

China is more exposed to semiconductor supply-chain disruptions than the US

2026-10-06

Contents

Overview	1
Introduction	1
Data	2
Analysis	3
Limitations	5
Conclusion	6
About the authors	6
Related topics	7

Overview

Semiconductor production is extremely concentrated. Most advanced logic chips, including those that power AI, are fabricated in Taiwan, and a large share of memory is produced in South Korea, so disruptions to semiconductor supply chains are a serious concern. However, standard international input-output data groups chips with broader electronics data, making it hard to understand which countries would be most affected by disruptions. We disaggregate semiconductors in international input-output data to better measure each country's exposure and to simulate the impact of a Taiwan supply shock and a China–West decoupling.

- **China is more exposed than the US to semiconductor supply disruptions, because semiconductors account for a larger share of the cost of the goods it buys.** In 2022, every \$1,000 of Chinese final demand generated \$15.2 in revenue for semiconductor producers, $2.7\times$ the \$5.7 generated by the same American spending.
- **A disruption in Taiwan combined with a China–West trade decoupling would hit China several times harder than the US.** In our model's combined shock scenario, which estimates impacts two to three years after shocks occur, Chinese advanced processor prices rise 17-fold and Chinese real gross national expenditure falls by 3%, against a roughly 20% price rise and a 0.6% fall for the US.
- **The gap survives our robustness checks.** Alternative parameter choices and data splits move our estimates. But while the exact size of the gap carries significant uncertainty, China's exposure and simulated losses exceed those of the US in every case.

Introduction

Semiconductors are essential inputs to servers and data centers, computers, phones, cars, and industrial equipment, and their production is extremely concentrated: most advanced logic chips are fabricated in Taiwan and a large share of memory in South Korea. That makes the industry a natural place to worry about disruption.

But chips are hard to study. They are a small share of the cost of the goods they end up in, and they sit high up in production chains that cross many borders. If the cost of Taiwanese fabrication went up, that might make iPhones more expensive in the US, but only after the affected chips have been packaged and tested, shipped to contract manufacturers in China or India, built into finished phones, and exported to the US.

A standard tool for tracing this kind of chain is an international input-output (IO) table, which records what each industry buys from every other industry across multiple countries. However, most international tables bundle semiconductors into a broader electronics category that is too coarse to trace the chips themselves.

In this study, we separate out the semiconductor industry from this broader category in the OECD’s Inter-Country Input-Output (ICIO) tables, adapting and expanding on an approach [developed](#) by the OECD team of Haramboure, Lalanne, Schwellnus, and Guilhoto. We also add granularity by imputing chip types, such as processors and memory, that make up cross-border flows. These updated tables let us trace semiconductors from where they are produced to the countries whose spending ultimately pays for them. We can then estimate a country’s exposure to semiconductor supply-chain disruptions as the revenue semiconductor producers earn for every dollar of that country’s final demand, the sum of its consumption and investment in goods and services, whether made at home or imported. The more of a country’s final demand, directly or indirectly, goes to chipmakers, the more its consumption and investment will be affected when chips become scarce or expensive.

The tables also let us calibrate trade models and estimate the impacts of supply-chain shocks, such as trade wars or natural disasters. We return to this in the Analysis section below.

Lastly, our estimates extend only into 2022, where the underlying data ends. Thus, we lack estimates for 2023–present, during which US investment in AI servers has grown dramatically. We return to this in the Limitations section.

Data

Multi-region input-output tables, such as the OECD’s ICIO tables, map transactions between industries across multiple regions. In the ICIO’s case, for each year it gives us a $4,050 \times 4,050$ matrix, covering 50 industries $\times$ 80 countries and a rest of world region, with each cell telling us the nominal sale of intermediate goods and services one country’s industry sold to another in that year. For example, the ICIO estimates that in 2022 the Senegalese “refined petroleum” industry spent \$329 million on intermediate inputs from the Nigerian “oil and gas extraction” industry. The tables also provide a series of vectors that break down final demand in each country across industries (both domestic and foreign), and show how much each industry paid in tax and how much value-added it generated.

Economists use these tables to answer empirical questions about global value chains. For example, by tracing purchases backward or forward through the production network, we can work out (1) how much sales in one sector rise when spending increases in another and (2) how much one sector’s costs depend on prices in another. Consider each in turn:

1. For every \$1,000 of US private consumption in 2022, \$1.50 went directly to Chinese electronics manufacturers, for example, when an American bought a Chinese-made television. But Americans also bought cars, appliances, and other goods that contain Chinese electronic intermediate inputs. Counting those too, each \$1,000 of US consumption generated about \$6 of revenue for the Chinese electronics industry.
2. The same logic runs in the other direction. If Chinese electronics become more expensive, how much of an American’s spending is affected? Directly, only 0.15% of the US consumption basket is Chinese electronics. But the price of every good or service that uses Chinese electronics as an input rises too, so the share of the basket exposed to the price increase, assuming full pass-through, is $4 \times$ higher, at 0.6%.

Both calculations assume that each industry keeps spending the same share of its budget on each input, which is reasonable for small changes in spending or prices. But when a disruption is large enough that buyers switch suppliers and those shares move, these simple calculations are no longer enough, and economists turn to more explicit models of how industries respond to price changes.

This kind of tracking is particularly important for semiconductor production chains. Consumers and investors rarely buy chips directly, but chips are found in servers, computers, mobile phones, cars, and other consumer goods.

The difficulty is that the ICIO cannot see semiconductors. To cover 80 countries, it splits each economy into only 50 industries. Semiconductor production sits inside a category for electronic components, which in turn sits inside a broader category for computer, electronic, and optical products, and only that broadest category appears in the ICIO. Haramboure et al. used detailed national input-output tables for key countries in the semiconductor supply chain, together with customs data, to split this broad industry into semiconductor production and other electronics manufacturing.

We update their method, creating a semiconductor-augmented ICIO for 2016–2022. We use the OECD’s bilateral customs data (BIMTS) and South Korea’s national customs data to proxy a further split of semiconductor products across processors, memory, and other chips. This product-level detail is only visible where chips cross a border, so we cannot see it for chips made and used within the same country.

These tables are only estimates. The ICIOs themselves involve significant imputation and balancing, and the semiconductor split of Haramboure et al. involves even more, so we measure the size of sectors and trends against external estimates, and flag and adjust when they differ. Despite these imputations, the tables give us a unique look at how semiconductor production chains fit into the international economy.

Analysis

Chinese final demand generates far more revenue for the semiconductor industries of Taiwan, South Korea, and Japan than American final demand does. We measure a country’s exposure to a national semiconductor industry as the revenue that industry earns for each dollar of the country’s final demand. By 2022, China was more exposed than the US to the semiconductor industries of Taiwan, South Korea, and Japan by factors of 4.7, 4.0, and 3.6, respectively. From 2016 to 2022, Chinese exposure to global semiconductor manufacturing remained above $2.5 \times$ US exposure. This exposure is shown in the figure below.

The period was also one of rapid growth for China’s own semiconductor industry. The share of Chinese semiconductor demand (direct and indirect) met by Chinese producers rose from 12% in 2016 to 23% in 2022. Globally, nominal sales of the semiconductor manufacturing industry grew 59% over the same period; by 2022, every \$1,000 of global final demand generated \$8.80 in industry sales, of which \$1.85 went to China’s producers, \$1.57 to Taiwan’s, and \$1.15 to South Korea’s.

Simulating a Taiwan shock and decoupling


To study large shocks, we need more detail than the table alone gives. Processors are not all alike: the advanced chips (and their components) that run AI servers come almost entirely from a handful of fabs in Taiwan, South Korea, and the United States, while mature chips are made in many more places. So we make a rough split of each country’s processor sales into advanced and mature, and separate out the South Korean memory line that high-bandwidth memory ships under. With those additions, the 2022 table calibrates a general-equilibrium trade model, which we use to simulate a large fall in Taiwanese productivity, a full trade decoupling between China and the West, and the two together.

Chinese spending relies on semiconductors about 2.7 times as much as American spending

Semiconductor output required, USD per \$1,000 spent

China Taiwan South Korea Other

USD per \$1,000 spent



EPOCH AI | CC-BY

epoch.ai

Figure 1: Stacked bar chart of semiconductor output required per \$1,000 of spending in the US and China, split by producer: China, Taiwan, South Korea and other. The US needed \$4.9 in 2016 and \$5.7 in 2022, and China needed \$14.9 and \$15.2, about three times and 2.7 times as much.

In the combined scenario, at our baseline calibration meant to represent a 2–3 year adjustment after the shock, China experiences a 17-fold rise in advanced processor prices, compared with about a 20% rise for South Korea and the US. Chinese real gross national expenditure (GNE) falls by 3% (or about 1% if the trade decoupling only affects semiconductor sales), while US real GNE falls by only about 0.6% (or nothing).

These results are sensitive to key parameters. The GNE numbers move most with how readily buyers switch between countries' goods in general (halving our assumed ease of switching raises China's loss to 8.3% and America's to 2.0%); the chip-specific choices, such as how readily advanced and mature processors substitute, move China's loss by a few points (rising to about 6% when we make that substitution harder but keep it plausible) and the US loss by almost nothing. Price effects are much less robust than GNE effects, and China's advanced-processor price, in particular, ranges from a 6% rise to a 100-fold increase across our robustness checks in the paper's appendix. Across them, China's losses consistently exceed America's.

Even with these checks, the model captures only one channel through which a disruption does damage: Goods become scarcer and more expensive, and buyers switch suppliers. It leaves out other factors, such as a recession triggered by the shock, financial fallout like sanctions and capital flight, and bottlenecks at individual firms that our aggregate data cannot see. Most of these would add to the cost, while chip inventories, which the model also ignores, would cushion buyers for a time. The figures are best read as the trade channel of a disruption over a two-to-three-year horizon.

Limitations

Two further limitations apply to our analysis.

The first limitation is timing. The tables stop in 2022 because the underlying data does, but spending on chips has changed and grown since.

To explore this, we apply our method to the Asian Development Bank's multiregional input-output tables for 2022 and 2024, which are more aggregated than the OECD's ICIO. On these tables, China's final-demand exposure to the global semiconductor industry falls from $2.8\times$ to $2.2\times$ US exposure, and, in the combined shock scenario of our simulations run on the 2024 table, China's real GNE loss falls by about 0.4 percentage points while the US loss falls by less than 0.1 percentage points. These tables do not separate electronics from electrical equipment, and the 2024 table must rest on significant imputation, so they should be read with caution, but they suggest that the direction of the paper's main findings holds through 2024. Given the rise in spending on AI servers after 2024, updating the tables as new data becomes available is a priority of ours.

The second limitation is aggregation. Each industry in the tables adds up many different firms, but the tables give them all the same mix of inputs. A Chinese firm that assembles phones for export using Taiwanese chips sits in the same industry as a camera maker selling to Chinese households. The tables see only their sum, a single industry that buys chips from Taiwan, exports some of its output, and sells some to Chinese final demand. Because China has many such assemblers, this likely overstates its exposure.

To test this, we rebuilt the 2022 table using the OECD's split of Chinese manufacturing into processing exporters (firms that import inputs, transform them, and export all their output) and other firms. The split narrows the gap between Chinese and US exposure but fails to close it. China's final-demand exposure to the global semiconductor industry falls to between $2.4\times$ and $2.6\times$ US exposure overall, between $4.2\times$ and $4.4\times$ for Taiwan's chips, and between $3.5\times$ and $3.8\times$ for South Korea's. The split is still coarse, and settling the question needs finer input-output data, ideally at the firm or plant level.

Conclusion

Even with these limitations, the picture is consistent across tables and the modeled scenarios. China buys far more of the world's semiconductors per dollar of final demand than the United States does, and its own capacity in advanced processors and stacked memory **remains well below** that of the US and its allies. That combination leaves it more exposed, both to disruption in Taiwan and to a trade decoupling. The full method, data, and results are in the **paper**.

About the authors



Daniel Carey

Daniel Carey is a researcher at Epoch AI. He is interested in the economics of transformative AI.

Related topics

[Chips](#)[Geopolitics](#)[Economic impact](#)