended before the required duration had elapsed.** If that reading is correct, the archives of this field already contain what is hardest to obtain elsewhere, namely a record of both completed and uncompleted excursions in the same system, which is precisely what the waiting-time predictions of Part II, Section 3 require in order to be tested.

What is not claimed. The correspondence above is between equations, not mechanisms, and the membership conditions of Part II, Section 4 are not satisfied at present. Two gaps are specific enough to state. First, **the non-decreasing state condition fails for the obvious choice of variable:** the velocity of a creeping slope can and does fall again, as the aborted accelerations themselves demonstrate. If this system belongs to the class at all, the ratcheting variable must be the irreversible one, the accumulated damage or the state of the crack network, with the velocity as its observable consequence; identifying that variable is a prerequisite, not a detail. Second, **it is not established which physical quantity fluctuates the exponent.** The same filter applies here as everywhere in this paper: rainfall and pore-water pressure act principally on the prefactor A , and prefactor fluctuations cannot produce a singularity; it is damage accumulation and the evolving crack network that would have to move α itself. Whether they do so, and whether they carry it across the boundary intermittently, is unresolved.

Why the section is included. Voight's relation is phenomenological: it is fitted, widely used, and not derived from first principles. A framework that specifies where such an exponent comes from, why it sits above 1, and what governs its variation would be addressing an open question in that field rather than importing a solved one. **That is the reason this direction is recorded here, and it is recorded as a correspondence between equations and nothing more.**

REFERENCES

- Wilson, H. R. & Cowley, S. C. (2004). *Theory for Explosive Ideal Magnetohydrodynamic Instabilities in Plasmas*. Physical Review Letters, 92(17), 175006. DOI: 10.1103/PhysRevLett.92.175006.
 - Cowley, S. C. & Artun, M. (1997). *Explosive instabilities and detonation in magnetohydrodynamics*. Physics Reports, 283(1–4), 185–211. DOI: 10.1016/S0370-1573(96)00060-0.
 - Cowley, S. C., Cowley, B., Henneberg, S. A. & Wilson, H. R. (2015). *Explosive instability and erupting flux tubes in a magnetized plasma*. Proceedings of the Royal Society A, 471(2180), 20140913. DOI: 10.1098/rspa.2014.0913.
 - Ham, C. J., Cowley, S. C., Brochard, G. & Wilson, H. R. (2016). *Nonlinear Stability and Saturation of Ballooning Modes in Tokamaks*. Physical Review Letters, 116(23), 235001. DOI: 10.1103/PhysRevLett.116.235001.
 - Biglari, H., Diamond, P. H. & Terry, P. W. (1990). Influence of sheared poloidal rotation on edge turbulence. Physics of Fluids B: Plasma Physics, 2(1), 1–4. DOI: 10.1063/1.859529.
 - Lang, P. T. et al. (2004). *ELM pace making and mitigation by pellet injection in ASDEX Upgrade*. Nuclear Fusion, 44(5), 665–677. DOI: 10.1088/0029-5515/44/5/010.
 - Webster, A. J. & Dendy, R. O. (2013). *Statistical Characterization and Classification of Edge-Localized Plasma Instabilities*. Physical Review Letters, 110(15), 155004. DOI: 10.1103/PhysRevLett.110.155004.
 - Weibull, W. (1951). *A statistical distribution function of wide applicability*. Journal of Applied Mechanics, 18(3), 293–297. DOI: 10.1115/1.4010337.
 - Voight, B. (1988). *A method for prediction of volcanic eruptions*. Nature, 332(6160), 125–130. DOI: 10.1038/332125a0.
 - Voight, B. (1989). *A relation to describe rate-dependent material failure*. Science, 243(4888), 200–203. DOI: 10.1126/science.243.4888.200.
 - Fukuzono, T. (1985). *A new method for predicting the failure time of a slope*. Proceedings of the 4th International Conference and Field Workshop on Landslides, 145–150.
 - Saito, M. (1965). *Forecasting the time of occurrence of a slope failure*. Proceedings of the 6th International Conference on Soil Mechanics and Foundation Engineering, 2, 537–541.
 - Helmstetter, A. et al. (2004). *Slider block friction model for landslides: Application to Vaiont and La Clapière landslides*. Journal of Geophysical Research: Solid Earth, 109, B02409. DOI: 10.1029/2002JB002160.
 - Chang, C. & Wang, G. (2022). *Creep of clayey soil induced by elevated pore-water pressure: Implication for forecasting the time of failure of rainfall-triggered landslides*. Engineering Geology, 296, 106461. DOI: 10.1016/j.enggeo.2021.106461.
 - Segalini, A., Valletta, A. & Carri, A. (2018). *Landslide time-of-failure forecast and alert threshold assessment: A generalized criterion*. Engineering Geology, 245, 72–80. DOI: 10.1016/j.enggeo.2018.08.003.
-