DRAFT: NOT FOR PUBLICATION

Structural Tightness in Benchmark Input-Output Data: A Minimum-Deformation Approach

Mei Yuan Mustafa Babiker

Abstract

We study a structural layer of benchmark input-output data that is present before prices, substitution, and equilibrium adjustment are introduced. Using the benchmark intermediate-use matrix from the official U.S. Bureau of Economic Analysis supply-use tables, we examine how much the benchmark flow structure must be reorganized when one row or column total is perturbed and the new accounting constraints must still be satisfied while remaining as close as possible to the original matrix. We implement this diagnostic with the iterative proportional fitting procedure and measure the required deformation with Kullback-Leibler divergence. Across benchmark years, the sectoral ranking implied by these deformation measures is relatively stable, with limited reshuffling within the leading group under both the 15-sector and 73-sector classifications. Sectoral disaggregation then sharpens that ranking by identifying more precisely which industries occupy the top of the distribution, first under the 73-sector classification and then under a more detailed classification. Finally, using equal-step shocks under the 15-sector classification, we find that raw KL deformation rises more than proportionally with shock size, indicating that larger shocks require disproportionately more reorganization of the benchmark intermediate-use structure. These results suggest that benchmark input-output data already contain a persistent and nonlinear structural layer before behavioral assumptions are imposed. This layer helps identify where structural tightness is concentrated, how it changes with sectoral detail, and how it intensifies as shocks grow.

1 Introduction

Computable general equilibrium (CGE) analysis is often understood through behavioral adjustment: relative prices move, substitution takes place, and the economy settles into a new equilibrium outcome. But those responses are built on a benchmark input-output (I-O) structure that already embodies an organization of production linkages. If that underlying organization already contains its own tensions, then some of what later unfolds through behavioral adjustment may already be conditioned by the benchmark itself. Does the benchmark network contain such a structural layer, one that matters even before behavioral assumptions are imposed? Put differently, is part of the economy's later adjustment already being shaped by the organization embedded in the benchmark data, before prices and substitution begin to do their work?

Focusing on this structural layer, we examine how much the benchmark flow pattern must be reorganized when a row or column total is perturbed and the new accounting constraints must still be satisfied while remaining as close as possible to the original structure. If only limited reweighting is needed, then the network can absorb that shock with relatively little internal change. If much more reweighting is required, then the shock exposes greater structural tension in the benchmark organization. We also trace where that reweighting falls, since some shocks spread the adjustment broadly across existing links, while others place it much more heavily on a smaller part of the network.

We implement this structural diagnostic using the iterative proportional fitting procedure (IPFP), in the biproportional fitting approach familiar from input-output table rebalancing. It yields two kinds of information. One is the overall extent of the required reorganization, or deformation, of the benchmark structure. The other is where that reorganization falls and how unevenly it is distributed across existing input-output linkages and sectors. We measure this deformation with Kullback–Leibler (KL) divergence because it provides a natural measure of minimum distortion from an observed benchmark matrix, summarizing how much the benchmark production network must be reweighted in order to accommodate a given shock while still satisfying the new accounting constraints. Together, these measures make the structural layer visible by showing both how much deformation is required under different shocks to row or column totals and where the resulting pressure becomes most strongly concentrated.

Using this framework, we find three main patterns. The first comes from sector disaggregation. Comparing a coarse sector aggregation with a finer sector aggregation shows that some subsectors whose structural importance is muted at the coarse level become much more visible under finer disaggregation. The second is multi-year continuity. Although some sectors move more than others, the broad ranking is relatively stable across years. The third is shock-size variation. As shocks become larger, the implied deformation rises disproportionately, suggesting that the amount of structural reorganization required by the benchmark network is not linear in the size of the shock. Taken together, these patterns point to a structural layer that is already present in the official benchmark tables before price-mediated adjustment begins. A later comparison with CGE results would need to recognize that the benchmark dataset used for CGE calibration is itself a modified version of the original BEA table. This makes it important to understand how the same patterns appear in the processed benchmark representation, since some may be preserved while others may be altered during the build-stream processing.

The rest of the paper proceeds as follows. Section 2 presents the approach and study design. In Section 3, we present the main results from the official BEA benchmark tables. Section 4 concludes with the implications for processed benchmark representations used in CGE model calibration.

2 Approach

To isolate the production-network layer of the benchmark I-O table, we work with the benchmark intermediate-use matrix, where intersectoral intermediate linkages can be examined on their own before final demand, value added, and other components are brought back in. Let z_{ij} denote the benchmark flow from input i into sector j in that matrix. We consider one shock at a time. A row shock changes the total of intermediate sales associated with one input, while a column shock changes the total intermediate input requirement of one sector. In that sense, each shock defines a direction of adjustment, asking how the benchmark matrix would have to respond if one row total or one column total were changed while the rest of the accounting structure still had to be reconciled.

For a given shock, we write the update as a constrained minimization problem, choosing an updated matrix x_{ij} to minimize its distortion relative to the benchmark matrix z_{ij} subject to the perturbed row and column totals. We measure that distortion with Kullback-Leibler (KL) divergence.

$$\min_{x_{ij}} \sum_{i,j} \left[x_{ij} \log\left(\frac{x_{ij}}{z_{ij}}\right) - x_{ij} \right] \tag{1}$$

subject to

$$\textstyle\sum_j x_{ij} = r_i \qquad \sum_i x_{ij} = s_j \tag{2}$$

where r_i and s_j denote the row and column totals implied by a given shock. In this formulation, the updated matrix is the minimum-distortion reweighting of the benchmark matrix that satisfies the new accounting constraints.

The associated Lagrangian is

$$L = \textstyle\sum_{i,j}\left[x_{ij}\log\left(\frac{x_{ij}}{z_{ij}}\right) - x_{ij}\right] + \sum_i \alpha_i\left(r_i - \sum_j x_{ij}\right) + \sum_j \beta_j\left(s_j - \sum_i x_{ij}\right) \tag{3}$$

The first-order conditions imply

$$\log\left(\frac{x_{ij}}{z_{ij}}\right) = \alpha_i + \beta_j \tag{4}$$

or

$$\frac{x_{ij}}{z_{ij}} = e^{\alpha_i}e^{\beta_j} \tag{5}$$

so that the updated matrix takes the biproportional form

$$x_{ij} = p_i z_{ij} q_j, \qquad p_i = e^{\alpha_i}, \qquad q_j = e^{\beta_j} \tag{6}$$

where the row and column scaling factors p_i and q_j are functions of the Lagrange multipliers and are chosen so that the perturbed row and column totals are satisfied. This shows that the biproportional updating structure arises directly from the minimum-distortion problem itself.

The scaling factors can be computed by the iterative proportional fitting procedure (IPFP), which alternates between row and column adjustments until the required totals are satisfied. In the two-way matrix setting used here, this is the familiar biproportional balancing logic commonly referred to as RAS in input-output applications.

The structural diagnostic retains both the magnitude and the distribution of the required update, providing overall deformation driven by a shock and where that deformation falls. On distribution of deformation, some shocks require relatively diffuse reweighting across many existing input-output linkages, while others concentrate the adjustment much more heavily in a smaller set of sectors or intermediate uses. This distinction matters because two shocks may require similar total deformation while differing sharply in how concentrated the associated pressure becomes.

We apply this procedure across four directional scans. For row totals, we consider both positive and negative shocks of the same proportional size. We do the same for column totals, again using matched proportional increases and decreases. This yields four directions in total. The resulting deformation measures are then compared across sectoral aggregation levels, across multiple benchmark years, and across shock sizes. These comparisons allow us to examine how the structural ranking changes with sectoral disaggregation, how stable it remains over time, and how deformation scales with the size of the shock.

The empirical analysis uses the official benchmark supply-use tables published by the U.S. Bureau of Economic Analysis (BEA). From each benchmark year, we extract the intermediate-use table and use it as the benchmark matrix. The sample spans 1997 through 2024, yielding 28 benchmark years. In what follows, we refer to the sector whose row or column total is perturbed as the trigger sector. We use carrier sector for a sector that absorbs a relatively large share of the resulting deformation once the benchmark matrix is reweighted to satisfy the new constraints.

3 Results

A structural diagnostic is useful only if it is not tied too narrowly to a single benchmark year, a single level of aggregation, or a single shock size. We therefore examine the results along three dimensions, benchmark years, sectoral aggregation, and shock size. These comparisons help clarify whether the pattern reflects a persistent feature of the benchmark production network, whether finer disaggregation reveals structure that is muted in broader aggregates, and whether larger shocks require qualitatively more reorganization than smaller ones.

3.1 Structural tightness shows strong continuity across benchmark years

We first examine the multi-year continuation results under the 15-sector classification. Figure 1 plots, for each benchmark year, the mean KL deformation per shock averaged across the four directional scans. The main pattern is the stability of the upper part of the ranking over time. The highlighted sectors remain near the top throughout the sample, while the rest of the ranking changes only gradually rather than being repeatedly reshuffled from year to year. This suggests that the benchmark intermediate-use structure preserves a persistent pattern of directional structural tightness over time, even though the exact levels move somewhat across years.

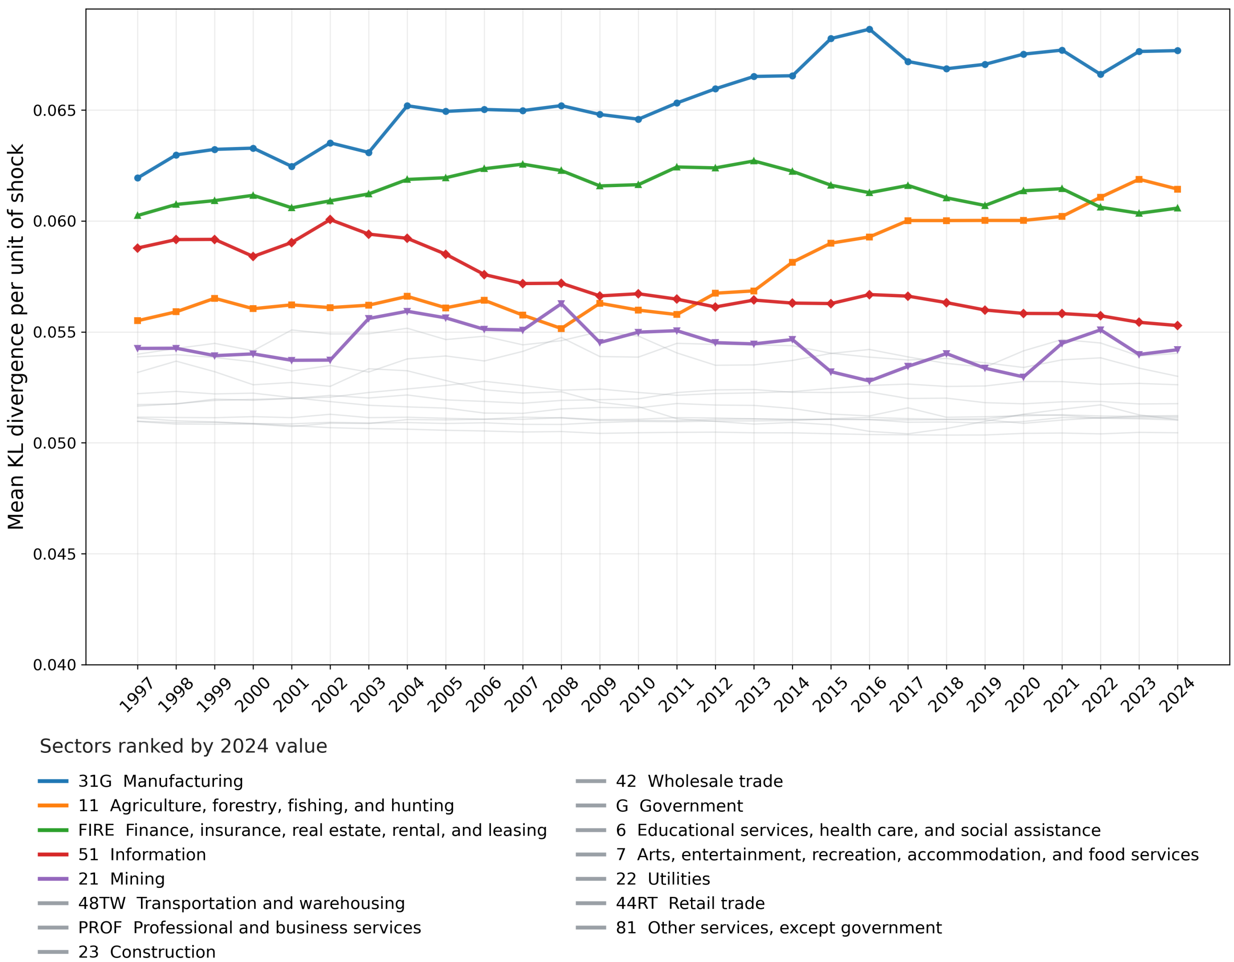

Figure 1: Multi-year continuation of mean KL divergence per unit of shock under the 15-sector classification, 1997-2024. The highlighted lines indicate the five highest-ranked sectors in 2024, while the remaining sectors are shown in gray. The ranked list below reports all 15 sectors in descending order of their 2024 values. The y-axis is truncated at 0.04 to improve readability.

Under the 73-sector classification, the continuation result remains clear, but the upper tail becomes much more sharply separated. Figure 2 shows that most sectors remain compressed in a relatively narrow lower background band, while a smaller set of sectors stands distinctly above it. This upper tail is also highly persistent across benchmark years. Using the stacked 1997-2024 results, adjacent-year top-10 overlap averages about 0.97, indicating very limited year-to-year reshuffling within the leading group. Sector 211 and sector 324 remain fixed at ranks 1 and 2 throughout the full sample, while several other sectors remain in the top 10 in every benchmark year.

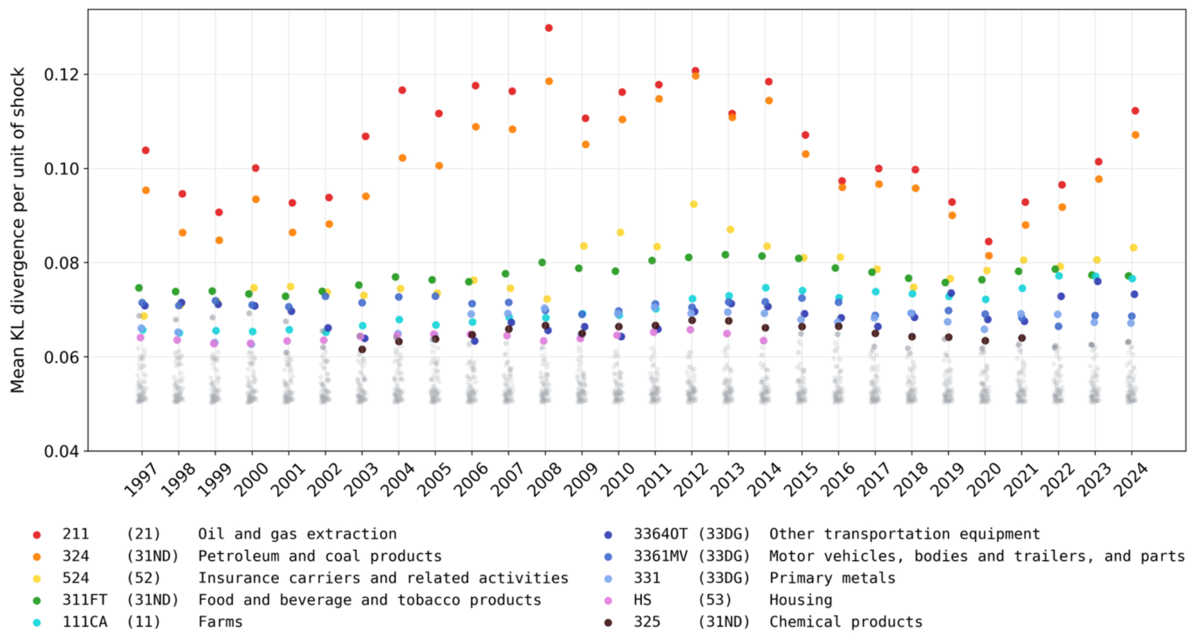

Figure 2: Multi-year continuation of mean KL divergence per unit of shock under the 73-sector classification, 1997-2024. Colored markers indicate the ten highest-ranked sectors based on their long-run position in the ranking, while the remaining sectors are shown in gray. Sector labels in the legend report the 73-sector code together with the corresponding 15-sector code in parentheses. The y-axis is truncated at 0.04 to improve readability.

This continuity result is important for interpretation. The ranking is not identical in every year, and some movement does occur within the upper group. Even so, the overall pattern is quite stable. Across both levels of sectoral aggregation, directional structural tightness appears as a persistent feature of the benchmark intermediate-use structure rather than as a one-year artifact.

3.2 Sector disaggregation sharpens the top of the ranking

Under the 15-sector classification, the higher ranks are occupied by groups such as Manufacturing, Agriculture, FIRE, Information, and Mining. Under the 73-sector classification, narrower sectors stand out more clearly near the top of the ranking, including Oil and gas extraction, Petroleum and coal products, Insurance carriers and related activities, Food and beverage and tobacco products, and Motor vehicles, bodies and trailers, and parts. Moving from 15 sectors to 73 sectors therefore makes it easier to see which subsectors are associated with the highest values of KL divergence per unit of shock.

The mining group provides one clear example. Under the 15-sector classification, sector 21, Mining, remains in the upper part of the ranking. Under the 73-sector classification, sector 211, Oil and gas extraction, separates clearly from the other mining-related subsectors and occupies the top position on average across benchmark years. A similar sharpening appears within

Manufacturing. At the 15-sector level, Manufacturing remains near the top of the ranking. At the 73-sector level, that position is reflected much more unevenly across subsectors such as Petroleum and coal products, Food and beverage and tobacco products, Motor vehicles, and Other transportation equipment.

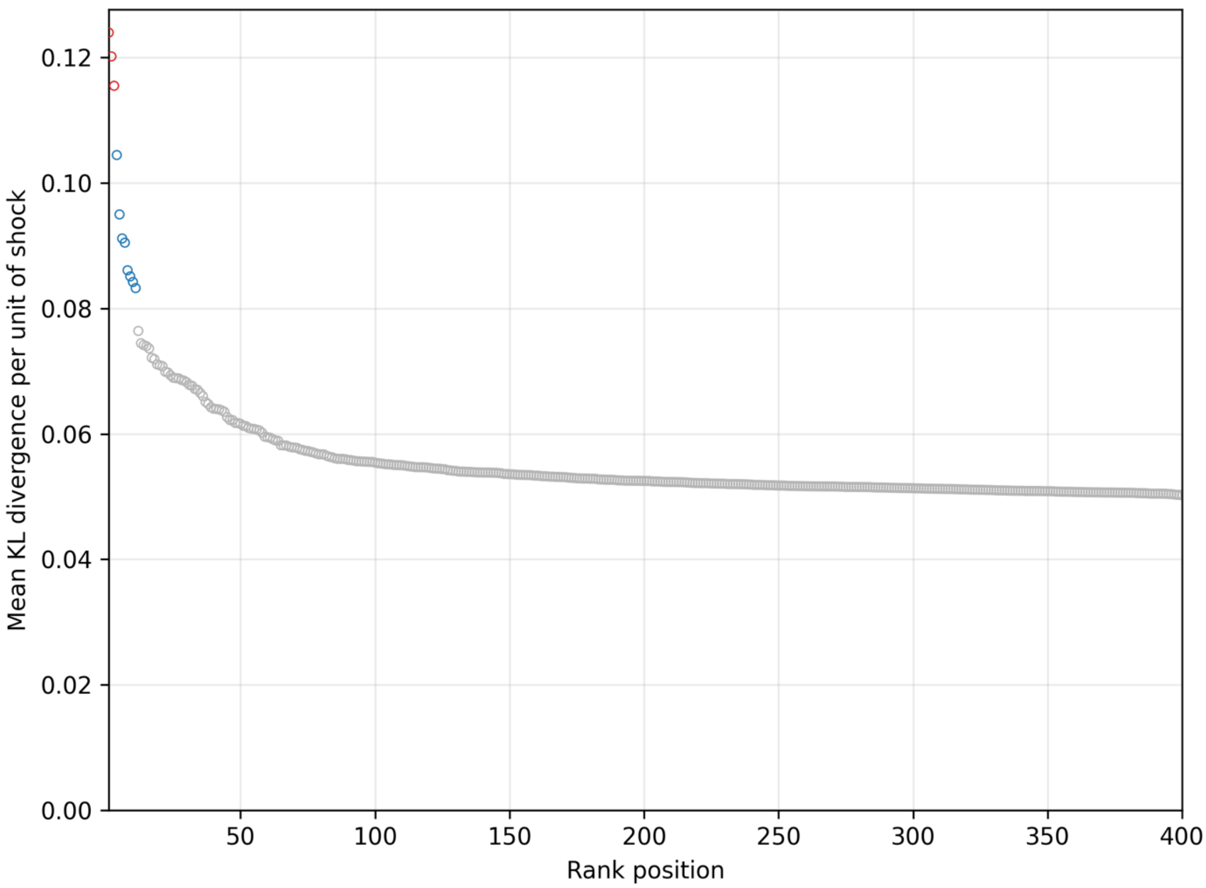


```
Top 11 sectors (ranked):

 1. 211000 (211)   Oil and gas extraction
 2. 1111A0 (111CA) Oilseed farming
 3. 324121 (324)   Asphalt paving mixture and block manufacturing
 4. 1121A0 (111CA) Beef cattle ranching and farming, including feedlots and
                   dual-purpose ranching and farming
 5. 112300 (111CA) Poultry and egg production
 6. 311230 (311FT) Breakfast cereal manufacturing
 7. 524200 (524)   Insurance agencies, brokerages, and related activities
 8. 311700 (311FT) Seafood product preparation and packaging
 9. 31161A (311FT) Animal (except poultry) slaughtering, rendering, and processing
10. 531HST (HS)    Tenant-occupied housing
11. 112120 (111CA) Dairy cattle and milk production
```

Figure 3: Ranked distribution of mean KL divergence per unit of shock in the 2017 detailed benchmark intermediate-use matrix. Sectors are ordered from highest to lowest value. Colored markers distinguish the top three detailed sectors, the next band through rank 11, and the remaining sectors. The list below the plot reports the highest-ranked detailed sectors together with their corresponding 73-sector codes.

The 2017 detailed ranking shows the same pattern at a finer level. In the ranked distribution of mean KL divergence per unit of shock, three detailed sectors stand clearly above the rest. These sectors are 211000, Oil and gas extraction, 1111A0, Oilseed farming, and 324121, Asphalt paving mixture and block manufacturing. They are followed by a second band of highly ranked sectors extending through about rank 11, including 1121A0, Beef cattle ranching and farming, 112300, Poultry and egg production, 311230, Breakfast cereal manufacturing, 524200, Insurance agencies, brokerages, and related activities, 311700, Seafood product preparation and packaging, 31161A, Animal (except poultry) slaughtering, rendering, and processing, 531HST, Tenant-occupied housing, and 112120, Dairy cattle and milk production. The visible drops after ranks 3 and 11 separate a very small top tier from a second band of highly ranked sectors, after which values decline more gradually across the rest of the ranking.

The mapping from 73 sectors into 402 sectors shows that not all highly ranked 73-sector sectors are built in the same way. Some are anchored by one or a few very highly ranked detailed industries. Oil and gas extraction is the clearest case, since sector 211 maps directly to 211000, which remains at rank 1 in the detailed ranking. Others spread into a much wider range of detailed industries. Farms, for example, contains some very highly ranked detailed sectors, including Oilseed farming, Beef cattle ranching and farming, and Poultry and egg production, but also many others that appear much lower in the detailed ranking. Food and beverage and tobacco products shows a similar pattern. The 402-sector exercise therefore shows which 73-sector leaders are associated with a narrow core of detailed industries and which reflect a wider and more uneven set of detailed industries.

3.3 Deformation rises disproportionately with shock size

To assess whether deformation rises proportionally with shock size, we evaluate raw KL divergence at equal steps from 2 to 10 percent, in increments of 2 percentage points, all applied to the same benchmark intermediate-use matrix. Using the 15-sector classification, we find that mean raw KL divergence rises more than proportionally with shock size. Figure 4 shows this convexity clearly, with the strongest responses in Manufacturing, followed by FIRE and Professional and business services, and weaker responses in other sectors. At the aggregate level, larger shocks require disproportionately more structural reorganization of the benchmark intermediate-use matrix in order to satisfy the new accounting constraints while remaining close to the original structure. This heterogeneity in nonlinear response across sectors is already present before prices, substitution, or equilibrium adjustment are introduced.

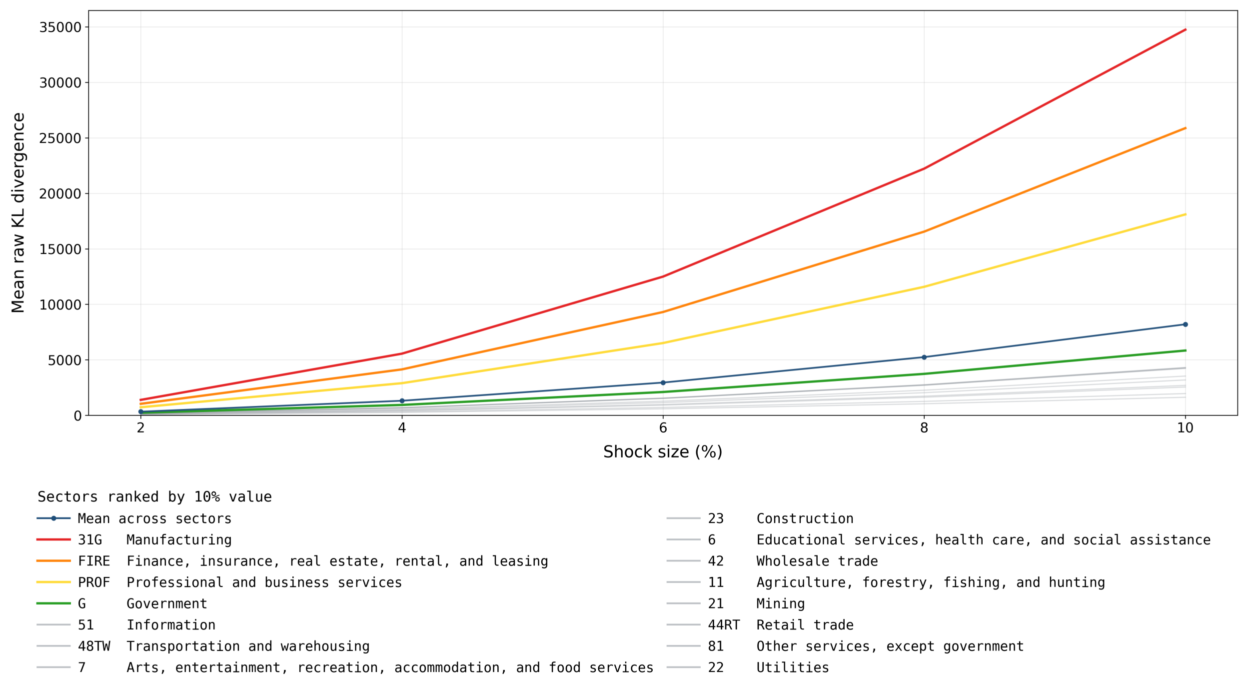

Figure 4: Shock-size variation in mean raw KL divergence under the 15-sector classification, using equal-step shocks of 2, 4, 6, 8, and 10 percent. Sectors listed below the plot are ordered by their values at the 10 percent shock size. Colored lines indicate the highest-ranked sectors, while the remaining sectors are shown in gray. The aggregate mean across sectors is reported by the line with markers.

Taken together with the multi-year continuity and sectoral disaggregation results in the preceding sections, these findings suggest that the benchmark intermediate-use matrix already contains a persistent and nonlinear structural layer before any CGE behavioral adjustment is imposed. Directional structural tightness is not an incidental feature of a single benchmark table. It persists across benchmark years, becomes sharper under finer sectoral aggregation, and intensifies disproportionately as shocks grow.