Emergent CES and the Quadruple Role of Curvature

Jon Smirl

2026-03-01

What Is an Aggregation Function?

Economists constantly face the problem of combining many different things into a single number. How do you summarize the output of an economy that produces cars, haircuts, and software? How do you measure the combined productivity of workers with different skills?

An **aggregation function** takes J inputs $x_1, x_2, \ldots, x_J$ and produces a single output $F(x_1, \ldots, x_J)$. The most familiar example is the simple average:

$$F = \frac{1}{J} \sum_{j=1}^{J} x_j$$

But a simple average treats all inputs as perfect substitutes – losing one unit of x_1 is exactly compensated by gaining one unit of x_2. In most economic settings, this is unrealistic. A factory needs *both* workers and machines; a portfolio needs *both* stocks and bonds. The inputs are *complements*, not substitutes.

The CES Function

The **Constant Elasticity of Substitution** (CES) function, introduced by (Arrow1961), generalizes the average by adding a single parameter ρ that controls how substitutable the inputs are:

> **Definition (CES Aggregation Function).**
>
> For J inputs with equal weights:
>
> $$F = \left(\frac{1}{J} \sum_{j=1}^{J} x_j^{\rho} \right)^{1/\rho}$$
>
> where $\rho \leq 1$ is the substitution parameter. The **elasticity of substitution** is $\sigma = 1/(1-\rho)$.

The parameter ρ controls the shape of the function:

- $\rho = 1$: Perfect substitutes. F is the arithmetic mean. Losing one input is fully compensated by gaining another.
- $\rho \to 0$: Cobb-Douglas. F approaches the geometric mean. This is the textbook production function $Y = K^{\alpha} L^{1-\alpha}$ (Solow1956).

- $\rho \to -\infty$: Perfect complements (Leontief). $F = \min(x_1, \ldots, x_J)$. Output is determined entirely by the scarcest input.

Most real economies operate in the complementary range $\rho < 0$ ($\sigma < 1$). Estimates from manufacturing data typically find $\sigma \approx 0.4$–0.8, meaning $\rho \in [-1.5, -0.25]$.

Why CES Is Not Just an Assumption

Most economics papers simply *assume* a CES production function because it is analytically convenient. But the deeper question is: *why should CES be the right functional form?* The answer, formalized as *Theorem: Emergent CES*, is that CES is not a choice – it is the *only* function satisfying natural economic requirements.

There are three independent proofs of this result:

Multi-Scale Aggregation Fixed Point

Imagine aggregating first within factories, then across factories within industries, then across industries within the economy. If the aggregation function is *consistent* across these levels – meaning the same functional form works whether you aggregate two inputs or two thousand – then it must be CES.

> **Theorem (Kolmogorov-Nagumo Aggregation).**
>
> If an aggregation function $M(x_1, \ldots, x_J)$ satisfies: 1. **Symmetry**: permuting inputs does not change the result 2. **Continuity**: small changes in inputs produce small changes in output 3. **Consistency**: $M(M(x_1, x_2), M(x_3, x_4)) = M(x_1, x_2, x_3, x_4)$
>
> then M must be a power mean: $M = \left(\frac{1}{J}\sum x_j^r\right)^{1/r}$ for some r, which is exactly CES with $\rho = r$ (Kolmogorov1930).
>
> *Proof.*
>
> The key insight is that condition (3) forces the existence of a continuous, strictly monotone function φ such that $M = \varphi^{-1}\left(\frac{1}{J}\sum \varphi(x_j)\right)$. This is the Kolmogorov-Nagumo theorem. Combined with the homogeneity requirement (constant returns to scale), φ must be a power function $\varphi(x) = x^\rho$, yielding CES.

The Aczel Functional Equation

A second proof comes from pure mathematics. (Aczel1948) showed that power means are the *only* means satisfying a natural bisymmetry condition. When combined with the economic requirement of constant returns to scale, this again uniquely determines CES.

Maximum Entropy Self-Consistency

A third proof uses information theory. Among all production functions that (a) have constant returns to scale and (b) are consistent with the maximum-entropy distribution over inputs, only CES survives. This connects the production function to the *statistical properties* of the economy.

The Curvature Parameter K

The substitution parameter ρ and the number of inputs J combine into a single curvature measure:

Definition (Curvature Parameter).

$$K = \frac{(1-\rho)(J-1)}{J}$$

This is the *curvature* of the CES isoquant at the symmetric point where all inputs are equal. It measures how "curved" the production possibilities are. When $K = 0$ (either $\rho = 1$ or $J = 1$), the isoquant is flat – inputs are perfect substitutes or there is only one input. As K increases, the isoquant curves more sharply, meaning balanced allocation becomes increasingly important.

The central result of this paper is *Theorem: Quadruple Role*: this single parameter K simultaneously controls four distinct economic phenomena.

Role 1: Superadditivity

When inputs are complements ($\rho < 0$, so $K > 0$), combining heterogeneous inputs produces *more* than the sum of their parts. This is the *superadditivity* premium:

Theorem (Superadditivity Bound).

For any input vector $\mathbf{x}$ with CES aggregate F:

$$F \geq \bar{x} \cdot \left(1 + \frac{K}{2} \cdot \mathrm{CV}^2\right)$$

where $\bar{x}$ is the arithmetic mean and CV is the coefficient of variation of inputs. The bonus is proportional to K times input heterogeneity.

Intuition: A team of diverse specialists (high CV) produces more than a team of identical generalists, and this bonus grows with K. This is why firms hire people with different skills rather than cloning their best employee.

Example: With $\rho = -1$ ($\sigma = 0.5$) and $J = 4$ inputs, $K = 1.5$. If inputs have $\mathrm{CV} = 0.3$, the superadditivity bonus is approximately $1.5 \times 0.09/2 = 6.75\%$ – heterogeneous inputs produce nearly 7% more than the homogeneous benchmark.

Role 2: Correlation Robustness

The same curvature that creates superadditivity also provides *robustness to correlated shocks*. When all inputs move together (as in a recession), a linear aggregate captures only the common movement. CES also captures the *idiosyncratic* variation – and this bonus is proportional to K^2:

Theorem (Correlation Robustness).

Under random shocks with compound symmetry correlation r:

$$\text{CES Bonus} \approx K^2 \cdot \mathrm{Var}(\text{idiosyncratic component})$$

CES extracts value from the idiosyncratic variation that linear aggregates miss.

This has a direct connection to the *estimation paradox*: the very curvature that makes CES harder to estimate statistically is what makes it economically valuable. If CES were easy to distinguish from Cobb-Douglas in data, the diversity premium it captures would be small.

This result also connects to **model collapse** in AI: when training data becomes self-referential (AI training on AI-generated text), the effective number of independent signals shrinks, reducing the CES diversity premium.

Role 3: Strategic Independence

The curvature parameter K also controls strategic behavior. Consider J agents each choosing their own input level, with payoffs depending on the CES aggregate:

> **Theorem (Nash Stability).**
>
> At the symmetric equilibrium where all agents contribute equally: - The equilibrium is a **Nash equilibrium** – no single agent can profitably deviate (Nash1950) - No **coalition** of agents can profitably redistribute among themselves - The stability penalty for deviation is proportional to K

Intuition: When K is large (strong complementarity), each agent's contribution is essential. Deviating hurts the deviator more than it helps, because the CES function penalizes imbalance. This is why cartels are harder to sustain in industries with complementary inputs – the incentive to deviate from the cartel agreement is reinforced by the curvature of the production function.

Role 4: Network Scaling

The fourth role connects K to *network structure*. In a network of J nodes with CES aggregation at each node, the **spectral gap** of the network Hessian (the gap between the largest and second-largest eigenvalues) is proportional to K:

$$\lambda_1 - \lambda_2 \propto K$$

This spectral gap controls how quickly information, prices, and shocks propagate through the network. A larger K means faster convergence to equilibrium and a cleaner separation between aggregate and idiosyncratic dynamics.

The *trade spectral gap test* confirms this prediction using bilateral trade data from COMTRADE: the Fiedler distance (a measure of network connectivity) correlates negatively with trade concentration ($r = -0.213$), consistent with the theory.

Why One Parameter Controls Four Phenomena

These four roles are not coincidences – they are **four views of the same geometric fact**. The CES isoquant at the symmetric point has a specific curvature tensor, and K is its scalar summary. Each role corresponds to a different way of probing this curvature:

Role	How K enters	Economic meaning
Superadditivity	$F - \bar{x} \propto K \cdot \mathrm{CV}^2$	Diversity premium

Role	How K enters	Economic meaning
Correlation robustness	Bonus $\propto K^2 \cdot \mathrm{Var}$	Idiosyncratic value
Strategic independence	Deviation cost $\propto K$	Equilibrium stability
Network scaling	Spectral gap $\propto K$	Information speed

The *CES isotropy test* using BEA input-output data confirms the geometric prediction: the relationship between sectoral concentration (Herfindahl index (Herfindahl1950)) and curvature is approximately isotropic, with $\beta = 0.54$ (the theoretically predicted value for isotropic curvature is 1, and the manufacturing subsample gives $\beta = 0.85$ with confidence intervals containing 1).

What Comes Next

This article has introduced CES emergence and the quadruple role of curvature in the simplest setting: equal-weight, symmetric equilibrium. The full theory extends in several directions:

- **General weights** and the secular equation
- The **CES potential** $\Phi = -\sum \log F_n$ as a landscape governing dynamics — see *The Economic Landscape*
- **Information friction** and market breakdown — see *The Economics of Not Knowing*
- **Hierarchical architecture** with multiple CES levels and timescale separation — see *The Economy Has Layers*
- **Empirical tests** of all predictions against real data — see *48 Predictions and Counting*

The key takeaway: CES is not an arbitrary modeling choice. It is a mathematical necessity for consistent aggregation, and its curvature parameter K is the single number that organizes production economics.

References