The Economics of Not Knowing

Jon Smirl

2026-03-01

You're Buying a Used Car

The seller smiles and tells you the engine is in great shape. Maybe it is. Maybe it is not. You cannot pop the hood and instantly know the difference between a reliable sedan and a money pit — not without paying a mechanic to inspect it, pulling the vehicle history report, and spending an afternoon you would rather spend doing almost anything else.

This is the famous "lemons" problem (Akerlof1970). Because buyers cannot easily tell good cars from bad, they assume the worst and offer low prices. Sellers of genuinely good cars refuse those prices and leave the market. What remains is a market dominated by lemons — not because good cars do not exist, but because *information is expensive*.

Akerlof won the Nobel Prize for this insight, and it is usually taught as a story about used cars. But the same logic shows up everywhere. A bank evaluating a loan application cannot perfectly observe the borrower's true risk. An employer interviewing a job candidate cannot see inside that person's head. An insurer offering a health plan does not know who will get sick. In every case, the cost of knowing — the gap between what you need to know and what you can cheaply observe — reshapes the market.

The question this article asks is simple: can we capture all of these information costs with a single parameter? The answer is yes, and it changes the way we think about production, markets, and crises.

One Parameter: Information Friction T

Think of T as a fog dial. At $T = 0$, visibility is perfect. Every buyer sees every product's true quality. Every bank knows every borrower's true risk. Decisions are sharp and markets are efficient.

Turn the dial up. At moderate T, agents squint through the fog. They can still tell obviously good from obviously bad, but the borderline cases blur together. Mistakes happen. Some good products get passed over; some bad ones get bought.

Crank T high enough, and the fog is impenetrable. Agents essentially pick at random. Markets collapse into the lemons equilibrium.

> **Definition (Information Friction).**
>
> The information friction $T \geq 0$ is the shadow price of distinguishing good from bad. When $T = 0$, agents are perfectly informed. As T increases, decisions become noisier and more costly to improve.

What makes T powerful is that it is not just a metaphor. The rational inattention literature, initiated by (Sims2003), shows that when agents have limited capacity to process information, their optimal decisions follow a precise mathematical form. The parameter governing the amount of noise in those decisions is exactly T — the inverse of the agent's information processing capacity.

One Parameter, Many Problems

Here is the payoff. Economists have spent decades studying information problems that look very different on the surface but share the same underlying structure. T unifies them:

Akerlof's lemons (Akerlof1970). The used car buyer faces high T because quality is hidden. The insurance company faces high T because health risk is private. In both cases, T measures the cost of closing the information gap between buyer and seller.

Spence's signaling (Spence1973). A job applicant gets a college degree not because college teaches useful skills, but because completing a degree signals ability. The cost of the signal is the applicant's way of *lowering* T for the employer. Expensive signals work precisely because they reduce the fog.

Stiglitz's screening (Rothschild1976). An insurer offers two plans — one cheap with a high deductible, one expensive with full coverage. Healthy people choose the cheap plan; sick people choose the expensive one. The menu of contracts is a device to *lower* T by getting customers to reveal their own type.

Sims's rational inattention (Sims2003). A consumer deciding between 47 brands of cereal does not carefully evaluate each one. She grabs the familiar box. Her inattention is rational: the cost of processing 47 options exceeds the benefit. That processing cost *is* T.

In every case, the same dial is being turned. What differs is who bears the information cost and how the market adapts to it.

> **Example.**
>
> Consider three markets with different levels of T:
>
> - **Online retail** (low T): Reviews, ratings, return policies, and price comparison sites make quality nearly transparent. T is close to zero.
> - **Hiring** (moderate T): Resumes, interviews, references, and probation periods reduce the fog, but significant uncertainty remains. T is moderate.
> - **Complex financial instruments** (high T): A collateralized debt obligation packages thousands of mortgages into tranches that even the issuing bank cannot fully evaluate. T is very high.

What Information Friction Does to Production

Emergent CES and the Quadruple Role of Curvature established that the curvature parameter $K = (1-\rho)(J-1)/J$ controls the diversity premium — the extra output you get from combining different kinds of inputs rather than using identical ones. A team of specialists outperforms a team of generalists, and the bonus is proportional to K.

But that result assumed $T = 0$. Everyone could see every input's quality perfectly. What happens when the fog rolls in?

The answer is the *effectiveCurvatureKeff*:

Theorem (Effective Curvature).

Under information friction T, the exploitable curvature is:

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

where T^* is a critical threshold that depends on the strength of complementarity. When $T \geq T^*$, effective curvature is zero: the diversity premium vanishes entirely.

Read that carefully. Information friction does not just add noise to an otherwise functioning economy. It *erodes the fundamental structure of production.* The diversity premium, the correlation robustness, the strategic stability — all the roles of K from *Emergent CES* — depend on agents being able to tell inputs apart. When the fog is thick enough, those roles disappear. Not gradually, but with a sharp cutoff at $T = T^*$.

And here is a key detail: T^* is higher when complementarity is stronger. A surgical team (ρ very negative, strong complements) can tolerate more information noise than a commodity market (ρ close to 1, near-substitutes) before the diversity premium vanishes. Strong complementarity provides a buffer — but even that buffer has a limit.

Why This Matters: The Same Fog, Different Damage

This framework explains something that has puzzled policymakers. The same increase in information friction can be harmless in one sector and catastrophic in another. A moderate rise in T — say, from a new regulation that makes disclosure harder — barely dents a tightly complementary manufacturing process where T^* is high. But it can destroy the functioning of a weakly complementary financial market where T^* is low.

The 2007–2008 financial crisis is a case in point. When investors realized they could not value the mortgages underlying complex securities, T spiked across the financial system. The *banking crisis severity test* confirms the prediction: across 147 countries, the interaction of substitutability and information friction predicts crisis severity ($p = 0.016$). Countries with higher pre-crisis opacity experienced worse outcomes, exactly as the effective curvature theorem predicts.

The Map So Far

Emergent CES gave us ρ (or equivalently K): the curvature of production, measuring how much diversity matters. This article adds T: the cost of seeing clearly. Together, (ρ, T) form a two-dimensional map. Every market, firm, and institution sits somewhere on this map.

- **Low T, low ρ** (transparent, complementary): High-performance teams, precision manufacturing. The diversity premium is large and fully exploitable.
- **Low T, high ρ** (transparent, substitutable): Commodity markets, online retail. Inputs are interchangeable and markets work well, but there is little diversity premium to capture.
- **High T, low ρ** (opaque, complementary): Complex financial instruments, developing-country supply chains. The production structure *could* generate a big diversity premium, but information friction erodes it.
- **High T, high ρ** (opaque, substitutable): The lemons market. Weak complementarity plus poor information. Akerlof's world.

The *CES potential* $\mathcal{F} = \Phi_{\text{CES}}(\rho) - T \cdot S_q$ formalizes this map by combining the production value (first term) with the information cost (second term) into a single object. The next articles explore what happens when the economy moves across this map — how crises unfold in a predictable sequence, and how the dynamics of adjustment determine whether recovery takes months or a decade.

References