Learning Curves: Why Costs Always Fall

Jon Smirl

2026-03-01

A Pattern That Keeps Showing Up

In 1936, Theodore Wright noticed something remarkable about airplane manufacturing. Every time the cumulative number of airframes produced doubled, the labor cost per unit fell by a consistent percentage — about 20 percent. This was not a one-time efficiency gain. It was a regularity: the hundredth plane was cheaper than the tenth, and the thousandth was cheaper still, all following the same predictable curve (Wright1936).

Wright was studying airplanes, but the pattern turned out to be nearly universal. Over the following decades, economists and engineers found the same relationship in shipbuilding, steel production, chemical processing, and consumer electronics. (Arrow1962) formalized the idea as "learning by doing" — the act of producing something teaches you to produce it more cheaply. (Nagy2013) tested it across 62 technologies spanning more than a century and confirmed that the power-law relationship holds with striking consistency.

This is Wright's Law, and it is one of the most reliable regularities in all of economics.

The Formula

The learning curve takes a simple mathematical form:

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

Here $C(Q)$ 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** — the rate at which costs decline with experience.

The learning elasticity α determines the **progress ratio**: every doubling of cumulative output reduces cost by a factor of $2^{-\alpha}$. For a typical value of $\alpha = 0.12$, this means each doubling cuts cost by about 8 percent. For $\alpha = 0.23$, each doubling cuts cost by about 15 percent. The difference between these two numbers may sound small, but compounded over decades and many doublings, it is enormous.

Three Examples

Ford Model T. Henry Ford's assembly line is perhaps the most famous example of learning-by-doing in action. The Model T launched at $850 in 1908. By 1925, after cumulative production had grown roughly 100-fold, the price had fallen to $260. Plotted on log-log axes — log cost against log cumulative production — the data points fall along a straight line, exactly as Wright's Law predicts.

Solar panels. In 1977, a watt of solar photovoltaic capacity cost about \$76. By 2023, it cost roughly \$0.20 — a decline of more than 99.7 percent. This 375-fold cost reduction tracks Wright's Law almost exactly, with $\alpha \approx 0.16$. Solar costs did not fall because of any single breakthrough. They fell because of relentless incremental improvement driven by cumulative production experience — the same force Wright identified in airplane factories.

DRAM memory. A gigabyte of dynamic random-access memory cost approximately \$1,000,000 in 1970. By 2023, the same gigabyte cost about \$0.002. That is a reduction by a factor of 500 million, achieved over roughly 25 doublings of cumulative production. The learning elasticity for DRAM is approximately $\alpha \approx 0.23$, making it one of the fastest-learning technologies ever measured.

Why Some Technologies Learn Faster

Here is a question that conventional treatments of Wright's Law leave unanswered: why is $\alpha \approx 0.23$ for semiconductors but only $\alpha \approx 0.12$ for most other industries? The standard answer — "it depends on the technology" — is not very satisfying. The CES framework provides a structural explanation.

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.

The geometric dimension d_n counts how many independent spatial directions the production process operates on simultaneously. Consider two contrasting cases.

A cable or wire ($d = 1$). The manufacturing process is fundamentally one-dimensional. You extrude material along a single axis. Improvements happen by making the process faster, more precise, or more efficient, but the optimization occurs along one dimension. The learning elasticity is $\alpha \approx 1 \times 0.12 = 0.12$.

A semiconductor chip ($d = 2$). Transistors are etched onto a flat wafer — a two-dimensional surface. When lithographic precision improves, the gains compound across both the x and y directions. A 10 percent improvement in resolution yields roughly 20 percent more transistors per unit area because the improvement applies in two independent dimensions. The learning elasticity is $\alpha \approx 2 \times 0.12 = 0.24$.

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

Why Learning Curves Determine Transition Timing

Learning curves are not just a curiosity of manufacturing economics. They are the clock that governs major economic transitions.

Consider the question: when will distributed AI hardware become cheaper than renting compute from centralized cloud providers? Or: when will solar power become cheaper than coal in a given region? In both cases, the answer is the same — follow the learning curve.

The incumbent technology typically has a flat cost curve. Coal plants built decades ago have already ridden down their learning curve; their costs are roughly constant. The challenger technology — solar, distributed AI chips, lithium batteries — is still in the steep part of its learning curve, with costs falling every year as cumulative production grows.

The **crossing point** is the moment when the challenger's declining cost curve intersects the incumbent's flat cost. Before the crossing, the incumbent is cheaper. After it, the challenger is. The date of the crossing is determined almost entirely by α and the rate of cumulative production growth.

For a high-α technology like semiconductors ($\alpha \approx 0.23$), each doubling of production delivers a 15 percent cost reduction, and the crossing arrives relatively quickly. For a generic $\alpha \approx 0.12$ technology, the same doubling delivers only 8 percent, and the crossing takes roughly twice as many doublings — which, at a given production growth rate, means roughly twice as many years.

This is why the geometric dimension correction matters so much. It is not an academic refinement. It determines whether a transition takes a decade or a generation. *wright-law*

The Paradox of Investment

Learning curves create a deep strategic paradox that the CES framework makes precise. Because costs fall with cumulative production, the technology that is produced the most today will be cheapest tomorrow. This means that large-scale investment in any technology — even a centralized one — accelerates the learning curve that will eventually make decentralized alternatives viable.

When a hyperscaler like Google or Microsoft spends billions on AI chips, that spending finances the semiconductor learning curve. The chips get cheaper. But cheaper chips do not only benefit hyperscalers — they benefit everyone, including the distributed competitors who will eventually undercut centralized pricing. The concentrated investment finances the very learning curve that undermines the investor's market position.

(Nordhaus2021) calls this the broader phenomenon of technology-driven cost decline reshaping economic structure. Within the CES framework, it is the mechanism by which Level 1 (semiconductor costs, riding down the learning curve over decades) determines the timing of transitions at Level 2 (market structure) and beyond.

What To Remember

Wright's Law is simple: $C(Q) = C_0 \cdot Q^{-\alpha}$. Costs fall as a power law of cumulative production. The learning elasticity α is the single most important parameter for forecasting technology costs.

The CES contribution is the geometric dimension correction: $\alpha_n = d_n \times \alpha_0$. Technologies that operate on surfaces ($d = 2$) learn twice as fast as technologies that operate on lines ($d = 1$). This is why semiconductors have outpaced nearly every other manufacturing learning curve for half a century, and why Moore's Law held as long as it did.

If you want to know when a transition will happen, measure α. If you want to know why α takes the value it does, measure d. Everything else is detail.

References