Seeing Crises Before They Hit

Jon Smirl

2026-03-01

A Marble in a Valley

Imagine the economy as a marble sitting at the bottom of a curved valley. When something bumps it — an oil price shock, a supply chain disruption, a shift in consumer confidence — the marble rolls up the side and then settles back to the bottom. How quickly it returns depends on how steep the valley walls are. A deep, narrow valley pulls the marble back fast. A shallow, flat valley lets it wander.

This is the CES potential landscape from *The Economic Landscape*. The valley's shape is controlled by the effective curvature $K_{\text{eff}} = K(1 - T/T^*)^+$. When information friction T is low, the valley is deep and recovery is quick. As T rises toward T^*, the valley flattens. The marble takes longer to come back. At $T = T^*$, the valley disappears and the marble rolls away — that is the crisis.

The remarkable fact is that this flattening is **measurable** before the crisis arrives. You do not need to wait for the marble to roll away. You can detect the valley getting shallower by watching how the marble behaves after small bumps. This is the basis of early warning signals for economic crises.

One Formula, Four Signals

Near the bottom of the valley, the economy's return to equilibrium after a small shock follows an exponential decay with a characteristic timescale:

> **Definition (Relaxation Time).**
>
> $$\tau(T) = \frac{\tau_0}{1 - T/T^*}$$
>
> where τ_0 is the baseline relaxation time when $T = 0$, and T^* is the critical information friction at which the valley vanishes.

As T approaches T^*, the relaxation time diverges — the economy takes longer and longer to recover from shocks. This single formula generates four observable consequences, all derived from the *static_vri*:

1. Autocorrelation rises. When the relaxation time τ is short, today's deviation from equilibrium is quickly corrected, and tomorrow's value is nearly independent of today's. When τ is long, today's deviation persists into tomorrow, and the day after, and the day after that. The lag-1 autocorrelation of any economic time series scales as $\exp(-\Delta t/\tau)$, which approaches 1 as $\tau \to \infty$. In plain terms: the economy starts repeating itself, stuck in a rut it cannot escape.

2. Variance increases. The economy is constantly buffeted by small random shocks — weather, policy surprises, demand fluctuations. In a steep valley, these shocks are quickly absorbed and the marble stays near the bottom. In a shallow valley, each shock pushes the marble further from center before it comes back (if it comes back at all). The variance of economic indicators scales as τ, which diverges as $T \to T^*$. Fluctuations get bigger.

3. Recovery time lengthens. This is the most direct manifestation: after a recession, how many quarters does it take GDP to return to trend? After a financial shock, how long before credit spreads normalize? The recovery time *is* τ, and it grows without bound as the landscape flattens.

4. Cross-sector correlation rises. In a steep valley, different sectors are buffeted by different shocks and behave relatively independently. As the valley flattens, all sectors begin responding to the *same* slow mode. Sectors that were previously independent start moving together. Correlation rises toward 1.

> **Theorem (Pre-Crisis Deceleration).**
>
> All four indicators — autocorrelation, variance, recovery time, and cross-sector correlation — diverge as $T \to T^*$ with the same critical exponent. They are not four separate phenomena. They are four measurements of the same underlying fact: the CES potential landscape is flattening.

Why This Is Not Hindsight

Every crisis looks predictable in retrospect. What makes the CES framework different is that these signals are **leading indicators** — they change *before* the crisis, not during or after it.

The valley does not flatten instantaneously. Information friction T rises gradually as complexity increases, trust erodes, or regulatory gaps widen (see *The Economics of Not Knowing*). As T creeps toward T^*, the relaxation time τ grows, and all four signals begin moving months before the valley actually disappears. Two empirical tests give concrete numbers.

Evidence: CPI Sectoral Dispersion

The first test uses CPI (Consumer Price Index) data disaggregated by sector. The theory predicts that as the economy approaches a turning point, sectoral price dispersion should rise — this is the variance signal applied to prices. Rising dispersion means sectors are being pushed apart by shocks that the economy can no longer absorb quickly.

> **Example.**
>
> The *test:cpi-leading-indicator* examines all business cycle turning points in the U.S. CPI data. The result: **18 out of 18 turning points** were preceded by rising sectoral dispersion. The median lead time was **12 months**. The correlation between dispersion and subsequent turning points is $r = +0.37$.
>
> Eighteen for eighteen is not a statistical fluke. It is consistent with the theory's prediction that dispersion must rise before a crisis, because dispersion is a direct measure of the lengthening relaxation time τ.

This test is particularly clean because CPI data is published monthly, in real time, at the sectoral level. There is no look-ahead bias. An analyst computing sectoral CPI dispersion in January 2007 would have seen it rising and could have flagged elevated risk, a full year before the September 2008 collapse.

Evidence: Semiconductor Dispersion

The second test uses a different sector at a different frequency. The WSTS (World Semiconductor Trade Statistics) data reports semiconductor sales across multiple product segments — memory, logic, analog, optoelectronics, and so on. The theory predicts that dispersion across these segments should lead semiconductor downturns.

> **Example.**
>
> The *test:dispersion-indicator-wsts* confirms this prediction: semiconductor segment dispersion leads industry downturns by approximately **3 quarters**. When semiconductor segments that normally move independently begin diverging in their growth rates, a downturn follows.

The semiconductor test operates at a completely different level of the economic hierarchy than CPI — consumer prices across broad sectors versus production within a single industry. The fact that the same early warning pattern appears at both levels is consistent with the hierarchical CES framework (see *The Economy Has Layers*): the same potential landscape and relaxation dynamics apply at every level, just on different timescales.

How to Use This in Practice

The recipe is straightforward: choose disaggregated indicators at the level of interest (sectoral CPI for macro crises, segment-level production for industry crises), compute both cross-sectional dispersion and lag-1 autocorrelation at regular intervals, and watch for *simultaneous* rises in both. A rise in dispersion alone might reflect a healthy structural shift. A rise in autocorrelation alone might reflect a persistent but benign shock. When both rise together, the CES framework says the landscape is flattening.

The signals do not tell you *what* will go wrong — only that something is about to. The specific trigger is unpredictable. But the vulnerability is measurable. As (Reinhart2009) documented across eight centuries of financial history, crises are rarely truly surprising in hindsight. The CES framework explains why: the landscape was already flat, and any sufficiently large bump would have sent the marble over the edge.

The Connection to Curvature

Why do these early warning signals relate to the curvature parameter K from *Emergent CES*? Because the depth of the valley is proportional to K_{eff}, and the *test:early-warning-by-k* confirms that sectors with higher structural curvature K (stronger complementarity) show stronger early warning signals. This makes sense: a deeply curved valley that flattens provides a more dramatic change in dynamics than a valley that was shallow to begin with.

The general lesson from (Scheffer2009) — that complex systems approaching tipping points exhibit characteristic deceleration — is not merely an analogy. In the CES framework, it is a theorem. The variance-response identity connects the curvature of the potential landscape to the statistics of observable fluctuations, turning a qualitative intuition into a quantitative measurement protocol.

The economy tells you when it is becoming fragile. The hard part is not detecting the signal — it is acting on it.

References