The Boundaries of the Firm: When to Make vs. Buy

Jon Smirl

2026-03-01

Why Does Apple Design Its Own Chips?

Apple sells over 200 million iPhones a year. Each one contains an A-series or M-series chip that Apple designed in-house, fabricated by TSMC under contract, assembled by Foxconn, and shipped in a box with screws Apple bought from commodity suppliers.

Why this particular division of labor? Apple could buy chips from Qualcomm (it used to). TSMC could design its own chips (it does not). Foxconn could make its own screws (it would never bother). These are not arbitrary choices. They follow a pattern, and the pattern has a precise mathematical explanation.

The question goes back to Ronald (Coase1937): why do firms exist at all? If markets are efficient, every transaction could happen between independent parties. The fact that firms bundle many activities under one roof means something about markets is costly. (Williamson1975) identified the key factor: when inputs are **relationship-specific** — tailored to a particular partner — market transactions expose you to hold-up. Your supplier knows you cannot easily switch, and prices accordingly.

But Williamson's framework is qualitative. It tells you *that* relationship specificity matters, not *how much* it matters or *where* the boundary should fall. The CES framework fills this gap with two numbers.

Two Parameters, One Boundary

Recall from *the CES potential* that every economic activity can be located in a two-dimensional space defined by:

- ρ (the substitution parameter): How easily can you replace one input with another? When ρ is strongly negative, inputs are tight complements — each one is essential. When ρ is close to 1, inputs are near-substitutes — any supplier will do.
- T (information friction): How costly is it to evaluate, monitor, and coordinate across an organizational boundary? Low T means transparent relationships. High T means fog — you cannot easily tell whether your supplier is delivering quality.

The *effective curvature* theorem from *The Economics of Not Knowing* tells us that the exploitable complementarity is:

$$K_{\text{eff}} = K \cdot \left(1 - \frac{T}{T^*}\right)^+$$

where $K = (1-\rho)(J-1)/J$ is the raw curvature (Arrow1961). When K_{eff} is high, there is a large premium from coordinating inputs tightly. When K_{eff} is low, inputs can be managed at arm's length.

This directly answers the make-or-buy question:

Definition (Governance Boundary).

A firm should **integrate** (make internally) an activity when its effective curvature K_{eff} under market governance is significantly lower than under internal governance — that is, when the information friction of the market (T_{market}) erodes the complementarity premium enough to justify the overhead of doing it in-house ($T_{\text{internal}} < T_{\text{market}}$, but with bureaucratic costs).

The governance boundary is the curve in (ρ, T) space where the net benefit of integration equals zero. Below the curve, the market works fine. Above the curve, bring it inside.

The Logic in Plain English

Here is the intuition, step by step.

Strong complements + high market friction = integrate. If your inputs must work together precisely (low ρ), and the market makes it hard to verify that an outside supplier is delivering the right quality (high T_{market}), then the hold-up problem is severe. A supplier who knows you are locked in will exploit that. Worse, you may not even be able to *detect* whether the supplier is cutting corners until it is too late. The combination of tight complementarity and opaque market conditions makes internal production the safer choice.

Near-substitutes + low market friction = buy. If you can easily switch suppliers (high ρ), or if market information is cheap and reliable (low T_{market}), there is no reason to bear the overhead of internal production. Let the market compete for your business.

The middle ground = hybrid governance. Long-term contracts, joint ventures, and strategic partnerships occupy the intermediate region. These arrangements reduce T below the pure market level (through relationship-building, shared information systems, contractual safeguards) without incurring the full bureaucratic cost of integration.

This is exactly what (Grossman1986) and (Klein1978) described in the property-rights theory of the firm — but now with a quantitative framework rather than case-by-case reasoning.

The Semiconductor Test Case

The semiconductor industry provides an unusually clean test because the production chain has well-defined stages with very different complementarity structures. Let us walk through each stage and see where it sits in (ρ, T) space.

Example.

Chip design ($\rho \approx -3$, low T within the firm). Circuit design is intensely complementary. The logic architecture, the memory subsystem, the power management, and the I/O interfaces must all work together at the transistor level. A mistake in one block can invalidate the entire chip. This is deep complementarity — ρ is strongly negative. Within a design team, information flows freely (low T), but across firm boundaries, sharing detailed circuit layouts creates IP risk and coordination nightmares (high T_{market}). **Prediction: integrate.** Reality: Apple, Qualcomm, NVIDIA, and AMD all design in-house.

Fabrication ($\rho \approx -2$, low T within the firm). Semiconductor fabrication requires thousands of precisely sequenced process steps. The lithography, etching, deposition, and doping stages are strong complements — a defect in any stage ruins the wafer. Within a fab, the process engineers share detailed metrology data (low T). Outsourcing individual process steps would require sharing proprietary recipes with external partners (high T_{market}). **Prediction: integrate.** Reality: TSMC, Samsung, and Intel all run their own fabs as vertically integrated operations.

Testing and packaging ($\rho \approx 0.2$, medium T). Once chips are fabricated, they must be tested and packaged. These steps are complementary with fabrication but less tightly coupled — a chip either passes the test or it does not, and the interface between fab and test is well-standardized. Information about test requirements can be communicated through specifications rather than tacit knowledge (medium T_{market}). **Prediction: mixed or outsource.** Reality: many firms outsource to OSAT (Outsourced Semiconductor Assembly and Test) partners like ASE and Amkor, though some advanced packaging stays in-house.

Materials procurement ($\rho \approx 0.8$, low T). Silicon wafers, photoresists, and specialty gases are near-commodities. Multiple qualified suppliers exist for each material. Specifications are standardized and quality is easily verified through incoming inspection (low T_{market}). Inputs are near-substitutes — switching from one silicon wafer supplier to another requires minimal adjustment. **Prediction: buy from the market.** Reality: every semiconductor firm purchases materials from external suppliers. No one integrates backward into silicon ingot production.

The pattern is striking. As we move down the production chain from design to materials, ρ increases (complementarity weakens) and T_{market} generally decreases (market information improves). The governance structure shifts from full integration to full market procurement, exactly as the theory predicts.

What the Data Say

The semiconductor example is vivid but anecdotal. Does the pattern hold systematically across industries?

The *test:governance-boundaries-contract-intensity* uses data from the Bureau of Economic Analysis on vertical integration patterns and contract intensity across U.S. industries. Following (Nunn2007), contract intensity measures the fraction of an industry's inputs that require relationship-specific investments — a direct proxy for low ρ.

Theorem (Governance Boundary Prediction).

The elasticity of substitution $\sigma = 1/(1-\rho)$ predicts the share of vertically integrated production across industries:

$$\beta = -0.120, \quad p < 0.001, \quad R^2 = 0.53$$

> Industries with lower σ (stronger complementarity, more negative ρ) have systematically higher rates of vertical integration, explaining over half the cross-industry variation in governance structure.

The negative coefficient confirms the theory: as substitutability rises (higher σ, higher ρ), firms integrate less and rely more on markets. The R^2 of 0.53 is remarkably high for a cross-industry regression with a single structural parameter. Most of the variation in how firms organize their production can be explained by one number: how complementary are their inputs?

Why This Matters

The boundaries-of-the-firm question is not merely academic. Every day, executives decide whether to build a capability in-house or source it externally. Management consultants charge enormous fees to advise on "core competency" and "strategic outsourcing." The CES framework says the answer is not about corporate culture or strategic vision — it is about (ρ, T).

This has practical implications:

When technology changes T, firm boundaries should shift. The rise of cloud computing lowered T_{market} for IT infrastructure. The framework predicts that firms should (and did) outsource data centers to AWS, Azure, and Google Cloud. The rise of large language models is currently lowering T_{market} for software development, legal research, and customer service. The framework predicts further outsourcing of these functions.

When regulation changes T, industry structure should shift. Disclosure requirements lower T_{market}. The framework predicts that industries with strong disclosure regimes (public utilities, pharmaceuticals) should be less vertically integrated than industries with weak disclosure (private equity, defense contracting). This is broadly consistent with observed patterns.

When complementarity changes, boundaries should shift. As AI makes previously complementary tasks more modular (ρ rises), the framework predicts disintegration. Activities that required tight internal coordination — because only human judgment could bridge the gaps between tasks — become outsourceable once AI standardizes the interfaces.

The Bigger Picture

This article has shown that the CES framework, through *effectiveCurvatureKeff* and the (ρ, T) *regime diagram*, provides a quantitative answer to Coase's 1937 question. Firms exist because market information friction erodes the complementarity premium. The boundary between firm and market is not a matter of managerial preference — it is determined by two measurable parameters.

The connection to the broader framework is direct. *The Regime Diagram* introduced the (ρ, T) diagram and *The Economics of Not Knowing* defined *effective curvature*. This article applies that machinery to one of the oldest questions in economics. *The Economy Has Layers* extends the analysis to hierarchies of firms — supply chains where the boundaries-of-the-firm question arises at every level, and the answer at each level depends on the answers at adjacent levels.

The key takeaway: the make-or-buy decision is not a judgment call. It is a calculation. Measure ρ (how complementary are the inputs?) and T (how costly is market information?), and the governance boundary tells you the answer.

References