How Long Do Transitions Take?

Jon Smirl

2026-03-01

The Question Everyone Asks

When a new technology is clearly better than the old one, why does the transition take so long? Electricity was demonstrably superior to steam power by the 1890s, yet factories did not fully convert until the 1920s. Personal computers outperformed mainframes for most business tasks by the early 1980s, but mainframe-centric IT departments persisted into the mid-1990s. In each case, the economy seemed stuck in the old arrangement well past the point where the new one was superior.

This is not stubbornness or irrationality. The CES framework provides a precise mathematical answer: the duration of a transition depends on exactly two parameters — the **geometric dimension** of the technology and the **learning elasticity** of the industry — and the economy can predictably remain in the old equilibrium even after the new one becomes dominant. The delay is not a bug. It is a measurable feature of how learning curves interact with the structure of production.

Learning Curves: The Engine of Transition

Every transition begins with a learning curve. As (Wright1936) first documented for airplane manufacturing, unit costs decline as a power law of cumulative production:

$$C(Q) = C_0 \cdot Q^{-\alpha}$$

where $C(Q)$ is the unit cost after Q cumulative units have been produced, C_0 is the cost of the first unit, and α is the **learning elasticity** — how rapidly costs fall with experience. A higher α means faster learning.

(Arrow1962) formalized this as "learning by doing": the act of producing something teaches you how to produce it more cheaply. This insight is now standard in economics. What the CES framework adds is the recognition that α is not a free parameter. It is determined by the physical structure of the technology.

Geometric Dimension: Why Some Technologies Learn Faster

The key formula is:

Definition (Geometric Learning Rate).

$$\alpha_n = d_n \times \alpha_0$$

where d_n is the geometric dimension of the production process and $\alpha_0 \approx 0.12$ is the baseline learning elasticity common across industries.

What is geometric dimension? It counts the number of independent directions along which the production process can be improved simultaneously. Consider two examples:

A wire or cable ($d = 1$). Improvements happen along one dimension — you make the process longer, faster, or more precise, but fundamentally the optimization is one-dimensional. The learning elasticity is $\alpha \approx 0.12$.

A semiconductor chip ($d = 2$). Improvements happen across a two-dimensional surface. Transistors are etched onto a flat wafer, and each generation shrinks features in both the x and y directions simultaneously. Every percentage improvement in lithographic precision yields gains across the entire planar area. The learning elasticity is $\alpha \approx 2 \times 0.12 = 0.24$.

This is not a metaphor. Semiconductor fabrication literally operates on planar geometry, and the doubling of learning rate relative to one-dimensional processes is an observable, measurable fact. Moore's Law — the observation that chip density doubles roughly every two years — is a consequence of $d = 2$ geometry applied to the Wright learning curve.

How Duration Depends on Learning Rate

The CES framework connects learning rates to transition duration through the concept of a **drift rate**. As the new technology accumulates production and rides down its learning curve, it gradually approaches cost parity with the incumbent. The drift rate $\varepsilon_{\text{drift}}$ measures how quickly the cost gap is closing:

$$\varepsilon_{\text{drift}} \propto d \times \alpha_0$$

The transition duration then scales as:

Theorem (Transition Duration Scaling).

$$\text{Duration} = O\left(\frac{1}{\sqrt{\varepsilon_{\text{drift}}}}\right)$$

The time an economy spends transitioning from an old equilibrium to a new one is inversely proportional to the square root of the drift rate.

This square-root scaling has an important consequence: doubling the learning rate does not halve the transition time. It reduces it by a factor of $\sqrt{2} \approx 1.41$. Learning faster helps, but with diminishing returns.

Plugging in the numbers:

Technology type	d	α	Predicted duration
Generic 1D industry	1	0.12	~11 years
Semiconductor (planar)	2	0.24	~8 years
Energy (intermediate)	1–2	0.12–0.24	~8–11 years

The Delayed Transition: Staying Past the Crossover

Perhaps the most counterintuitive prediction is that the economy does not switch to the new technology the instant it becomes cheaper. Instead, there is a **delayed transition** — a predictable period during which the old equilibrium persists even though the new one is already superior.

Think of it this way. Imagine a ball sitting in a shallow valley. A second, deeper valley is forming nearby. Even after the new valley becomes deeper than the old one, the ball stays put — it needs a push, or the ridge between the two valleys needs to erode away, before it rolls over. The delay between "the new valley is deeper" and "the ball actually moves" is not random. It depends on how fast the landscape is changing, which is exactly $\varepsilon_{\text{drift}}$.

This is why economic transitions feel frustratingly slow to technologists and frustratingly fast to incumbents. The technologist sees cost curves that crossed years ago and wonders why adoption has not happened. The incumbent sees a stable market that suddenly collapses. Both are observing the same delayed transition from different vantage points.

Three Historical Examples

Steam to electricity (~20 years, 1900–1920). (David1990) documented the "productivity paradox" of electrification: factories adopted electric motors but kept the old layout designed for centralized steam-driven shafts. Productivity gains arrived only when factories were redesigned around distributed electric power. This is a $d \approx 1$ process (linear manufacturing), predicting $\alpha \approx 0.12$ and a transition on the order of 11–20 years. The actual duration of roughly 20 years sits at the upper end, consistent with the additional organizational restructuring required.

Mainframes to PCs (~12 years, 1981–1993). The IBM PC launched in 1981, and by the early 1990s client-server architectures had displaced mainframe-centric computing for most business applications. Semiconductor learning rates ($d = 2$, $\alpha \approx 0.24$) drove the cost decline of microprocessors, predicting transitions around 8 years. The actual 12 years reflects the fact that the transition involved not just hardware costs but also software ecosystems and organizational practices — slowing the effective drift rate below the pure hardware prediction.

Centralized to distributed AI (~8 years, CES prediction). Current AI infrastructure is dominated by centralized data centers with massive GPU clusters. The CES framework predicts that as inference hardware rides down the semiconductor learning curve ($d = 2$, $\alpha \approx 0.23$), cost parity for distributed inference will arrive and a transition to mesh-style distributed AI will follow. The predicted duration: approximately 8 years from the crossing point, with a delayed transition period during which centralized providers maintain market share even after distributed alternatives are cost-competitive.

Why This Matters

The practical value of *transition_duration_scaling* is that it replaces hand-waving about "disruption timelines" with a formula. If you know the geometric dimension of the relevant technology and can estimate the baseline learning elasticity, you can bound the transition duration.

For policymakers, this means that interventions to accelerate transitions should target the drift rate $\varepsilon_{\text{drift}}$. Subsidizing early production to steepen the learning curve (increasing effective α) will shorten the transition, but only as the square root — massive subsidies produce modest acceleration.

Reducing switching costs, by contrast, can shrink the delayed transition window directly by lowering the ridge between the old and new valleys.

For investors and strategists, the framework says: do not expect the crossing point to coincide with adoption. Budget for the delay. Semiconductor-driven transitions ($d = 2$) will take roughly 8 years from crossing to completion. Transitions driven by one-dimensional learning processes will take roughly 11 years. And transitions that require organizational or institutional restructuring on top of the technology shift — as electrification and mainframe displacement both did — will take longer still.

The economy does not jump. It climbs, predictably, at a pace set by the geometry of learning.

References