Technology Constraints = Welfare Losses: The Eigenstructure Bridge

Jon Smirl

2026-03-01

The Bottleneck Is the Problem

During the 2020–2022 global chip shortage, automakers lost an estimated $210 billion in revenue. Consumers waited months for cars, appliances, and electronics. Governments scrambled to subsidize new fabrication plants. Everyone agreed on two things: (1) semiconductor fab capacity was the technological bottleneck, and (2) the shortage was causing enormous welfare losses — higher prices, lost output, delayed products, cascading disruptions through supply chains.

These two observations felt like separate facts. One was about technology (we cannot make enough chips). The other was about welfare (people are worse off). But the deepest result in the CES hierarchy framework says they are not separate facts at all. They are **the same fact**, viewed from two different angles.

The sector that is hardest to change technologically is always the sector causing the most welfare loss. The bottleneck *is* the welfare problem. This is not an approximation or a rule of thumb. It is a theorem.

The Bridge

The formal statement connects three objects. The first is the CES potential Φ, which encodes the economy's technological structure — how sectors combine inputs, how substitutable they are, and where capacity constraints bind. Its Hessian $\nabla^2\Phi$ (the matrix of second derivatives, restricted to the slow-moving variables) measures how *curved* the technological landscape is at each point. High curvature means a sector can adjust easily; low curvature means it is stuck.

The second is the welfare loss function V, which measures how far the economy is from its efficient allocation. Its Hessian $\nabla^2 V$ measures how sensitive welfare is to displacements in each sector. Large values mean that being slightly off-target in that sector costs a lot of welfare.

The third is the institutional supply-rate matrix W. This is the new ingredient. W measures how efficiently institutions — markets, regulators, firms, trade agreements — translate technological *possibility* into actual *adjustment*. A sector might have the engineering capacity to expand, but if permitting takes five years, or if a single monopolist controls all supply, the effective adjustment rate is low.

Theorem (Eigenstructure Bridge).

$$\nabla^2\Phi\big|_{\text{slow}} = W^{-1} \cdot \nabla^2 V$$

The curvature of the technological landscape equals the inverse of the institutional supply-rate matrix times the curvature of the welfare loss function.

Read from left to right: where technology is rigid ($\nabla^2\Phi$ has a large eigenvalue), welfare is sensitive ($\nabla^2 V$ is large) — scaled by how efficiently institutions transmit adjustment (W). Read from right to left: welfare losses are just technological constraints filtered through institutional capacity.

What This Means Without the Math

Forget the matrices for a moment. The bridge theorem says something simple and powerful:

You do not need to measure welfare directly.

Measuring welfare is notoriously hard. You would need to know everyone's preferences, compute counterfactual allocations, and aggregate across millions of heterogeneous agents. Economists have debated how to do this for a century, since (Samuelson1947).

But measuring technological bottlenecks is *easy*. Which sectors have capacity constraints? Where are lead times longest? Which inputs are in shortage? This information is observable in real time from industry data, purchasing manager surveys, and price signals.

And measuring the institutional matrix W is *feasible*. How concentrated is supply? How long do permits take? What trade restrictions are in place? How quickly can investment flow to the constrained sector? These are estimable quantities.

The bridge theorem says: multiply the technological bottleneck measure by the institutional friction measure, and you get the welfare loss. Technology times institutions equals welfare. No preference estimation required.

The Semiconductor Example

Consider the COVID-era chip shortage in detail. The technological landscape had extremely low curvature in the semiconductor direction — building a new leading-edge fab takes 3–5 years, costs $20+ billion, and requires specialized equipment from a handful of suppliers. The Hessian of Φ in the semiconductor direction was very large (steep penalty for deviating from existing capacity).

The institutional matrix W for semiconductors was also unfavorable. TSMC controlled over 90% of advanced logic production, concentrated on a single island. Export controls between the U.S. and China restricted technology transfer. Environmental and zoning approvals for new fabs added years to timelines. Capital markets, while willing to fund expansion, could not accelerate the physical construction.

The bridge theorem predicts that the welfare loss Hessian in the semiconductor direction should be correspondingly enormous — and it was. The auto industry alone lost the equivalent of 5–7 million vehicles. Consumer electronics faced 6–12 month delays. Downstream industries from medical devices to defense systems experienced cascading shortages. The welfare cost was concentrated exactly where the theorem says it should be: in the direction of the hardest-to-move technological constraint, amplified by institutional rigidities.

Now contrast this with another sector that also faced COVID disruptions: restaurant services. Restaurants closed, reopened, adapted to delivery, and adjusted capacity within weeks to months. The technological curvature in the restaurant direction was low (easy to adjust), and the institutional matrix was favorable (minimal regulatory barriers to reopening, competitive market structure, low capital requirements). The bridge theorem predicts smaller welfare losses per unit of disruption in this direction — and indeed, while restaurant closures were painful, the sector adjusted far faster than semiconductors, and the aggregate welfare cost, while significant, did not cascade through the economy in the same way.

The Damping Cancellation Surprise

The bridge theorem has a startling corollary known as the *damping_cancellation_algebraic*. Suppose a government wants to reduce the welfare loss in a constrained sector. The obvious policy is to increase local regulation or oversight — tighten standards, add monitoring, force faster adjustment within that sector. This increases the diagonal entry of W for that sector.

The theorem says this backfires in a precise way. Increasing W_{nn} (the institutional adjustment rate for sector n) does two things: it speeds up convergence toward equilibrium, but it also *lowers* the equilibrium output level. These two effects exactly cancel. The net welfare impact is zero.

This is not a loose approximation. It is an exact cancellation that follows from the eigenstructure of the bridge equation. Making sector n adjust faster internally does not help, because the binding constraint is not sector n's internal efficiency — it is the *upstream* bottleneck feeding into sector n.

The implication for policy is direct: to relieve sector n's welfare contribution, do not reform sector n. Reform sector $n - 1$ — the upstream supplier, the input market, the slower-moving layer of the hierarchy. During the chip shortage, subsidizing automakers to use chips more efficiently (reforming the auto sector) would have been futile. The binding constraint was fab capacity (the upstream sector). The CHIPS Act, which subsidized fab construction, was the correct policy direction — it targeted the upstream bottleneck, not the downstream symptom.

Why the Bridge Is Deep

Most economic results connect *either* technology *or* preferences to outcomes. The bridge theorem connects *both simultaneously*, through an intermediary (W) that is itself an object of policy. This three-way connection — technology, institutions, welfare — in a single matrix equation is what makes the result fundamental.

It also explains why economic crises feel so intractable. During a crisis, technological bottlenecks tighten ($\nabla^2\Phi$ grows), institutional capacity is strained (W shrinks), and welfare losses compound ($\nabla^2 V$ explodes). The bridge theorem shows these are not three separate problems compounding by coincidence. They are one problem, expressed in three vocabularies. Fix the bottleneck, improve the institutions, and welfare follows — not as a hopeful side effect, but as a mathematical identity.

References