Hidden Cycles in a Century of Data

Jon Smirl

2026-03-01

Five Names, One Century of Argument

In the 1920s and 1930s, economists proposed that capitalist economies move in waves. (Kitchin1923) identified short inventory cycles of roughly 3 years. (Juglar1862) had already documented investment cycles of 7–11 years. (Kuznets1930) pointed to building cycles of 15–25 years. (Kondratiev1925) claimed to see long waves of 40–60 years spanning entire technological eras. (Schumpeter1939) tried to synthesize all of them into a grand theory of capitalist evolution.

Are these cycles real, or artifacts — patterns the human eye finds in noisy data? When (Burns1946) built the empirical foundations of modern business cycle research at the NBER, they deliberately avoided committing to any fixed set of cycle lengths, precisely because the evidence was so contested.

A century later, we can settle this. The tool that settles it is called Empirical Mode Decomposition, and it lets the data speak for itself.

A Sieve for Signals

Imagine pouring a mixture of sand, gravel, and pebbles through a series of sieves, each with a finer mesh than the last. The finest sieve catches the sand. The next catches the gravel. The coarsest catches the pebbles. You never told the sieves what size to look for — the separation happened because of the physical differences in particle size.

Empirical Mode Decomposition, or EMD, does the same thing to a time series (Huang1998). It takes a signal — in this case, the Federal Reserve's index of US industrial production from 1919 to 2024 — and separates it into components called Intrinsic Mode Functions (IMFs). Each IMF captures oscillations at a characteristic timescale. The fastest oscillations come out first. Then progressively slower ones. Finally, a long-run trend remains.

The critical feature of EMD is that it is **adaptive**. Unlike Fourier analysis, which decomposes any signal into fixed sine waves regardless of whether the data actually contains them, EMD extracts whatever oscillations are genuinely present. It handles non-stationary data and nonlinear oscillations. You do not tell EMD what frequencies to look for. It finds them.

This matters enormously for the cycle debate. Fourier analysis always returns power at every frequency, making "real cycle or not?" a matter of arguable significance thresholds. EMD sidesteps this entirely. If a cycle is there, it appears as a distinct IMF. If it is not there, no IMF emerges at that frequency.

What the Data Shows

Applied to 105 years of monthly US industrial production data, EMD extracts seven IMFs. The first two (IMF1 and IMF2) capture high-frequency noise — fluctuations shorter than about a year that reflect monthly volatility and seasonal residuals. The last (IMF7) is a very slow oscillation at the boundary of what a century of data can resolve. The five in between are where the action is:

IMF	Period (years)	Classical Cycle
IMF3	2.46	Kitchin inventory cycle
IMF4	5.56	NBER business cycle
IMF5	~11	Juglar investment cycle
IMF6	~22	Kuznets building cycle
IMF7	~50	Kondratiev long wave

Every one of the named cycles that economists have debated for a century shows up as a distinct component. They are not artifacts of a particular decomposition method or a particular sample window. When the analysis is repeated across 18 different rolling 20-year subsamples, the same five timescales appear with stable periods. The cycles are robust.

But the most striking finding is not the individual cycles. It is the relationship between them.

A Geometric Ladder

Divide each cycle period by the one below it:

$$r_1 = \frac{5.56}{2.46} = 2.26, \quad r_2 = \frac{11}{5.56} \approx 1.98, \quad r_3 = \frac{22}{11} \approx 2.0, \quad r_4 = \frac{50}{22} \approx 2.27$$

The median ratio is $r^* \approx 2.19$. The ratios are remarkably close to each other, spanning from about 2.0 to 2.3. This means the five cycle periods form an approximate geometric progression:

$$T_n \approx T_0 \cdot (r^*)^n$$

Each cycle is roughly 2.2 times longer than the one below it. This is not something anyone assumed or imposed. It fell out of a purely data-driven decomposition. The economy has organized itself into layers, and those layers are evenly spaced on a logarithmic scale.

This geometric spacing is what makes the finding theoretically consequential, not just empirically interesting.

Why a Geometric Ladder?

The CES hierarchy framework (see *emd-timescale*) predicts that a multi-level economic system — one where fast processes (inventory management) are nested inside slower processes (investment planning), which are nested inside still slower processes (infrastructure building) — will naturally separate into layers with a characteristic timescale ratio.

The theoretical prediction is that the number of effective layers should be $N_{\text{eff}} \approx 4.5 \pm 1.0$, and the timescale ratio between adjacent layers should be determined by the curvature parameter K of the

CES aggregation at each level. The empirical finding of $N_{\text{eff}} = 5$ layers with $r^* = 2.19$ falls squarely within this range.

> **Theorem (Hierarchical Timescale Ratio).**
>
> In a CES hierarchy with timescale separation, adjacent levels maintain a characteristic period ratio r^* determined by the curvature of the CES potential at each level. The ratio is approximately constant across levels when the structural parameters are similar.

The intuition is straightforward. Each level of the economy acts as a filter for the level above it. Inventory managers respond to monthly demand, but their aggregate behavior over years drives investment decisions. Investment accumulated over a decade shapes building cycles. The CES aggregation at each level determines how much smoothing occurs — how much fast variation is absorbed before passing through. When the aggregation structure is similar across levels, the smoothing factor is similar, and you get a geometric progression.

What This Settles

The EMD results resolve three long-standing questions.

First, the named business cycles are real. They are not statistical artifacts or confirmation bias. They appear in a method with no prior knowledge of economic theory and no reason to produce five components rather than three or twelve.

Second, the cycles are not independent phenomena requiring separate explanations. The geometric ladder says they are manifestations of a single hierarchical architecture. You need one theory for how economic layers aggregate, and the cycles follow.

Third, the economy has a characteristic depth. Five effective layers, not two (micro/macro) and not fifty. A model with fewer layers will miss dynamics. A model with many more will overfit.

The *test:emd-hierarchy-timescale* provides the full statistical details: rolling-window stability tests, confidence intervals on the period estimates, and comparisons with alternative decomposition methods (wavelets, bandpass filters) that yield consistent results.

From Observation to Architecture

A century of data, filtered through an algorithm that imposes nothing, reveals five layers and a geometric ladder. This is an empirical fact. The CES hierarchy framework provides a theoretical explanation: the layers arise because economic aggregation has a characteristic curvature, and that curvature sets the timescale ratio between adjacent levels.

The economy is not a single system oscillating at one frequency, nor a collection of unrelated oscillations. It is a **layered architecture** where each level operates on its own timescale, constrained from below by slower levels and driven from above by faster ones. The geometric ladder is the signature of that architecture, visible once you use a tool that does not impose its own structure on what it finds.

Kitchin, Juglar, Kuznets, and Kondratiev were not wrong. They were each looking at one layer of a five-layer system. The ladder was always there. It just took a century, and the right sieve, to see it clearly.

References