NBER WORKING PAPER SERIES

INDUSTRIAL POLICY IN THE GLOBAL SEMICONDUCTOR SECTOR

Pinelopi K. Goldberg
Réka Juhász
Nathan J. Lane
Giulia Lo Forte
Jeff Thurk

Working Paper 32651
http://www.nber.org/papers/w32651

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2024, revised February 2026

We are grateful to Ralph Ossa, Bob Staiger, and Dan Trefler for suggesting this project. We thank Panle Barwick, Chad Bown, Simon Evenett, Chang-Tai Hsieh, Brian Kovak, Ryan Michaels and John Minnich for helpful conversations, Pranay Kotasthane and Abhishek Krishnan for sharing their expertise on industrial policy in semiconductors in India, Shrikant Lohokare and Jessica Mueller at the Global Semiconductor Alliance for providing the data used to estimate the model, and participants of the "Expert Dialog on Subsidy Reform" Conference in Washington, DC, as well as several seminar participants for many useful comments. We are grateful to Kisho Hoshi, Richard Jung, Stephen Xu, and Esha Vaze for their excellent research assistance. Réka and Nathan gratefully acknowledge funding from the Alfred P. Sloan Foundation. The views expressed herein are those of the authors and do not necessarily reflect the views of views of the Board of Governors of the Federal Reserve System or any other person associated with the Federal Reserve System, or of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Pinelopi K. Goldberg, Réka Juhász, Nathan J. Lane, Giulia Lo Forte, and Jeff Thurk. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Industrial Policy in the Global Semiconductor Sector
Pinelopi K. Goldberg, Réka Juhász, Nathan J. Lane, Giulia Lo Forte, and Jeff Thurk
NBER Working Paper No. 32651
July 2024, revised February 2026
JEL No. F13, F61, L63, N60, O38

ABSTRACT

The resurgence of subsidies and industrial policies has raised concerns about their potential inefficiency and alignment with multilateral principles. Critics warn that such policies may divert resources to less efficient firms and provoke retaliatory measures from other countries, leading to a wasteful "subsidy race." However, subsidies for sectors with inherent cross-border externalities can have positive global effects. This paper examines these issues within the semiconductor industry: a key driver of economic growth and innovation with potentially significant learning-by-doing and strategic importance due to its dual-use applications.

Our study aims to: (1) document and quantify recent industrial policies in the global semiconductor sector, (2) explore the rationale behind these policies, and (3) evaluate their economic impacts, particularly their cross-border effects, and compatibility with multilateral principles. We employ historical analysis, natural language processing, and a model-based approach to measure government support and its impacts. Our findings indicate that government support has been vital for the industry's growth, with subsidies being the primary form of support. They also highlight the importance of cross-border technology transfers through FDI, business and research collaborations, and technology licensing. China, despite significant subsidies, does not stand out as an outlier compared to other countries, given its market size.

Model estimates suggest the presence of learning-by-doing at the firm-product level as well as economies of scope within a firm and substantial cross-border learning spillovers. These spillovers likely reflect cross-country technology transfers and the role of fabless clients and input suppliers in disseminating knowledge globally through their interactions with foundries. Such cross-border spillovers are not merely accidental but result from deliberate actions by market participants that cannot be taken for granted. Firms may choose to share knowledge across borders or restrict access to frontier technology, thereby excluding certain countries. Future research will use model estimates to simulate the quantitative implications of subsidies and to explore the dynamics of a "subsidy race" in the semiconductor industry.

Pinelopi K. Goldberg
Yale University
and NBER
penny.goldberg@yale.edu

Réka Juhász
University of British Columbia
Vancouver School of Economics
and NBER
reka.juhasz@ubc.ca

Nathan J. Lane
University of Oxford
Department of Economics
nathaniel.lane@economics.ox.ac.uk

Giulia Lo Forte
Board of Governors of the Federal
Reserve System
giulialoforte@gmail.com

Jeff Thurk
University of Georgia
Department of Economics
jeff.thurk@gmail.com

1 Introduction

The recent resurgence of subsidies and other forms of industrial policy is widely perceived as a significant departure from the principles of multilateralism and international cooperation. Economists are generally skeptical of industrial policy as it can lead to inefficiency by diverting resources to less efficient firms. Internationally, industrial policies may have negative cross-border effects if resources are reallocated towards less efficient domestic producers, prompting retaliation and similar policies from other countries. The resulting "subsidy race" can become a "race to the bottom," a wasteful competition of resources that fails to achieve policy goals.

However, some subsidies, such as those for new goods, green technologies, or sectors with inherent cross-border externalities, can have positive cross-border effects and benefit other countries. Assessing these cross-border effects is challenging when comparative advantage is dynamic. Subsidies today might have minimal immediate impact on other countries but could lead to significant future effects. Conversely, some countries may not currently use subsidies, yet their past subsidies have helped them secure a dominant position in the global market. This raises broader questions: (a) Can the new subsidies be economically justified? (b) Should current WTO subsidy rules be updated to address the complexities of new and rapidly evolving industries?

This paper is part of a larger project aimed at understanding these issues within key sectors of the global economy. It focuses on semiconductors, one of the most dynamic and globally integrated industries, heavily targeted by industrial policies worldwide. The semiconductor industry is a key driver of economic growth, enabling all facets of modern digital life and accelerating scientific innovation through high-performance computing. Learning-by-doing and dynamic comparative advantage are considered to be crucial features of the industry, leading to "infant industry protection" arguments in favor of government support. Additionally, the industry has strategic importance stemming from semiconductors being "dual-use" goods, with both civilian and military applications, prompting national security-based arguments for government intervention.

Our study has three primary objectives:

1. To document and quantify, where possible, the industrial policies that have been recently implemented in the global semiconductor sector.
2. To explore the rationale behind these policies.
3. To evaluate their economic impacts, particularly their cross-border effects, and assess their compatibility with multilateral principles.

Our first objective focuses on policy description and measurement. Accurate measurement of implemented policies is a crucial prerequisite for analyzing their effects. However, as our analysis will reveal, this task is far from straightforward. Policymakers and industry

experts frequently cite various figures in their statements and reports, but the sources of these numbers are often unclear, and the approaches used to measure government support are inconsistently applied across countries. An exception is the excellent study by the OECD (2019), which we incorporate into our discussion. To advance measurement, we combine three different approaches: (a) historical analysis of state support of the industry worldwide based on existing sources; (b) the natural language processing (NLP) approach pioneered by Juhász, Lane, Oehlsen, and Perez (2023) to identify industrial policies in the *Global Trade Alert* (*GTA*) database between 2010 and 2022; (c) a model-based approach to identify (unobserved) production subsidies as a residual factor lowering the marginal cost of production after controlling for all other relevant cost determinants.

We compare our findings with those of the OECD study throughout our analysis. The two studies differ in many important respects, both conceptually and in terms of firm and year coverage, so it is unsurprising that our findings vary in some dimensions. Nonetheless, a common conclusion of both studies is that government support in the semiconductor industry is exceedingly difficult to quantify due to the diverse instruments and targets used by different governments across the value chain. Measurement difficulties are particularly pronounced in the case of China. The *GTA* database does not cover financial support provided by the Chinese government, much of which is provided by subnational governments. One exception is China's "Big Fund" in 2014. Estimates from other sources, such as *JW Insights* reported in the *Financial Times*, do not break down support by year or firm, making economic analysis of the policies challenging. This is partly why we also consider the model-based approach, inspired by Kalouptsidi (2018a) and Barwick, Kalouptsidi, and Zahur (2023).

Our model-based approach involves specifying a model of the semiconductor industry, estimating key demand and cost parameters using firm- or industry-level data, and then using the model structure to back out marginal costs for each product. In this framework, production subsidies are identified as a residual factor that reduces the marginal cost of all firms in a country, i.e., a factor that leads to lower marginal costs in a country relative to a benchmark country after accounting for all other relevant cost determinants. The disadvantage of this approach is that the estimates depend on modeling and identification assumptions, though it is possible to test their robustness against alternative assumptions. Nevertheless, a significant advantage of this approach is that it allows not only for measurement, but also for the evaluation of the subsidies' economic effects, which is impossible without a theoretical framework. From a measurement perspective, the model-based approach is particularly useful when its findings can be contrasted and potentially corroborated with evidence from other sources. In our context, we use confidential data from the Global Semiconductor Alliance (GSA) to implement the model-based approach. However, the overlap between this data, which covers 2004-2019, and the *GTA* data, which covers 2010-2022, is limited, making the comparison less informative.

Despite these caveats, the combined use of the three approaches leads to some clear takeaways:

First, government support has been critical for the semiconductor industry's growth, particularly during its initial development phase. This support is evident across all major segments of the value chain, benefiting established leaders at the technology frontier, such as Korea and Taiwan, countries seeking to advance their industry, such as China and the U.S., and countries attempting to enter the market, such as India. In more mature markets, governments have traditionally allowed the private sector to take the lead. However, since 2020, there has been a significant increase in government intervention, with China, the United States, Japan, Korea, and India notably ramping up financial support for the industry.

Second, subsidies are the primary form of government support, manifesting as financial grants, state aid, tax incentives, loans and loan guarantees, and equity injections. This trend aligns with the findings of the OECD report. These policies primarily target production improvement and research, development and innovation.

Third, China has been a prominent user of subsidies. However, our estimates do not pinpoint China as an outlier in its subsidy use; rather, its level of support is comparable to other countries, when considering the size of its market. This conclusion is supported by both our NLP-based analysis of *GTA* data and the subsidy estimates from our model-based approach.

Fourth, cross-border technology transfer from more advanced to less advanced firms has been as crucial as state support for the industry's development. This transfer has occurred through foreign direct investment (FDI), research collaborations, and technology licensing. Outside of the United States, our analysis found no instance where a domestic semiconductor industry developed without substantial foreign technology. This underscores the difficulty of developing the industry without foreign partners willing to share technology. It also explains why China has struggled to reach the technological frontier despite pursuing similar policies to other Asian economies that, as U.S. allies, had better access to foreign technology.

Fifth, policymakers aim to achieve several goals through their support of the semiconductor industry, including economic growth and development, international competitiveness, resilience, and national security. Implicit in these objectives is the belief in strong learning-by-doing effects in chip manufacturing, which lead to dynamic comparative advantage, economies of scale, and high industry concentration. In the presence of learning-by-doing, subsidies have a multiplier effect, amplifying and accelerating cost reductions as experience accumulates. Even when firms internalize the benefits of learning spillovers, so that private production is socially optimal, governments may still have valid reasons to support their domestic semiconductor industry. This support helps counter the natural tendency toward industry concentration, diversify the supply chain, and enhance its resilience. These objectives become even more critical when considering national security concerns, given that crucial

segments of the semiconductor supply chain are concentrated in a few geopolitically critical countries. Learning spillovers across technologies and firms provides additional justification for subsidies.

Sixth, our estimates indicate strong learning-by-doing in line with industry lore (around 22%). We find substantial learning internal to each manufacturing facility (foundry) and product (technology), as well as across technologies within a foundry, pointing to economies of scope in learning. Finally, we find strong evidence of cross-border learning spillovers. When these international spillovers are not taken into account, the estimated learning effects appear to be much smaller. We show that these results are driven by the price patterns in the data rather than by modeling assumptions.

Although our data and model cannot pinpoint the sources of these international spillovers, we hypothesize that they stem from the highly fragmented yet concentrated nature of the global supply chain and the importance of firm-to-firm relationships between buyers, input suppliers, and manufacturers. Chip design and manufacturing require close cooperation between buyers, manufacturers and their input suppliers. Manufacturers pool orders across buyers to exploit economies of scale, leading to standardization of designs and processes. Buyers also place orders with manufacturers in different countries, thus intermediating knowledge acquired through interactions with manufacturers across borders. Another likely source of international spillovers is foreign technology transfer, including through foreign direct investment (FDI), research collaborations, or the cross-border recruitment of engineers and professionals. As previously noted, foreign technology transfer is a crucial driver of growth in the industry. Our findings on international spillovers are consistent with this premise.

These results have interesting implications for the cross-border effects of subsidies. International learning spillovers imply positive cross-border effects. However, since these spillovers may result from foreign technology transfers or business and research collaborations that require deliberate actions by market participants, they are neither automatic nor inevitable. Just as firms can share knowledge across borders, they can also restrict access to frontier technology, cutting off certain countries. In such cases, cross-border learning spillovers will be small, and non-subsidized firms in other countries will suffer from business stealing. Nevertheless, positive cross-border effects may still occur if subsidies lead to product innovation and lower costs at subsidized foundries, which benefit international buyers. In the future, we plan to use our model estimates to assess the quantitative implications of subsidies through counterfactual simulations. Additionally, we intend to explore the implications of a "subsidy race," where multiple countries subsidize their semiconductor industries. This will help us determine whether subsidies act as global complements or substitutes when learning-by-doing is present.

The remainder of the paper is structured as follows. Section 2 provides a brief overview of the industry. Section 3 summarizes the history of industrial policy in the semiconductor sector

to offer context and background for the policies we analyze later. Section 4 presents the main findings from the NLP-based analysis of the *GTA* database and compares them with the results of the OECD report. Section 5 outlines the primary objectives pursued by policymakers as identified in the *GTA* database. Section 6 introduces the model-based approach, the additional proprietary data used to estimate the model, the rationale for using subsidies, and presents the empirical findings. Section 7 concludes.

2 The Semiconductor Industry

Semiconductor firms are firms engaged in the design and/or fabrication of semiconductors—any material whose electrical conductivity ranges between that of a conductor and that of an insulator. The bulk of industry revenue is accounted for by integrated circuits (ICs), which are networks of transistors fabricated on a surface to process binary data by switching them on and off. The semiconductor industry forms the backbone of the hi-tech industry and is considered a prime driver of economic growth.

2.1 The manufacturing process

Semiconductor fabrication begins with designing chips using software that lays out and simulates transistor networks on semiconducting material. Manufacturing occurs in a fabrication facility, colloquially known as a "fab." Transistors are etched onto a wafer via photolithography, a process akin to photographic development that repeats to form multiple chip copies. After layers are added, the wafer undergoes "wafer probe" testing to identify defects. Successful chips are then cut from the wafer, and each of these pieces is called a "die." The dies are then encased and prepared for integration with other components.

Technological advancements in fabrication manifest in a number of product characteristics; these include increased wafer sizes and reduced line widths, enhancing the number of chips per wafer, and the density of transistors on each chip. For instance, larger wafers reduce costs per die (a single chip cut from a wafer) by about 30%, while a 30% reduction in line width roughly doubles transistor density, making chips smaller, faster, and more efficient. Despite initial higher costs and the increased potential of defects, refinements in fabrication confer benefits such as increased output and performance.

2.2 The semiconductor supply chain

Table 1 summarizes the four broad categories of products within the semiconductor industry. First, firms such as Intel and Advanced Micro Devices (AMD) produce microprocessors,

which amount to integrated circuits (ICs) containing one or more central processing units (CPUs). This first category of products are used in personal computers, tablets, and servers. The second category, system on a chip (SoC), is the newest type of semiconductor chip, which combines all the necessary components of an entire system on a single chip. These products are popularly used in small devices, such as smartphones, as they integrate CPUs with graphics, cameras, as well as audio and video processing. Key firms in the SoC market include names such as Nvidia, Broadcom, and Qualcomm.[1] The third category includes commodity integrated circuits commonly used in simple technological devices, such as barcode scanners. Finally, the fourth category includes memory chips, particularly flash memory, which are produced by large conglomerates such as IBM and Samsung.

Table 1: Categories of Semiconductor Products and Major Firms

Semiconductor Products	Example Firms	Example End-Products
1. Microprocessors	AMD, Intel	Computers, servers
2. "System on a Chip"	Broadcom, Nvidia, Qualcomm	Mobile phones
3. Commodity integrated circuits	Analog Devices, Xilinx	Barcode scanners
4. Memory	IBM, Samsung, Toshiba	Computers, flash drives

Comparative advantage comes through innovation, and innovation is a fast-paced, cumulative effort in which tomorrow's new product depends heavily on a broad set of today's products and ideas. In a 1965 paper, Intel co-founder Gordon E. Moore noted that the capabilities of the integrated circuit doubled roughly every 18 months.[2] This prediction is popularly known as "Moore's Law" and has held up remarkably well for over 50 years since its inception. Additionally, Moore's Law speaks to the short life span of any current product and the need to develop tomorrow's great idea today. Accordingly, R&D comprises a significant—and increasing—component of firm expenses.

There are four primary components of the value chain: design, fabrication, testing, and sales/distribution. During the design stage, skilled design engineers construct prototypes of next-generation chips using expensive, high-end electronic design automation (EDA) software. When completed, these plans are delivered to a (potentially external) fabrication facility, where the chip circuits are constructed in successive layers on the surface of flat silicon wafers. Firms engaged in the fabrication stage must incur a large fixed capital investment ($\approx$\$2 billion) to build a plant, or fab, which consists of a wide variety of expensive equipment capable of

1. To illustrate the importance of these products: In early 2018, Singapore-based Broadcom attempted a hostile takeover of Qualcomm ($120 billion), which was later blocked by the Trump Administration due to national security concerns. Intel has since expressed interest in acquiring Qualcomm to solidify its position to deliver 5G mobile services in the future.
2. At the time, he was referring to the number of transistors a firm could inexpensively place on a single silicon wafer. Today, advancement generally refers more generically to processor speed.

building the chips under extreme environmental requirements for cleanliness. During the assembly stage, the wafers are split into individual dies for distribution to customers.

During the 1970s and 1980s, the industry was dominated by vertically integrated device manufacturers (IDMs), which managed all components of the value chain (design, fabrication, testing, and distribution). The establishment of Taiwan Semiconductor Manufacturing Company (TSMC) in 1987 by the Taiwanese government marked the beginning of an alternative business model, where production was outsourced to low-cost, third-party foundries. Firms choosing this business model lack an internal fabrication facility or follow a "fabless" model. Today, fabless firms account for roughly 90% of all semiconductor firms and generate one-third of industry revenue. The fabless business model provides two advantages to the traditional IDM. First, outsourcing fabrication enables these firms to avoid the substantial capital investment required to build fabrication facilities. Second, outsourcing—often overseas—enables these firms to take advantage of scale to lower input costs, pooling production with other fabless firms in third-party foundries. Outsourcing overseas also enables firms to benefit from lower foreign wages and weaker environmental standards. Thus, the fabless business model is thought to decrease upfront costs and ongoing production costs.

Outsourcing fabrication is particularly attractive in this industry as semiconductors tend to be high value, weigh little, and thus have low transportation costs. As many low-wage countries lack the technical expertise and capital infrastructure to establish viable foundries, the majority of foundries are located in a specific set of markets: mainly China, Europe, Japan, Taiwan, and the United States. The vast majority of outsourcing is done in Asia (Taiwan accounts for approximately 59% of all outsourced wafers produced), while U.S. foundries account for four percent of all third-party wafers produced. Hence, global supply chains have played an important, though not necessarily pivotal, role in facilitating growth of the fabless business model.

Today, the semiconductor supply chain is globally integrated, with different countries dominating different parts of the supply chain. Advanced economies such as the United States and the Netherlands lead in research and design. East Asia, particularly Taiwan and China, dominates fabrication, producing the vast majority of wafers in "pure play" foundries. Testing, assembly, and distribution are primarily controlled by East Asian countries (OECD, 2019; Bown, 2020). Each of these segments is highly concentrated, so that while manufacturing a chip requires cooperation among several firms across different countries, there is minimal competition at each stage of the process.[3] The global nature of the semiconductor chain has important implications for the assessment of subsidies or any other form of government support by individual countries, as these actions resonate through global linkages across the

3. For instance, in the EDA software market, three US-based companies – Synopsys, Cadence Design Systems, and Mentor Graphics — control ca. 85% of the global market (Bown, 2020).

entire supply chain. In other words, cross-border effects of industrial policies are inherent in this industry, and the key question is whether these effects are positive or negative.

2.3 Learning by doing

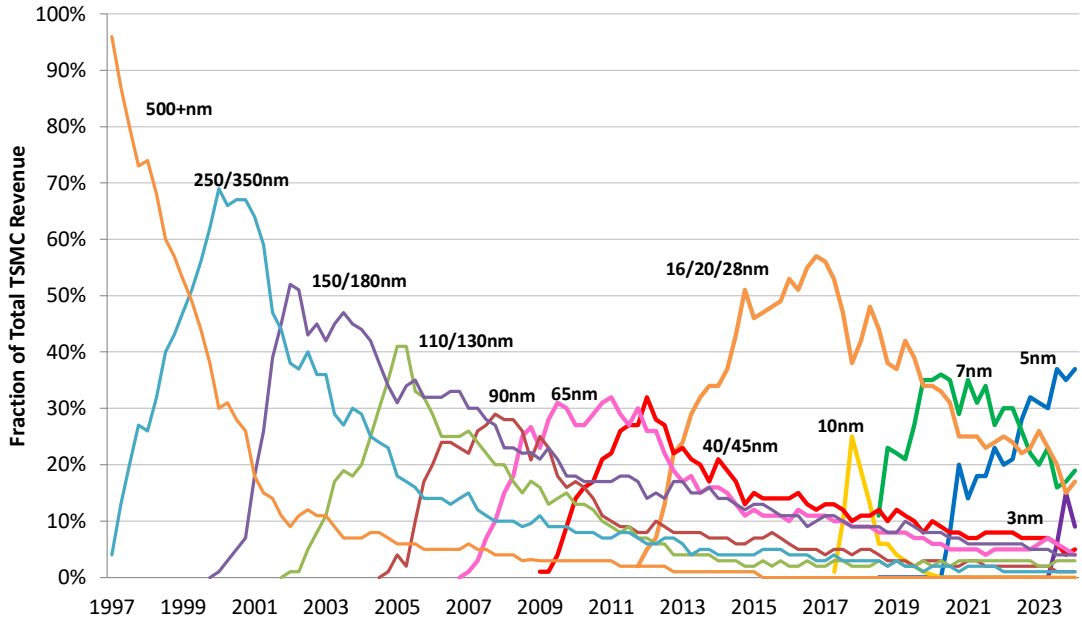
Semiconductor manufacturing includes several intricate steps, such as wafer fabrication, lithography, etching, and doping, each demanding precise control. Given the complexities and microscopic scale of these processes, mastering them requires more than just theoretical knowledge; it necessitates practical, hands-on experience. The fast-paced nature of the industry means that there is a large opportunity for continuous process improvement via the accumulation of knowledge through hands-on experience in manufacturing operations.

Learning by doing can show up in several areas. First, there is scope to optimize operational processes. Regular involvement in fabrication processes helps engineers and technicians identify process inefficiencies and subtle variations, which are often missed in theoretical training. This practical experience is vital for developing more efficient manufacturing techniques, enhancing yield rates and product quality. Second, semiconductor manufacturing firms – both IDMs and fabless – promote a culture of innovation in which production floor personnel are expected to experiment with new ideas in real-time. This hands-on adaptability is crucial in an industry characterized by rapid technological changes and shifting market demands. Third, semiconductor firms are constantly re-evaluating and improving worker skills. The advanced skills required to manage and optimize cutting-edge semiconductor manufacturing equipment are best developed through on-the-job learning, enhancing the technical proficiency of the workforce.

Intel Corporation, for example, has integrated learning by doing into its operational ethos, enabling it to maintain leadership in process technology. Specifically, Intel's development of its 14-nanometer and 10-nanometer technology nodes involved extensive iterative testing and process refinement. This experiential approach allowed the company to overcome significant manufacturing challenges and achieve high yield efficiencies, which were crucial during the competitive surge from rivals like AMD and Global Foundries.

Similarly, TSMC actively promotes rapid learning cycles in semiconductor manufacturing. As the world's leading semiconductor foundry, TSMC mastered the art of quick iteration across its advanced technology nodes, notably its 7-nanometer and 5-nanometer processes. From Figure 1 we observe that TSMC operates several line-widths at a given time and that the company quickly ramps-up production followed by a long period of reduced sales as a percent of total TSMC revenue. TSMC's approach involves close collaboration with equipment suppliers and customers to speed up problem-solving and process optimization, resulting in high yield rates and enabling faster commercialization of new technologies. This focus on

Figure 1: TSMC Sales by Line Width



continuous improvement through hands-on problem solving has enabled TSMC's market dominance and attractiveness to global clients like Apple and NVIDIA.

3 A Brief History of Industrial Policy in Semiconductors

Historically, governments have played an important role in shaping the semiconductor industry. In this section, we describe the different ways in which governments have facilitated the *establishment* of the industry, both at the technology frontier and in follower countries. The historical evidence suggests that the early phase of the industry's development may be the period where the government has the most important role to play.

We start with the U.S., the birthplace of the industry. Although most accounts of the industry's early days focus on the scientists and entrepreneurs who developed and commercialized the technology, the government was also involved in numerous ways. First, through the 1960s, the US military bought most integrated circuits produced (72% of all integrated circuits produced in 1965 (Miller, 2022, p. 31)). This major source of demand allowed producers to learn how to commercialize and mass produce a technology that was still in its infancy.

More subtly, large-scale government procurement meant that public agencies played an outsized role in shaping *how* the industry developed. As one example, the U.S. Army, Navy, and Air Force funded distinct approaches to solving the technical problem of soldering together more transistors before silicon integrated circuits emerged as the best solution (Berger, Khan, Schrank, and Fuchs, 2023). Importantly, the US military required that contractors

share technology, a policy that promoted the flow of knowledge. This, together with rigorous antitrust policies pursued in the postwar period, allowed for the rapid dissemination of technology (Choi, 2008; Berger et al., 2023).

For technological follower countries, governments have used a variety of policy levers to foster a domestic semiconductor industry. Particularly in those economies where the technology gap was large, the state played a dominant and sometimes leading role in establishing the industry. This is seen to different degrees in all the economies which subsequently became major producers at the technology frontier (Japan, South Korea and Taiwan), in economies such as China that have pursued industrial policy in the sector for decades with mixed results, and also in economies such as Israel (Breznitz, 2007), which successfully entered the industry, but command a smaller share of the global market. This is important, as while we will focus our discussion on the establishment of the industry in East Asian economies, the economic forces at play seem to extend beyond "East Asia specific" factors.

We now turn to East Asian economies which have succeeded in catching up with the technology frontier. Although each pursued the objective of establishing a semiconductor industry differently, their policies also share broad similarities. Crucially, beyond making the industry more attractive to entrepreneurs by changing relative prices (through tariffs, subsidized loans, tax breaks etc.), each state was directly involved in the process of international technology acquisition, absorption, and diffusion. Put differently, developing a domestic industry seems to have required some infant industry promotion and, critically, also extensive international technology transfer facilitated by public agencies.[4]

In Japan, the famous pilot agency, MITI (Japanese Ministry of International Trade and Industry), directly controlled technology import policy. As Japanese firms began entering semiconductors in the 1950s, MITI determined which firms received a technology license and, once approved, the government guaranteed that payments would be made (Lynn, 1998). There were a number of potential advantages to Japan's coordination. First, in the 1950s, Japan was emerging from two decades of isolation from the world economy. Its firms were behind the technology frontier, had few foreign links, and limited experience with international agreements. Thus, by coordinating technology licensing directly, the government pooled knowledge and lent the credibility of the Japanese government to firms that were unknown to the West. The weight of the Japanese state arguably also helped firms' bargaining power; in 40% of all cases submitted for approval, the government intervened and forced changes to the agreement that favored Japanese firms (Lynn, 1998; Choi, 2008).

In Korea, the government began pushing the electronics industry to move beyond the assembly stage in its Fourth Five-Year Development Plan (1977-1981), when its per-capita

4. Another important element was the availability of human capital. Taiwan established a semiconductor laboratory to study microelectronic physics when it only had semiconductor assembly capabilities (Liu, 1993).

income stood at $459, roughly on par with Guatemala's (Amsden, 1989). Semiconductors were one of the strategic products named in this plan; according to Amsden, beyond the usual incentive package for prioritized sectors, the government established an industrial estate for the production of semiconductors and computers, it protected the domestic market, and, most famously, it established the Electronics and Telecommunication Research Institute (ETRI).

In Korea, ETRI would play a key role in working together with the private sector to coordinate R&D efforts between public researchers and the private sector. In fact, the government accounted for the vast majority of the R&D spending in semiconductor technology through the 1980s, with private firms contributing a smaller share (Chen and Sewell, 1996, Table 5). At least as importantly however, through ETRI, the state pushed the *chaebol*, large conglomerates, to upgrade their production and move into more design-intensive, higher-return markets (Evans, 1992). In some cases, it did so by subsidizing and coordinating efforts by the private sector; in others, it undertook the basic research itself, working closely with the private sector (Evans, 1992).

One case in particular illustrates the extent to which public agencies took the lead in pushing the domestic technological frontier: electronic switching systems.[5] The Ministry of Communications (MOC) decided to develop indigenous switching technology, a hugely risky proposition that private firms would not have been willing to undertake on their own. The MOC however used its control of telecommunications public procurement to persuade foreign firms supplying the Korean market such as Ericsson to transfer some of their switching technology (a form of *quid pro quo*), and train ETRI personnel. ETRI then used this knowledge, as well as generous funding from the MOC, to design an electronic switching system that was better suited to developing country contexts with large rural populations. Importantly, however, private firms were intimately involved with ETRI at the research stage, and production was undertaken by private firms from the very beginning.

As this case illustrates, much like Japan, Korea did not develop the technology in isolation. Instead, it relied extensively on technology alliances with foreign, mostly US, firms. Chen and Sewell (1996) document how each of the three *chaebol* involved in semiconductors (Samsung, Goldstar and Hyundai) had multiple technology alliances for a variety of products with foreign firms (the authors list twenty-two distinct alliances for three years between 1983-1985). Moreover, each acquired US based firms "providing them with direct access to highly qualified scientists and engineers, advanced technologies and major markets" (Chen and Sewell, 1996, p. 765).

While Korea entered the industry at a relatively low level of economic development, its concentrated industrial structure meant that there were large, incumbent private sector firms that the state could work with to develop the industry. Taiwan also entered the industry at a similar level of development, yet its economy was characterized by small and medium-sized

5. This paragraph is based on the account in Evans (1992, p. 142).

firms. Thus, among the three cases we examine, the government's singular role in establishing the industry is most evident here. Much like Korea and Japan, however, Taiwan also relied extensively on foreign technology at the outset.

Liu (1993) denotes three stages in the initial development of the industry from the mid-1970s. First, through the newly founded Electronics Industry Research Center (reorganized later as Electronics Research and Service Organization, ESRO) the government set up a public pilot plant that developed technologies increasingly close to the frontier (in cooperation with RCA, a US-based firm). In this way, the government assumed the entirety of the risk of high-tech R&D. Over time, ESRO moved into all areas of the supply chain including computer-aided design (CAD) and IC mask making. Later, in the 1980s, it closed the technology gap further by moving into Very Large Scale Integration Technology (VLSI) – i.e., more than 100,000 elements in a chip.

In the second stage, the government transferred the technology it had developed at ESRO to the private sector through licensing, spin-offs (including both United Microelectronics Corp. and Taiwan Semiconductor Manufacturing Co.) and technology diffusion policies. In the third stage, the government provided a variety of incentives to encourage the private sector to enter the industry. Like Korea, the government set up an industrial park and also offered a large range of generous fiscal incentives (cheap loans, tax breaks, accelerated depreciation of R&D equipment and low-cost land).

These historical accounts suggest that the government was instrumental, and, in some cases, the driving force in the initial development of the industry in successful follower countries. However, observers agree that as the industries matured, government involvement took more of backseat to the private sector, "shift[ing] from that of father to friend" (Liu, 1993, p. 304). This is worth noting as we move to measuring more contemporary industrial policy in some of these same countries below in Section 4.

It is important to highlight the likely critical role US policymakers and firms played in making the technology available to these economies (all of whom were key US allies in the Cold War context). The technology became accessible to foreign firms partly because of U.S. antitrust policies (Choi, 2008), and partly because U.S. firms themselves were willing to share technology. This was not a given, and some authors note that Japanese firms shared much less technology with Korean firms in anticipation of future competitive threats (Chen and Sewell, 1996). Given the importance of international technology acquisition for all three economies, it is difficult to imagine them achieving similar success without some access to frontier technologies.

We conclude this section by examining the history of semiconductor industrial policy in China. Like other East Asian economies, China has sought to incubate a domestic semicon-

ductor industry for decades, with the earliest efforts dating back to the 1960s.[6] Following Mao's death and the economic reform pursued by Deng Xiaoping, the Chinese leadership has consistently prioritized, guided, and generously financed the development of a domestic semiconductor industry, often using policy levers similar to other East Asian economies discussed above (Minnich, 2023, p. 194 Figure 6.1). Notably, from the 1980s onwards, Chinese industrial policy has sought to establish joint ventures or partner with foreign firms in an attempt to transfer the foreign technology that leaders understood was necessary to close the gap with the technology frontier.

Yet, unlike other East Asian economies above, China has not yet succeeded in reaching the technological frontier in any part of the semiconductor value chain. Based on data compiled by researchers at Georgetown University's Center for Strategic and Emerging Technology, China accounts for less than one percent of the market for software tools used to design chips, two percent of the market for core intellectual property, four percent of the market for silicon wafers and other materials used to make chips, one percent of the market for the tools used to fabricate chips, five percent of the market for chip designs, and seven percent of market share in fabrication (none of the latter capacity involves frontier technology). This has led one observer to conclude that "China is staggeringly dependent on foreign technology, almost all of which is controlled by China's geopolitical rivals" (Miller, 2022, p. 249).

China's modest progress in semiconductors, after decades of industrial policy, merits further discussion. Why has China's efforts yielded different results, particularly when China has been remarkably successful at absorbing foreign technology in other sectors. One compelling explanation comes from Minnich (2023), who argues that until the 2010s, due to low domestic final demand as opposed to production for exports, China lacked the bargaining power to effectively use *quid pro quo* policies—such as trading market access for foreign technology—to facilitate technology transfer. The author provides multiple examples where China repeatedly failed to secure the desired foreign partners for its national champions. Moreover, consistent with their low bargaining power in semiconductors, Minnich shows that China did not formalize technology-sharing requirements in their semiconductor industrial policy in the post-WTO period. This is in stark contrast to other sectors, such as automobiles (Bai, Barwick, Cao, and Li, 2020) or wind turbines (Minnich, 2023), where China has made extensive use these policy tools.

Importantly, this argument is consistent with a potentially more successful industrial policy pursued by China recently. Researchers suggest that China's latest push, which many date to the launch of the National Integrated Circuit Investment Fund in 2014 (see Section 4.3), has appeared more promising (VerWey, 2019; Miller, 2022; Minnich, 2023). This is partly due to the fact that as China has developed, the domestic market has become more lucrative,

6. This summary draws from Minnich (2023). A full description of Chinese industrial policy is beyond the scope of this paper; we refer the reader to Minnich (2023) for an excellent discussion.

increasing China's bargaining power *vis-à-vis* foreign firms. As a consequence, China has appeared more assertive in its efforts to acquire foreign technology (Minnich, 2023).[7]

Lessons from China's efforts are also consistent with lessons from earlier East Asian experiences. In particular, China's efforts only highlight the importance of technology transfer for successful industrial policy; outside of the U.S., our historical analysis has not uncovered a single instance where a domestic semiconductor industry developed without substantial foreign technology. This highlights that developing the industry, even with generous support, is difficult without foreign partners willing to share technology. The importance of foreign technology also provides insights into recent efforts by the Biden Administration to stymie China's technological advancement in semiconductors through export restrictions targeted at advanced software and advanced capital equipment.

4 Quantifying Recent Government Support in Semiconductors

Having reviewed the main government policies employed in the early phase of the industry, we now turn to more recent years. In this section, we discuss two approaches to measuring government support in the semiconductor value chain post-2010. The first, employed by the OECD in its 2019 report on semiconductors, applies the institution's longstanding work measuring *market distortions* to semiconductors (OECD, 2019). The second, which we implement and refine in this paper for the semiconductor value chain, builds on Juhász et al. (2023)'s approach to identifying *industrial policy*, henceforth JLOP. The two perspectives take different, disciplined approaches to quantifying policy and use distinct concepts of government support and implementation. We first provide a short description of each, compare their relative strengths and weaknesses, and then describe their findings. We summarize the main features of each approach in Table 2.

4.1 OECD's (2019) methodology for quantifying government support

The OECD (2019) methodology aims to quantify *market distortions* relative to an unobserved *laissez-faire* counterfactual. More precisely, the OECD uses an expansive definition of government support, which includes "any financial or regulatory measures that can affect

7. Due to the backlash surrounding the use of *quid pro quo* measures from foreign governments and firms, these policies have become more informal over time, making them harder to detect (Minnich, 2023). In addition, outward FDI (much of which was tied to state venture capital funds) emerged as one of the more visible forms of technology acquisition, although this has come under increasing scrutiny since the mid-2010s.

Table 2: Comparison of OECD (2019) and this paper

	OECD (2019)	This paper
Aim	Quantify market distortions	Identify industrial policies
Years covered	Average for 2014-2018	2010-2022
Unit of analysis	Firms	Policies
Support reported for	21 firms	Countries
Input data	Companies' own publications and financial disclosures	JLOP (2023) for classifying industrial policies in *GTA*

costs, prices or the profitability of market actors in any portion of the value chain, wherever they operate" (p. 62).

Practically, the report quantifies two main forms of government support for 2014–2018: i) budgetary government support (e.g., tax incentives, financial grants etc.), and ii) government support provided through the financial system. The former, budgetary government support, is primarily taken from firms' financial reports and disclosures. For the latter, the authors estimate below-market debt and the below-market provision of equity. Quantifying these requires comparing financial support to a hypothetical, unobserved market-based benchmark, which we discuss in Section 4.4.

Government support is then estimated from the "bottom-up," with firms as the unit of analysis. The report estimates the two forms of government support for twenty-one firms in the global semiconductor industry. These firms account for two-thirds of global industry revenue and cover different parts of the value chain. Given data limitations, this process excludes non-publicly traded firms (e.g., HiSilicon). Multinational firms play an important role in the sector and are common in the sample. Since multinationals receive government support from multiple jurisdictions, this complicates country-level calculations.

4.2 Our method: quantifying government support from Juhász et al. (2023)

This study takes a different approach to quantifying government support, building off the industrial policy dataset of Juhász et al. (2023). We do this in two steps: i) we use data from JLOP to identify industrial policies affecting the semiconductor industry; ii) we then systematically quantify the government spending for each, using a combination of policy text and source documents. We discuss both steps below, but first describe the JLOP dataset.

JLOP uses machine learning to identify industrial policies from policy text and produce counts of new industrial policies at the sector-country-year level for 2010-2022. The authors define industrial policies as those policies that "seek to change the relative prices across sectors

or direct resources towards certain selectively targeted activities (e.g., exporting, R&D), to shift the long-run composition of economic activity" (p. 5). They take this definition to data using natural language processing (NLP) techniques to automatically classify industrial policies from the *Global Trade Alert (GTA)* database, an international inventory of commercial policy (Evenett, 2009).

By construction, the *GTA* database focuses on economic policies that meaningfully affect foreign commercial interests relative to domestic ones (capturing both positive and negative discrimination).[8] Hence, the policies identified by JLOP cover industrial policies that differentially impact domestic commercial interests, such as preferential subsidies and lending to domestic firms.

For this study, we identify and quantify semiconductor industrial policy in two steps. First, we select industrial policies from the JLOP dataset impacting the semiconductor sector. Although the source data provides Harmonized System (HS) industry codes, HS codes are insufficient to identify semiconductor industrial policies (Bown, 2020). Categories may be too broad; for instance, "semiconductor devices" (HS 8541) encompasses not only discrete semiconductors but also solar panels and light-emitting diodes (LED). Moreover, they may not account for complex, evolving supply chains, which span chemical inputs (photographic chemicals HS 370790) and machinery for manufacturing wafers (HS 848610). Additionally, the *GTA* is missing HS codes for about one-third of the policies.

Hence, we combine "dictionary-based" methods from text analysis with HS code to identify semiconductor industrial policies. Specifically, we identify policies from a dictionary of terms associated with the semiconductor industry, such as *"integrated circuit"*, *"chip"*, or *"semiconductor wafer fab"* in the textual descriptions of the policies. This process produces an initial set of semiconductor industrial policies. We then manually verify that each industrial policy is related to the semiconductor sector and collect contextual information using additional sources (e.g., primary source statements, economic reporting, and trade industry associations). For all countries, we employ research assistants with local knowledge and language skills.

Second, we quantify government spending associated with each policy using source text and validate the data using research assistants with area expertise. This approach identified fifty-eight industrial policies affecting semiconductor manufacturing in seventeen countries. Table 3 contains the summary statistics.

8. Specifically, the *GTA* Handbook defines its scope as covering "credible announcement of a meaningful and unilateral change in the relative treatment of foreign versus domestic commercial interests" (Evenett and Fritz, 2022, p. 1).

Table 3: Summary statistics

	Mean	Std. Dev.	Min.	25p	75p	Max.
All policies ($n = 58$)						
Policies per country	4.00	3.52	1	1	7	11
Policies per country per year	0.31	0.64	0	0	0	3
Billions USD per country	12.79	22.19	0	0.10	16.06	77.77
Billions USD per country per year	0.98	6.72	0	0	0	76.82
National policies ($n = 35$)						
Policies per country	2.24	2.25	0	1	4	7
Policies per country per year	0.17	0.48	0	0	0	3
Billions USD per country	12.44	22.15	0	0	15.96	77.77
Billions USD per country per year	0.96	6.72	0	0	0	76.82

Notes: Count of policies and aggregate expenditure for the seventeen countries identified as having any semiconductor industrial policies for the period 2010 – 2022. National policies exclude policies targeting specific firms. European countries are reported individually. The value of policies jointly implemented by a group of countries is equally split across the participating jurisdictions. This table excludes the aggregate expenditure of Chinese provinces.

4.3 Chinese semiconductor policy

Chinese industrial policy is difficult to quantify (Kalouptsidi, 2018a), and it is worth describing the details and limitations of accounting for Chinese semiconductor policy. In particular, our method does not allow us to ascribe monetary values to new Chinese industrial policy programs, with the exception of one policy: the National Integrated Circuit Investment Fund (typically referred to as the "Big Fund").[9] The core issue follows from the provincial nature of Chinese industrial policy governance. Specifically, in China, national policy *guidelines* are deployed by provinces, which design and fund provincial industrial policies. Thus, unlike other countries in the sample, in China there are typically no monetary values associated with national industrial policy announcements (which is what we capture through the *GTA*).

Figure 2 shows the importance of subnational policy in China and presents the number of provinces implementing each national policy among the semiconductor policies identified in our data. We see substantial take-up among China's thirty-one provinces across the various industrial policies. All three Five-Year Plans (FYPs) covered by our sample period include semiconductor policy, and around two-thirds or more of the provinces are implementing industrial policies to develop semiconductor capacity locally, with an increasing trend over time. The two dedicated integrated circuit development plans have more limited take-up,

9. China's "Big Fund" is a public venture capital fund set up in 2014 with a capitalization of USD 47 billion and a mandate to invest in semiconductor firms (dollar value obtained from the Financial Times based on market intelligence from *JW Insights*).

with about one-third of the provinces implementing industrial policies linked to national guidelines.

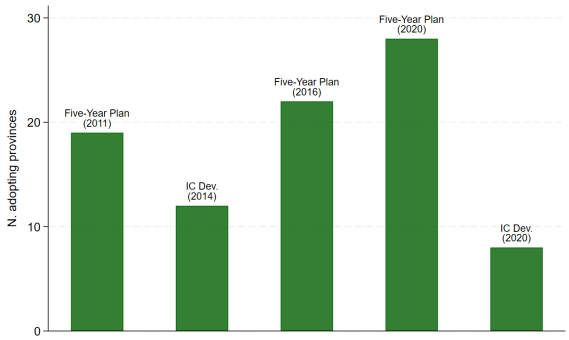


Figure 2: Take-up of national semiconductor industrial policies across Chinese provinces

Notes: The figure plots the number of Chinese provinces implementing each major national semiconductor industrial policy. Each province is coded as having taken up a particular national industrial policy if we could identify a specific provincial implementation plan. Source: authors' own research.

The provincial implementation of Chinese industrial policy means that we are almost certainly underestimating the value of semiconductor support if we do not include provincial spending. Figure 2 shows provincial spending is an issue; the question is how to quantify it.

We supplement our data with estimates of provincial spending on semiconductor policies, using information from the Financial Times based on market intelligence from *JW Insights*, a Chinese ICT market research firm.[10] A number of caveats are in order. We do not know how these estimates were constructed, nor what they include. Specifically, we do not know whether they include equity investments through the "Big Fund," which we do include in our initial, non-provincial estimates. Moreover, we do not know whether the estimates include spending beyond what was disbursed by implementing national-level industrial policies. Hence, additional information is required to make comparisons with other jurisdictions. For this reason, although we report numbers for China, we caution readers that they are uncertain and are not constructed like the other countries in our sample.

10. In the article titled "How Huawei surprised the U.S. with a cutting-edge chip made in China" dated November 29 2023, the Financial Times reports two types of data from *JW Insights* that we use to construct our estimates. First, they report spending by major provinces for 2016-2022. Second, for 2021-2022, they provide a breakdown of spending by the following categories: foundry, materials, testing, memory/storage, equipment, IDM, display panel, and others. We exclude the latter two categories (other and display panel) when disaggregated spending data are available.

4.4 Contrasting our approach and OECD (2019)

It is important to highlight the differences between the OECD's approach and ours. The OECD takes the firm as the unit of analysis and builds its estimation from the bottom up. Its goal is to quantify distortions relative to the unobserved laissez-faire. This approach has the important advantage of yielding estimates for both budgetary support and below-market financing. However, this method comes with some assumptions and some practical challenges.

First, the OECD method assumes that the policymaker changes incentives for *individual firms* at the margin. Put differently, this method is not well-suited to quantifying market distortions from "big push," or "moonshot" type government interventions that seek to catalyze large-scale change in an industry. To illustrate the issue, consider the case of the government investing in a portfolio of firms in the semiconductor industry, similar to the mandate of China's public semiconductor venture capital fund launched in 2014 (the "Big Fund"). The OECD methodology evaluates ex-post, firm-by-firm, whether the required return on capital was met, and classifies those that did not as receiving below-market equity investment. Yet precisely because the methodology evaluates outcomes firm-by-firm, it misses the fact that the portfolio of investments needs to be evaluated jointly, with the potential losses of many firms covered by the large gains of a few.

Second, relatedly, by quantifying distortions relative to a laissez-faire benchmark, researchers make an implicit assumption that a purely market-based counterfactual exists. This is certainly a plausible assumption in some cases, but in the presence of market failures, certain socially desirable economic activities do not exist without government intervention (see Juhász, Lane, and Rodrik (2024)). Practically, the OECD method assumes these cases away by positing a required return on equity that is the same for all semiconductor firms, irrespective of which market that activity takes place in.

Third, in terms of practical challenges, this approach is inherently suited to capturing support for large and successful firms, particularly publicly listed ones required by law to release financial statements, the key data source for the OECD study. Consequently, the approach will miss quantifying the effects of government support for firms that failed or did not become large enough to be enumerated. These issues may mean government support is underestimated.

Fourth, a distinct practical challenge relates to estimating below-market financing, which requires imposing some strong assumptions. This is particularly true for estimating the provision of below-market equity. Given that this source of support is important for the report's findings, we consider it in more detail. Broadly, the OECD compares the estimated required rate of return in the semiconductor industry (RRR) to a firm's net operating profits each year.[11] A firm is said to receive below-market equity any year in which the estimated

11. For a full description, see pp. 81-85 of the report.

cost of capital is lower than operating profits. The approach requires the authors to estimate RRR, which they do using a capital asset pricing model (CAPM).

Most importantly, we note that returns are evaluated on an annual basis instead of a net present value (NPV) basis. Although flow-based evaluations may be more straightforward to calculate, they may not be well-suited to projects where cash inflows and outflows vary through time. High-capital investment projects (e.g., a chip fab) may require large upfront expenditures with the anticipation of eventual, long-run payoffs. Evaluating annual profits at early stages of a project's life cycle may bias returns downwards for firms undertaking investment projects in anticipation of large future profits. In fact, the government's ability to provide this type of "patient" capital is often invoked as a justification for government equity investment.

Likewise, the approach taken in this paper has its own advantages and disadvantages. Here, the state and, in particular, its industrial policies are the unit of analysis. The advantage is that this allows us to more accurately capture *ex-ante* what the state is trying to do, regardless of whether that effort is successful. Another distinct advantage of working at the policy level is that it allows us to capture national-level policies that do not entail spending (e.g., import tariffs) and those that require expenditure.

An obvious disadvantage of JLOP is that it relies on count-based measures rather than values. We deal with this in two ways. First, *GTA* source data contains both major national policy announcements and, for some countries, detailed, firm-level data documenting disbursements of support. As these two types of policies are typically very different in magnitude, we present a version of our count-based measure of policies with and without disbursements of support to individual firms. Second, we estimate the magnitude of government support by assigning monetary values to semiconductor industrial policies.

That said, one should keep in mind a number of issues in interpreting values. First, in the case of instruments such as tax credits, total government support depends on the take-up of the policy. As such, our measure captures ex-ante (intended) support for the semiconductor industry. This may differ from ex-post support measured at the firm level after take-up. If the policy is uncapped, ex-post support may be larger than the estimated ex-ante support. If, however, take-up is lower than what the government intends, ex-ante support may be larger than ex-post support. Related to take-up, a spending package announced in a particular year may be disbursed over many years. The U.S. CHIPS and Science Act is a good example. While the bill was passed in 2022, support is being disbursed at the time of writing (2024).

Second, for instruments such as public loans and equity, we do not attempt to isolate the implicit subsidy from the full value of the loan or equity injection. This is not feasible in our context, given we do not observe disaggregated information, nor is it necessarily the object of interest. The reason is the different aims of the two measures. The OECD seeks to quantify market distortions resulting from government support, while we seek to characterize

the pattern of industrial policy spending. In that sense, a dollar of government support for semiconductors is a dollar not spent by the government elsewhere, holding total spending constant. Thus, for understanding patterns of industrial policy, taking the monetary value of the spending package is informative, though it measures a slightly different object than the OECD study.

There are also practical measurement challenges with our approach. One is that it is often difficult to pinpoint the budget allocated for supporting semiconductors when part of a larger spending package.[12] Moreover, spending may be adjusted or updated over time.[13] We deal with these cases through individual validation exercises and follow-ups.

Finally, we re-iterate the limitations surrounding Chinese industrial policies discussed in the previous section, specifically issues related to decentralized implementation. In these cases, our general approach is insufficient and requires supplemental data. Given the importance of China's semiconductor policy, we use external estimates of budgetary support for the sector to capture the possible magnitudes of the policy. These numbers should be taken as rough approximations.

Beyond the challenges of assigning monetary values to industrial policies, there are other limitations to consider with the JLOP data. First, by construction, JLOP captures policy flows. The *GTA* source data only tracks *new* policy announcements and does not capture the stock of pre-existing policies (e.g., any longstanding tax-incentives or budgetary support). Second, there is a risk of double-counting budgetary support in the *GTA*. For example, for China, the *GTA* includes disclosures of government support received by publicly listed firms. Thus, we may count support twice if we measure it once at the policy level and once at the firm level. Partly for this reason, we report two count-based measures below; one including all policies, and one including only national policy announcements (excluding firm policies). Finally, we note that the *GTA* is a living dataset. Our sample period ends in 2022, meaning that we likely undercount policies in more recent years, as there is substantial backfilling of data.

4.5 Patterns of Government Support in the Global Semiconductor Industry

We now describe the basic patterns of support for the global semiconductor industry based on our methodology. Where possible, we compare our findings to those in the OECD study. We emphasize two points, which are detailed above (Section 4.4). One is that the OECD reports data at the firm level, while we report data at the country level. In general, there is

12. For example, in 2021 the Japanese government announced a large COVID-19 recovery plan with USD 489 billion in spending. However, the portion of funding allocated to semiconductors was only USD 6 billion.

13. For example, the *GTA* entry for the "Big Fund" set up by China in 2014 states that by September 2015, the fund had exceeded USD 16.3 billion. However, the estimates of *JW Insights* reported in the Financial Times assessed the overall spending of the fund to be USD 47 billion.

no mapping from firms to countries in the OECD report. Second, while our data are annual, covering 2010 – 2022, the OECD data is not time-varying and instead captures averages for the 2014 – 2018 period.

4.5.1 Supporting jurisdictions

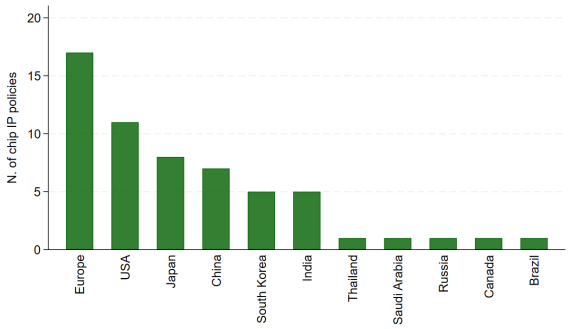
Figure 3 shows the jurisdictions that implemented semiconductor industrial policies during our sample period. Given the measurement challenges described above, we provide three different measures: Panel A provides the counts of all policies, Panel B plots the counts of national policies (excluding policies for individual firms), and Panel C plots the estimated monetary value of support.

The number of distinct jurisdictions reporting *any* semiconductor industrial policy is small. Specifically, there are seventeen jurisdictions that implemented at least one new industrial policy during our sample period (eleven if we group European countries). This may be a result of the capital and R&D intensity of the sector, which leads to both high levels of concentration and high barriers to entry.

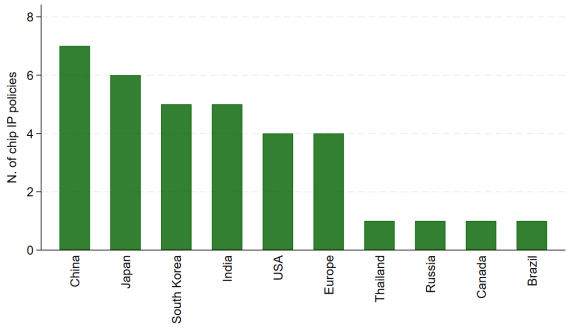
We find industrial policy concentrated predominantly among current major producers: China, Europe, Japan, South Korea, and the U.S. The notable exception to this is Taiwan, where we have identified no industrial policy. India stands out as a jurisdiction spending heavily to create a semiconductor industry without an established presence in the sector. This is notable; in terms of monetary value, India ranks fourth among all countries. Likewise, this result illustrates the advantages of an ex-ante measure of support, as the state appears to be trying to facilitate the establishment of the industry. Apart from India, we find no other country attempting to enter the industry with large-scale government involvement.

Which jurisdictions appear to provide the most support? Although not definitive, our analysis suggests the following tentative patterns. First, all measures point to six jurisdictions that provide the most intensive support: the five established producers (China, Europe, Japan, South Korea, and the U.S.) and India. Across all measures in Figure 3, industrial policy activity outside of the top-six jurisdictions is negligible.

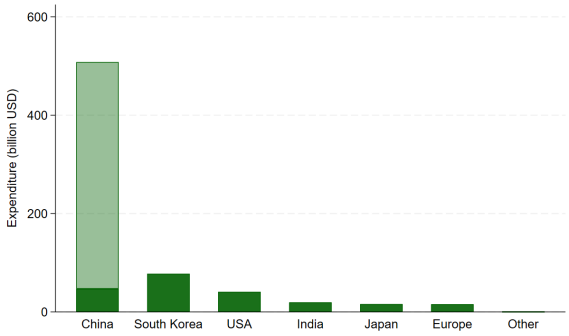
Second, our results are consistent with conventional wisdom that China has undertaken a substantial state-led push in the semiconductor industry over the past decade. China tops our list, marginally, if we compare the count of national policies across countries (Figure 3 Panel B). In absolute terms, for spending, our estimates suggest a substantial difference between China and other countries (Panel C). However, this does not take into account differences in the size of the economies. Although these results should be treated with caution, we believe they are broadly plausible. We explained in Section 4.3 that all three of China's latest Five-Year Plans have prioritized developing the semiconductor supply chain. Moreover, these national guidelines have been widely implemented across Chinese provinces (Figure 2). Alongside Five-Year Plans, there have also been two major integrated circuit policies, which



(a) Counts (all)



(b) Counts (national policies only)



(c) Monetary values (billions nominal USD)

Figure 3: Industrial policy by country

Notes: Panel A includes the count of all semiconductor policies. Panel B excludes policies deploying funds to individual firms and only enumerates national-level policies. Panel C assigns monetary values to all policies. The monetary value for China should be interpreted as follows: the dark green shading is spending on the "Big Fund". Light green shading refers to provincial spending on semiconductors for 2016-2022 as estimated by *JW Insights* and reported in the Financial Times. We group Europe together, which includes policies implemented by the following countries: Czechia, France, Germany, Italy, Netherlands, Malta, and the UK.

have been taken up by many provinces. Taken together, these sources suggest substantial support for semiconductors in China over the past decade. Yet, is China an outlier?

Although these estimates support the popular notion that China's government is substantially funding its domestic semiconductor industry, it is harder to assess whether this level of support is extraordinary. Indeed, an absolute, dollar-for-dollar comparison suggests China's support is overwhelming relative to other countries. Yet, comparing monetary support relative to the size of the economy, South Korea dominates China's spending.[14]

Thus, a judicious interpretation of this evidence suggests that while China provides substantial support to its industry, so do other major producers, with the notable exception of Taiwan (currently, according to our estimates). That is, most major producers actively support their domestic industries with large-scale national policies, especially when viewed as a share of total economic activity. Moreover, China's state-led effort in semiconductors is not historically unique. Our historical analysis above (Section 3) shows that every successful follower country we are aware of pursued aggressive industrial policy at the early stages of their domestic industry's development. In light of this, India's recent semiconductor push is also consistent with historical precedent.

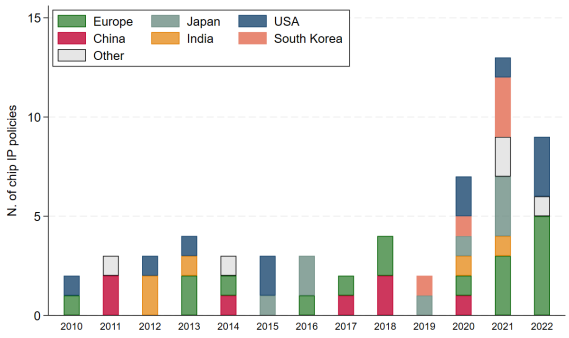
Next, in Figure 4 we examine the time trend of support over the sample period using the same three measures. Panel A of Figure 4 reports the count of all industrial policies by country and year. Panel B reports the same but excludes firm-level policies, and Panel C reports monetary values. Note that for monetary values in Panel C, we do not have a temporal breakdown of Chinese provincial spending, and only the value for the "Big Fund" is reported. Panel C also highlights the lumpiness of major spending packages in our measure; we report the date the policy was announced, not when funds are disbursed. For major policy packages, the dispersal typically takes place over multiple years.

Figure 4 suggests semiconductor activism has increased post-2020. While all countries have offered support to semiconductor firms through the sample period, the post-2020 period is notable. Table 4 details that since the COVID-19 pandemic, multiple countries have introduced major national policy packages specifically aimed at their semiconductor industries.

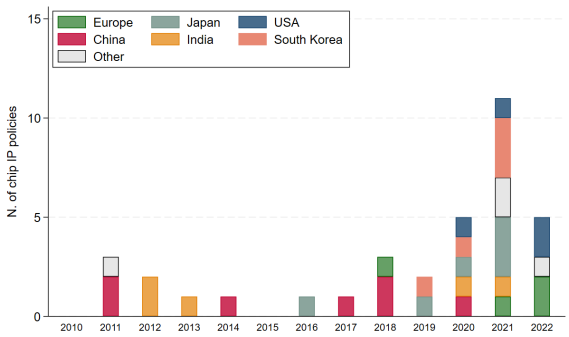
We next consider which distinct parts of the semiconductor supply chain are targeted by industrial policy. We distinguish three parts of the value chain: i) inputs; ii) design and manufacturing; and iii) assembling, packaging, and testing.[15] Table 5 shows that each part of the value chain has received government support. While design and manufacturing

14. Over 2010-2022, South Korea and China had an average GDP of USD 1.5 trillion, and USD 12 trillion respectively. Thus, for every USD 1,000 in GDP, South Korea and China have each spent USD 51, and USD 42 on semiconductor policies.

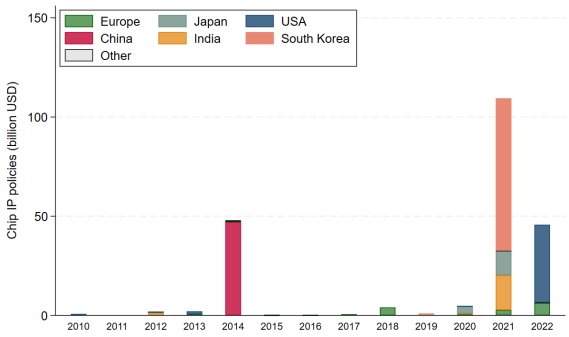
15. We use the classification from Bown (2020). We exclude ten policies that generically target the entire industry. These policies do not provide sufficiently granular information to assign monetary values to distinguish parts of the value chain. Among the ten excluded policies, six are Chinese policies and four are South Korean policies.



(a) Counts (all)



(b) Counts (national policies only)



(c) Monetary values excluding Chinese provincial spending (billions nominal USD)

Figure 4: Industrial policy by country and year

Notes: Panel A includes the count of all semiconductor policies. Panel B excludes policies deploying funds to individual firms and only enumerates national-level policies. Panel C assigns monetary values to all policies. For China, we lack the temporal breakdown of provincial spending. Only the monetary value of the "Big Fund" is included for China. We group Europe together, it includes policies implemented by the following countries: Czechia, France, Germany, Italy, Netherlands, Malta and the UK. "Other" comprises Brazil, Canada, Saudi Arabia, Thailand, and Russia.

Table 4: Timeline of major policies

Year	Policy	Country	Value
2011	12th Five-Year Plan	China	-
2012	National Policy on Electronics	India	-
2012	Modified Special Incentive Package Scheme	India	$1.6 billion
2014	Integrated Circuit Development Outline	China	-
2014	National Integrated Circuit Industry Investment Fund (Big Fund)	China	$47 billion
2016	13th Five-Year Plan	China	-
2018	Microelectronics research and innovation project	France, Germany, Italy, and UK	$2 billion
2019	Artificial Intelligence Strategy	South Korea	$1 billion
2020	14th Five-Year Plan	China	-
2020	Integrated Circuit Development Outline	China	-
2020	Program for Promoting Investment to Strengthen Supply Chains	Japan	$3.5 billion
2021	National Recovery and Resilience Plan	Germany	$1.7 billion
2021	Modified Programme for Semiconductors and Display Fab Ecosystem	India	$17.5 billion*
2021	Specified Semiconductor Funding Program	Japan	$5.4 billion
2021	Economic Measures for Overcoming COVID-19 and Opening Up a New Era	Japan	$6 billion
2021	Korean New Deal	South Korea	$72 billion
2021	Export control licensing	USA	-
2022	Electronique 2030	France	$5 billion
2022	CHIPS and Science Act	USA	$39 billion

Notes: The table is not a comprehensive list of all policies, only major policies are included.
* This value includes $7.5 billion approved under the Production Linked Incentives for Large Scale Electronics Manufacturing, the Production Linked Incentives for IT Hardware, the Scheme for Promotion of Manufacturing of Electronic Components and Semiconductors, and the Modified Electronics Manufacturing Clusters Scheme.

seem to have received the most support, inputs, as well as the most downstream, assembly, packaging and testing have also been targeted. This last step is typically targeted by emerging economies (Brazil in 2011, Russia and Thailand in 2021), though more recently also by advanced economies (Japan in 2019, and the US and Canada in 2022).[16]

Table 5: Industrial policy by part of the value chain targeted

	All	National
i) Inputs	**18**	**9**
Semiconductor manufacturing equipment	13	8
Specialized chemicals and materials	8	6
Electronic Design Automation software	7	5
ii) Design and Manufacture	**35**	**22**
Integrated Device Manufacturers	23	15
Semiconductor designers	15	14
Foundries	21	16
iii) Assemble, Package, and Test	**10**	**9**
Too generic	10	10

Notes: The same policy can target multiple parts and subparts of the value chain. "Too generic" refers to policies that contain no information on which part of the value chain is targeted (six Chinese and four South Korean). The first column reports all policies, the second column excludes firm-specific policies.

In comparison, the OECD findings can be summarized as follows. First, the OECD also find that almost all firms in their sample receive support, mostly in the form of budgetary support. Second, in terms of the absolute magnitudes of support received (which are most comparable to our estimates), they find large levels of support for some of the largest firms in their sample, such as Samsung and Intel.[17]

Third, a number of Chinese firms also receive large amounts of support, although these results are driven by below-market equity provision, which relies on the assumptions discussed in the previous sections (see Section 4.4).

Fourth, the report is able to identify the source of some budgetary support (though not consistently). Based on partial information, China, Singapore, Taiwan, and the U.S. were

16. Appendix Figure A.1 plots the time trend of policies by the part of the value chain targeted.
17. In the case of Samsung, the jurisdiction of where the budgetary support came from could not be identified. For Intel, budgetary support came from a variety of jurisdictions, the major ones being the U.S. and China.

identified as major supporters. This completes the picture painted from our estimates; the OECD finds that Taiwan, the only major producer missing from our estimates, also provides major support for semiconductor firms. Moreover, another discrepancy between our estimates and the OECD's is the large-scale support we identify for Japan and South Korea. This is likely due to the fact that Japan and Korea announced their major support packages after the OECD sample period (post-2018).

Taking all the findings together, we argue that government support in this industry seems to be a fairly general feature. Government support can be identified at all major parts of the value chain. Support can be identified amongst established producers at the technology frontier (notably, Korea and Taiwan), in countries hoping to develop their industry (China and the U.S.), and in countries attempting to enter the industry (India).

4.5.2 Policy instruments

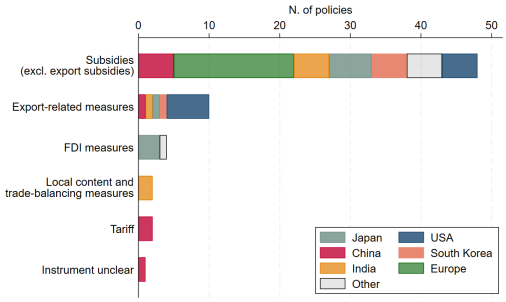
We now turn to the question of how government support for the industry is provided. Unlike in the previous subsections, it is not possible to assign monetary spending at this level of disaggregation.

Subsidies are the dominant form of support. In Figure 5, we plot the policy instrument used to provide support.[18] We classify policy instruments using two taxonomies. Panels A and B break down policies by the UN MAST (Multi-Agency Support Team) Chapter classification for non-tariff measures. Panels C and D do so using *GTA*'s more disaggregated classification scheme.

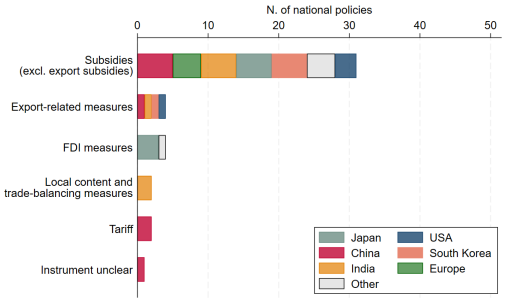
Subsidies are by far the most common type of support offered by countries (see panels A and B), which aligns with OECD findings. Consistent with the main forms of government support covered by the OECD report, we find subsidies predominantly taking the form of financial grants, state aid, tax measures, loans, loan guarantees, and equity injections to be prevalent.

Alongside these, we can also see some other forms of support that may be harder to identify in firm-level data, including FDI incentives (Japan), local content incentives (India), as well as trade policy instruments such as import tariffs (China) and export licensing requirements (U.S.).

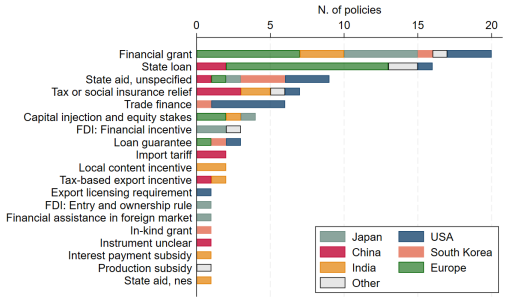
18. We note that a single policy is often composed of many types of measures. For example, in 2021 India introduced the Programme for Development of Semiconductors and Display Manufacturing Ecosystem. This policy covers two UN MAST categories: subsidies, and local content and trade-balancing measures.



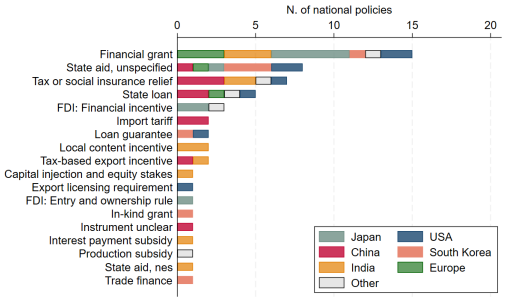
(a) MAST classification (all)



(b) MAST classification (national)



(c) *GTA* classification (all)



(d) *GTA* classification (national)

Figure 5: Type of support by country

Notes: Panels A and B show policies broken down by type according to the UN MAST Group. Panels C and D show policies broken down by type according to the *GTA* internal classification. Panels A and C include all semiconductor policies, while Panels B and D exclude policies deploying funds to individual firms and only enumerate national-level policies. We group Europe together, which includes policies implemented by the following countries: Czechia, France, Germany, Italy, Netherlands, Malta, and the UK. "Other" comprises Brazil, Canada, Saudi Arabia, Thailand, and Russia.

5 Policy Objectives

Why do governments target semiconductors with industrial policy? Having established that government support for the industry is ubiquitous among semiconductor-producing countries, we now turn to understanding more about the intentions of the policymakers.

An emphasis on growth, competitiveness, and, increasingly, resilience. For each policy, we have manually labelled both the stated *goal* of the policymaker, and the stated *means* to achieving that goal.[19] Table 6 shows that economic growth and development goals are by far the most prevalent. Indeed, this goal shows up across most jurisdictions and attests to the fact that domestic participation in the industry is widely seen as economically desirable. Similarly, the goal of enhancing international competitiveness is also widespread. Strikingly, resilience to supply chain disruptions was already an objective of the Indian and Chinese governments in the early 2010s. National security, not shown in Appendix Figure A.2, appears as a goal in European countries, the United States, Japan, and Canada starting in 2018 (Appendix Figure A.2 contains the time trends for the top three goals).

By what means do policymakers try to achieve these goals? Table 7 shows that policymakers are predominantly targeting the manufacturing process by subsidizing capacity expansion and the construction of new fabs, as well as supporting research and development across different parts of the value chain. We see fewer policies that try to achieve their goal by affecting international trade and investment decisions.

In comparison, the OECD finds that the vast majority of the budgetary support they identify is provided through research grants or the tax treatment of R&D spending (p. 9). They find smaller support for what they call investment incentives (production improvement in our terminology). This may be related to the fact that incentivizing investment is a more lumpy form of support, though it's also likely that production improvement policies became more important since the pandemic, and hence are not covered in the OECD report's time frame (for evidence of this, see Appendix Figure A.3, which shows the uptick in policies supporting production improvement since 2020).

Government documents may reveal growth, competitiveness, and resilience as the stated goals of policymakers, but government intervention in the semiconductor sector, and the use of subsidies in particular, require further justification. We turn to this next, with our model-based approach.

19. While there are fifty-eight total policy measures, a single measure can mention multiple means and multiple goals.

Table 6: Policies mentioning specific goals

Goal (broader aim)	All	National
Economic growth and development	38	25
International competitiveness	23	14
Resilience	10	8
National security	6	5
Green economy	4	3
COVID-19 support	3	3
No goal available	1	1

Notes: The first column reports the number of policies mentioning a specific goal. The second column excludes firm-specific policies, thus reporting only national-level policies. There are a total of 58 semiconductor industrial policies, 35 of which are national-level policies, as opposed to supra- or sub-national policies. Each policy can mention multiple goals. Economic growth and development include goals tied to industrial or technological development, economic growth, satisfying domestic demand, and addressing sub-optimal investment levels. International competitiveness includes the goals of expanding exports, becoming a global hub, catching up with the international frontier, or improving international competitiveness. Resilience includes the goals of self-sufficiency, supply chain resilience, and resilience to climate change or health pandemics. We use green economy for policies aimed at promoting the next generation of energy efficiency, resource saving, or the decarbonization and renovation of semiconductor production facilities. We further add the goal of providing economic support after the COVID-19 pandemic. (Note that we do not have enough information to identify the goal for one Korean policy.)

6 Model-Based Analysis

Industrial policy in the semiconductor sector is often justified based on the presumption of learning-by-doing (see the discussion in Section 2.3), which is a source of dynamic comparative advantage, high industry concentration, and market power. High concentration may make the industry vulnerable to disruptions, especially since the few firms dominating the industry are concentrated in a handful of countries. This section considers the evidence around learning-by-doing and government intervention; we do so with the help of a model-based analysis.

Our model serves three goals. First, it helps illustrate the industry features potentially justifying government intervention. Second, it provides an alternative way of measuring a subset of subsidies, namely production subsidies, beyond the data-driven analysis of Section 4. This analysis demonstrated the difficulties around measuring subsidies in the industry. China, in particular, is a complicated case. Hence, we complement this analysis using a model-based

Table 7: Policies mentioning specific means

Mean (direct objective)	All	National
Production improvement	34	24
Research, Development, and Innovation	30	19
International Trade and Investment	10	5
Other	1	0
No means available	3	3

Notes: The first column reports the number of policies mentioning a direct objective, or mean. The second column excludes firm-specific policies, thus reporting only national-level policies. There is a total of 58 chips-related industrial policies, of which 35 are national-level policies. Each measure can mention multiple means. "Production improvement" includes capacity expansion, creation of new fabs, production subsidies, preferential market access, tax incentives, and capital injection. "Research, Development, and Innovation" includes product innovation, investment in R&D, development of high-tech machines, and training of the workforce. "International Trade and Investment" includes export tax rebates, export licensing requirements, loans to export, and restrictions or encouragements to inward FDI. The category "Other" includes a generic capital increase for a specific company. Of the measures with no available means, all are South Korean.

approach for measuring subsidies, similar to the approach taken by Kalouptsidi (2018a) and Barwick et al. (2023). Our approach focuses on estimating a firm's marginal cost function using techniques from empirical Industrial Organization (IO). We identify "subsidies" as a residual in this function—i.e., a term contributing to a marginal cost decline, which cannot be explained by any other factors measured in the analysis. The advantage of this approach is that it provides a means of measuring the policy without direct policy data. A disadvantage, however, is that our results can be sensitive to modeling and identification assumptions. Third, the model helps us assess the effects of subsidies on both the implementing country and the rest of the world via counterfactual simulations.

6.1 A Simple Industry Model

This section describes the model, the estimation procedure, identification assumptions, and empirical results, and discusses their implications for knowledge spillovers and industrial policy. The theory combines several areas of economic literature. First, we follow the industrial organization literature in using the theory to generate demand- and supply-side conditions to detect industrial subsidies (e.g., Kalouptsidi (2018b); Barwick et al. (2023); Miravete, Moral, and Thurk (2018)). Second, our demand system leverages the international

trade literature's approach to estimating tractable demand models that account for cross-country quality differences (e.g., Hummels and Klenow (2005); Khandelwal (2010); Hallak and Schott (2011); Kugler and Verhoogen (2011); Johnson (2012); Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2020); Fajgelbaum, Goldberg, Kennedy, Khandelwal, and Taglioni (2024)). Finally, we incorporate dynamics by leveraging the empirical learning-by-doing literature (e.g., Irwin and Klenow (1994); Benkard (2000); Besanko, Doraszelski, and Kryukov (2019)).

The model focuses on the relationship between foundries where chips are produced— the suppliers in the model (e.g., TSMC in Taiwan or SMIC in China)—and the fabless firms who purchase chips—the buyers in the model (e.g., NVIDIA)—from foundries. We seek to identify and analyze production subsidies to the foundries, i.e., the transfers or other forms of implicit financial support that lower the cost of chip production. Of course, this is only a subset of policies that may be employed by the semiconductor sector. For instance, such subsidies do not include R&D subsidies. However, production subsidies are first order and arguably one of the more politically sensitive policy instruments targeted towards the sector. We think of these subsidies as tools which lower foundries' marginal costs, leading to lower prices for buyers.

Each period buyers (e.g., NVIDIA) source their chips from countries around the world. We assume that each country is endowed with a single firm which produces chips (e.g., TSMC in Taiwan and SMIC in China).[20] All period t information is known to buyers and firms but firms form expectations given period t information. Dynamics matter only through the supply-side via learning by doing. We further assume that buyers are myopic.[21]

6.2 Demand

We assume that a representative buyer (customer) maximizes a nested CES objective function over technologies k and suppliers j in each quarter t:

$$Q_t = \left[\sum_{k \in K} (\zeta_{kt} \; Q_{kt})^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}$$

where

$$Q_{kt} = \left[\sum_{j \in J} (\xi_{jkt} \; q_{jkt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$$

20. In practice, this assumption is less restrictive as it may seem, due to the high concentration in this industry – for instance, in Taiwan, TSMC has been the dominant player; in China, SMIC is the largest foundry by a wide margin.
21. See Sweeting, Jia, Hui, and Yao (2022) and Deng, Jia, Leccese, and Sweeting (2024) for the case of strategic buyers.

"Technology" k here refers to wafer diameter and the line width of the IC chips. More advanced chip generations have smaller line widths (measured in nanometers), which translate into more transistors (i.e., logic gateways), more computing potential, and greater energy efficiency. q_{jkt} and q_{kt} are the quantity consumed of product j and the quantity index for technology k respectively; ξ_{jkt} is product-level quality; ζ_{kt} is technology-node-level quality; σ is the within-group elasticity of substitution; and η is the across-group elasticity of substitution. We expect (but do not impose) that $\sigma > \eta > 1$, meaning that chips within a technology are more substitutable than chips across technologies.

We model product quality (ξ_{jkt}) as a function of observable product characteristics, and supplier fixed effects. Specifically, we model quality using the following functional form:

$$\xi_{jkt} = \bar{z}_{jkt}^{\delta} \times \nu_j^D \times \Delta\xi_{jkt}$$

where $\bar{z}_{jkt}$ denotes a "quality index" constructed on the basis of observable characteristics X_{jkt}, and whose construction is described in Appendix B.1. The characteristics included in X_{jkt} are the number of masks, metal layers, and polysilicon layers.[22] The term ν_j^D denotes demand-side supplier fixed effects, and $\Delta\xi_{jkt}$ is the residual unobserved quality.

Utility maximization leads to a demand equation where demand for chips quantity q_{jkt} of supplier j's technology k in period t is a function of price p_{jkt}, the quality index $\bar{z}_{jkt}$, a set of supplier fixed effects, technology-time fixed effects interactions, and an idiosyncratic error term ε:

$$\begin{aligned}\log(q_{jkt}) = &-\sigma\log(p_{jkt}) + \delta(\sigma-1)\log(\bar{z}_{jkt}) + \underbrace{(\sigma-1)\log(\nu_j^D)}_{\text{Supplier FEs}} \\ &+ \underbrace{\log(R_{kt}P_{kt}^{\sigma-1})}_{\text{Technology/Time FEs}} + \underbrace{(\sigma-1)\log(\Delta\xi_{jkt})}_{\text{Idiosyncratic shocks } (\varepsilon^l_{jkt})}\end{aligned} \tag{1}$$

The term R_{kt} is total expenditure on technology k in period t while the term $P_{kt} = \left[\sum_j \left(\frac{p_{jkt}}{\xi_{jkt}}\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}$ is the CES price index for technology k in period t.[23]

22. These are the key design and process choices that strongly influence both the quality and cost of a semiconductor device. A higher mask count generally reflects more complex circuitry and finer patterning, enabling greater transistor density and performance. Each additional mask, however, requires extra lithography, etching, and deposition steps, significantly raising manufacturing cost and cycle time. Similarly, increasing the number of metal layers expands the chip's interconnect capacity, reducing signal congestion and enabling faster, more power-efficient designs. This adds to fabrication complexity, yield risk, and cost due to additional deposition and chemical-mechanical polishing stages. Polysilicon layers, typically used to form transistor gates, play a critical role in defining device performance and electrical characteristics; while the count is usually low (one or two), additional poly patterning can improve circuit functionality and adds modestly to cost.

23. Though the technology-time fixed effects have a specific interpretation in the theoretical model, their interpretation in the empirical implementation of the demand model is broader as they also

Similarly, we model technology quality (ζ_{kt}) as a function of technology and time fixed effects:

$$\zeta_{kt} = \phi_k^D \times \rho_t^D \times \Delta\zeta_{kt}$$

where ϕ_k^D and ρ_t^D are demand-side technology, and time fixed effects, respectively. The term $\Delta\zeta_{kt}$ is the residual unobserved quality. This specification allows for end-customer demand to vary across chip technologies. Moreover, allowing demand to vary across time for all suppliers, in all countries and technologies, accounts for aggregate demand shocks that impact the entire fabrication industry—for instance, as customers increase their demand for digital devices.

The demand for technology k is given by:

$$\log(Q_{kt}) = -\eta \log(P_{kt}) + \underbrace{\eta log(P_t) + \log(Q_t) + \eta log(\rho_t^D)}_{\text{Time FEs}} + \underbrace{\eta \log(\phi_k^D)}_{\text{Technology FEs}} + \underbrace{\eta \log(\Delta\zeta_{kt})}_{\text{Shocks } (\varepsilon_{kt}^u)} \quad (2)$$

where $P_t = \left[\sum_k \zeta_{kt}^{\eta-1} \; P_{kt}^{1-\eta}\right]^{\frac{1}{1-\eta}}$ is the CES price index in period t.

6.3 Supply

As noted earlier, we assume that each country has only one supplier j (e.g., TSMC in the case of Taiwan, SMIC in the case of China). This assumption is driven by data constraints; in the data available to us, we can identify only the country in which a foundry is located, but not the foundry itself.

Cost Function

We model marginal costs as a function of the quality index $\bar{z}_{jkt}$, manufacturing wages in each supplier country, and a set of supplier, technology, and time fixed effects. The time fixed effects on the supply side capture among other things the prices of inputs that are internationally traded (e.g., silicone) and hence do not vary across countries. In addition, we follow Irwin and Klenow (1994) and assume period t marginal cost evolves with cumulative experience H_{jkt}:

$$c_{jkt} = \nu_j^S \times \phi_k^S \times \rho_t^S \times \bar{z}_{jkt}^{\beta} \times w_{jt}^{\kappa} \times H_{jkt}^{\gamma} \times e^{u_{jkt}} \quad (3)$$

$$H_{jkt} = 1 + Y_{jkt} + \underbrace{\alpha \times (Y_{jt} - Y_{jkt})}_{\text{Intra-Firm}} + \underbrace{\mu \times (Y_{Wkt} - Y_{jkt})}_{\text{Across Firms}} \quad (4)$$

account for time-specific preference shifts towards a particular technology, e.g., strong preference for a newly available technology.

where errors u_{jkt} are iid, w_{jt} denotes manufacturing wages in country j in period t, and ν_j^S, ϕ_k^S and ρ_t^S are supplier, technology, and time fixed effects, respectively. The set of fixed effects is similar to that on the demand side, but the rationale is different: Supplier fixed effects capture differences in marginal costs due to quality differences across foundries; technology fixed effects account for the fact that different technologies may require different inputs or processes; and quarter fixed effects capture changes in costs across time as prices of internationally traded inputs change or input manufacturers change their prices either because they become more efficient (i.e., input prices fall) or because they gain pricing power.

Y_{jkt} is cumulative past output of technology k sold to all buyers ($Y_{jkt} = \sum_{\tau=0}^{t-1} q_{jk\tau}$) in previous years. The terms Y_{jt} and Y_{Wkt} are the cumulative past output within the firm across the K technologies and the cumulative past output of technology k in the world, respectively. We refer to H_{jkt} as the "experience" of firm j producing technology k at period t and assume it is linear in past production. The parameter γ controls the importance of learning-by-doing as a function of a firm's experience. When $\gamma = 0$ there exists no learning while for $\gamma < 0$, the firm's cost decreases as it gains experience (i.e., as $H_{jkt} \uparrow$). When $\alpha > 0$, learning is internalized by the firm across the K technologies so there exists economies of scope in the learning process. When $\mu > 0$, learning spills over across firms (countries) due to, for example, technology transfers and improvements in inputs (e.g., photolithography equipment).

Firm Pricing

Each supplier maximizes expected discounted profits using quantity as the strategic variable. In Appendix B.4, we consider alternative assumptions regarding firm behavior (price competition; monopolistic competition; quantity competition with homogeneous products within technologies) and show that our results are robust to these alternative modeling assumptions.[24]

$$\max_{\{q_{jkt}\}} \left\{ E_0 \sum_{t=0}^{\infty} \left[\left(\frac{1}{1+d} \right)^t \sum_{k \in K} \left(p_{jkt} q_{jkt} - c_{jkt} q_{jkt} \right) \right] \right\}$$

The expectations operator E_0 indicates that a firm chooses quantity conditional on information available at time zero, d is a fixed exogenous risk-free rate, and c_{jkt} is the firm's marginal cost which varies across time because of learning-by-doing effects. The first-order condition shows that firms will set price as a markup over *dynamic* marginal cost:

24. An alternative modeling strategy would be to assume that foundries and buyers negotiate over individual orders. It is not clear how restrictive our modeling approach is relative to this alternative. Modeling quantity choice or pricing as a Nash-in-Nash Bargaining problem would replace our assumption with one where we assume negotiations are done independently (Collard-Wexler, Gowrisankaran, and Lee, 2019) which is also strong. Furthermore, we would not be able to credibly identify the parameters of such a bargaining model given that we do not observe the identity of buyers.

$$p_{jk0} = \frac{\varepsilon(s_{jk0})}{\varepsilon(s_{jk0})-1}\overbrace{\left(c_{jk0} + E_0\left[\sum_{t=1}^{\infty}\sum_{r\in K}\left(\frac{1}{1+d}\right)^t \times \frac{\gamma\,\Delta_{r,k}\,c_{jrt}\,q_{jrt}}{H_{jrt}}\right]\right)}^{\text{Dynamic Marginal Cost}} \tag{5}$$

$$\text{where } \Delta_{r,k} \equiv \begin{cases} 1-\mu & \text{if } r = k \\ \alpha & \text{if } r \neq k \end{cases}$$

or, in recursive form:

$$p_{jkt} - \frac{\varepsilon(s_{jkt})}{\varepsilon(s_{jkt})-1}\,c_{jkt} \;=\; \frac{1}{1+d} \;\times E_t\Bigg[\overbrace{\sum_{r\in K}\left(\frac{\gamma\,\Delta_{r,k}\,c_{jr,t+1}\,q_{jr,t+1}}{H_{jr,t+1}} + p_{jr,t+1} - \frac{\varepsilon(s_{jr,t+1})}{\varepsilon(s_{jr,t+1})-1}\,c_{jr,t+1}\right)}^{\text{Dynamic Incentive}}\Bigg] \tag{6}$$

where suppliers choose period t quantity conditional on information available at that time.

$\varepsilon(s_{jkt})$ denotes the elasticity of demand facing the supplier, which will generally be a function of its market share s_{jkt}. Note that this formulation is general and can accommodate several competition models popular in empirical work. In our model, which is based on the assumption of quantity competition, $\varepsilon(s_{jkt}) = \left[\frac{1}{\sigma}(1-s_{jkt}) + \frac{1}{\eta}s_{jkt}\right]^{-1}$. We consider various alternatives in Appendix B.4, and show that our results are robust to the mode of competition.

From (5), we observe that optimal price is a function of the standard markup and a dynamic component. Learning-by-doing makes the cost dynamic, so that the firm internalizes future cost reductions conditional on period 0 information. This is Boston Consulting's famous "Experience Curve." Importantly, the firm considers how technology k quantity choices impact the future marginal costs for all technologies (r) in its portfolio. We assume firms do not internalize the future competitive effects of their quantity decisions.[25]

If $\gamma = 0$, there exists no scope for learning-by-doing and we get the standard markup over static marginal cost. As learning-by-doing effects grow ($\gamma \downarrow -\infty$), however, the supplier optimally chooses to decrease its price from the static optimum, leading to higher period t output and lower future costs. Similar intuition holds for $\alpha > 0$ as the supplier internalizes cross-technology experience leading it to decrease prices across all of its K technologies. When $\alpha = 0$, there is no cross-technology spillover and production of technology r has no dynamic incentive for technology k. The extent to which suppliers benefit from cross-country learning comes via μ where we have assumed suppliers do not internalize their impact on

25. This is done for tractability as estimating the strategic quantity or pricing decisions in dynamic games is complex due to multiple equilibria. See Deng et al. (2024) for a two-firm numerical example.

the stock of global knowledge (captured via Y_{Wkt}). The learning channel amounts to a cross-border externality, where $\mu > 0$ amounts to free learning so the dynamic incentive to decrease price diminishes. Thus, the supplier will set price closer to the static profit-maximizing price when μ is large. We think of cross-border learning as including –among other things – improvements in supplier inputs such as photolithography equipment and technology transfers – in this sense, it may reflect general "learning" rather than "learning-by-doing".

The strength of learning-by-doing effects also depend upon the supplier's expectations about future demand and cost shocks via $E_t\left[c_{jr,t+1}q_{jr,t+1}\right]$. If the supplier thinks costs next period will be higher (i.e., high u), it chooses to decrease price today to increase production and reduce future costs by moving down the learning curve. On the other hand, if the supplier thinks there will be a positive demand shock next period, it reduces price to increase margins next period. These dynamic price incentives fall as the firm gains experience (i.e., $H \uparrow$), however.

6.4 Implications for subsidies

Consider an industrial policy that decreases the firm's marginal cost, such as a per unit subsidy. Under static pricing, such a policy has standard static price effects. However, with learning-by-doing, a policy in period t which reduces a firm's marginal cost increases demand in period t via lower prices, and subsequently decreases the firm's marginal cost in period $t+1$. The net effect is an industrial policy multiplier—one which is amplified (or mitigated) by the firm's strategic pricing response to the policy.

Yet, learning-by-doing *per se* does not necessarily justify the use of subsidies. If firms fully internalize the benefits of learning, then their private production decisions are socially optimal rendering subsidies redundant in the absence of other market failures. However there may be other considerations that justify the use of subsidies, even when learning-by-doing is internal to the firm —especially in the case of semiconductors.

One such consideration is that learning-by-doing generates dynamic comparative advantage and dynamic economies of scale, which can lead to high industry concentration. This is particularly significant in the case where learning-by-doing spills over across technologies (i.e., economies of scope). Anecdotal evidence suggests that within-firm learning is not exclusive to each chip generation; a firm's experience producing earlier generation chips impacts newer generation chips. For instance, TSMC's accumulated experience in producing chips has made it easier for the company to manufacture the newest, most advanced chip generation of <3nm line width. In the language of our model, this corresponds to the case of $\alpha > 0$. These internal-to-the-firm, but external-to-the-technology, spillovers can amplify the effects of learning-by-doing on market concentration. This is illustrated by TSMC's dominance of the advanced logic chip market segment. While such concentration may be desirable from

an efficiency perspective, it can lead to excessive reliance on a few firms within the world's semiconductor industry. Any disruption in the activities of these firms may jeopardize not only semiconductor production, but also the broader high-tech industry.

To the extent that economies of scope lead to an undesirable concentration of production making it vulnerable to shocks, subsidies to the domestic industry may increase diversification and enhance supply chain reliability and resilience (see Elliott, Golub, and Leduc (2022) for an analysis of how reliability externalities in complex, relationship-specific supply chains such as semiconductors may justify permanent subsidies to the industry). This argument becomes even stronger when coupled with geopolitical or national security concerns, especially if the countries where chip production is concentrated are vulnerable to major disruptions or considered antagonistic.

Beyond resilience and national security argument, policymakers might also be motivated by a profit-shifting motive, given that high concentration and market power imply supra-normal profits. However, profit-based arguments are more rare. A more common justification used by policymakers is that semiconductor manufacturing may create incentives for people to invest in skills and education, thereby upgrading the labor force and leading to positive spillovers to human capital. While anecdotal evidence from countries like Taiwan and Korea supports this claim, it is unclear what the direction of causality was. As discussed in Section 2, countries like Taiwan also invested heavily ex-ante in skills and education to make chip production possible within their borders.

Alternatively, if part of the learning-by-doing is external to each firm, then we have the classic case for industrial policy to increase production. In our model, this happens when $\mu > 0$. In this case, production of foreign companies across the global semiconductor industry may produce spillovers that benefit domestic firms, and vice versa, domestic production benefits foreign producers. The case for subsidies is less clear here however, as the benefits of domestically funded subsidies may in part accrue to foreigners.

In summary, the strength of the argument in favor of subsidies depends on the nature of learning-by-doing, specifically whether there are strong spillovers across technologies and whether learning is entirely internal to the firm. The empirical estimates of the model can inform these questions.

Cross-Border Effects and Welfare

Our model can be used to highlight several potential cross-border effects.

To the extent that subsidies accelerate or magnify learning-by-doing, they have positive effects on buyers in other countries benefiting from lower prices. We note that the buyers in this case are other firms (IDMs or fabless firms), most of which are located in countries other than the producing countries. In fact, the growth of the fabless business model, and the recent

evolution of the industry more broadly, might not have been possible without initial subsidies that fostered learning-by-doing (see Section 3).

However, the positive effects of subsidy-accelerated learning can be mitigated if dynamic economies of scale lead to market power and higher markups. Similarly, they may be diminished if subsidies shift production towards less efficient producers.

Additionally, if the industry is concentrated in a few countries, the rest of the world may become dependent on economic and political developments in these countries, reducing the industry's resilience and potentially raising national security concerns.

Unilateral subsidies will generally have a negative effect on chip producers in other countries through market stealing—we term this the "market-stealing effect." However, if there are learning spillovers across borders, it is possible that other chip producers benefit from subsidies to their competitors. Yet, even where the market-stealing effect dominates cross-border spillovers (i.e., unilateral subsidies make chip producers in the rest of the world worse off), it is possible that such subsidies *increase* global welfare if the subsidy-induced learning-by-doing sufficiently lowers global prices.

Finally, it is interesting to consider what happens when multiple countries subsidize their semiconductor sectors, engaging in a "subsidy race." Depending on the strength of the learning-by-doing and the shape of the learning curve (modulated by γ in the marginal cost function 3), it is possible that a subsidy race increases aggregate world production. This is more likely when the marginal return of experience is decreasing, i.e., costs reduce the most with a little bit of experience. In such a case, it might be optimal to spread experience across countries so as to foster competition that will counter the tendency towards high industry concentration and high markups. Similarly, when cross-border experience spillovers exist, competition in subsidies may improve welfare and the the subsidy race may become a "race to the top."

In summary, subsidies to the semiconductor sector can potentially have positive cross-border welfare effects. However, the extent of these effects is an empirical question. This empirical case hinges on the strength and nature of learning-by-doing and the initial efficiency of competing producers. At the same time, the economic forces that drive these welfare gains can also result in high industry concentration. Such concentration is problematic from both an economic and a geopolitical perspective, particularly if the dominant firms are located in only one or a few countries.

6.5 Identifying (Unobserved) Subsidies

The model introduced above provides a way of identifying production subsidies. Specifically, we can interpret the supply-side fixed effect ν_j^S from the marginal cost function (3) as the *average* (across years) subsidy in each country j.

The fixed effect ν_j^S captures all factors, beyond those controlled for in the marginal cost specification, that reduce (or, alternatively, increase) the marginal cost of production in country j relative to the benchmark country. The credibility of interpreting this fixed effect as subsidy depends on the extent to which all other factors affecting marginal cost are accounted for in equation (3).

Two aspects of this subsidy identification strategy are worth noting:

First, in an ideal setting, we would want to identify individual foundries in the data. In this case, we would control for foundry-specific characteristics that may impact marginal cost. However, in practice, this may not be a significant issue, given the high degree of market concentration in this sector. In many countries, a single firm holds a dominant market share, notably in the case of TSMC in Taiwan and SMIC in China.

Second, in theory, we can compare our inference-based subsidies to the subsidies calculated from our data in Section 4. For many countries (excluding China), we provide fairly reliable estimates of subsidies announced in policy documents. Thus, the subsidy estimates from the model could – in principle – be cross-validated against this information. Although this may not be possible for China, the strength of non-China comparisons provides a plausible way to evaluate the model-based inference. By extrapolation, this could enhance the credibility of the model-based estimates for China.

6.6 Estimation and Identification

We recover estimates of the demand and supply parameters in two steps. First, we estimate the demand-side parameters $[\hat{\sigma}, \hat{\eta}, \hat{\delta}, \hat{\nu}_j^D, \hat{\phi}_k^D, \hat{\rho}_t^D]$ using 2SLS (the instruments are described below). Given that buyers maximize static utility, the demand-side parameters are not affected by the learning parameters on the supply-side.

Next, we estimate the supply-side parameters $\hat{\theta} = [\hat{\gamma}, \hat{\alpha}, \hat{\mu}, \hat{\beta}, \hat{\kappa}, \hat{\nu}_j^S, \hat{\phi}_k^S, \hat{\rho}_t^S]$ using a generalized method of moments (GMM) estimator using the data of the GSA panel and the variables of period t, i.e.,

$$\hat{\theta} = \text{argmin} \left\{ g(\theta)' Z^s W Z^{s'} g(\theta) \right\} , \tag{7}$$

where $g(\theta)$ is a stacked vector of the supply-side structural errors (u_{jkt}) and W is the weighting matrix, representing a consistent estimate of $E[Z^{s'} g g' Z^s]$. The GMM estimator exploits the fact that at the true value of the learning-by-doing parameters, $(\gamma^\star, \alpha^\star, \mu^\star)$, the supply-side instruments Z^s are orthogonal to the structural errors because any ex-post errors are not correlated with ex-ante information available to the firms.

Recovering the underlying model parameters requires addressing price endogeneity in Equation 1 and Equation 2 and unobserved expectation of future endogenous variables in Equation 5.

On the demand side, simultaneity bias arises from the potential correlation of prices with *unobserved* quality ($\Delta\xi_{jkt}$ in Equation 1 and $\Delta\zeta_{kt}$ in Equation 2). We remind the reader that we control for observable quality through product characteristics used to construct a "quality index" as well as a large set of fixed effects (supplier, technology, and time fixed effects). Our supply-side model suggests a natural instrument for addressing potential correlation between unobserved quality and price: a firm's "experience" or stock of knowledge. The stock of knowledge affects price in our framework only through its effect on cost. Accordingly, we use a firm's cumulative past output in technology k, ($Y_{jkt} = \sum_{\tau=0}^{t-1} q_{jk\tau}$), as well as its cumulative past output in all other technologies ($Y_{jt} - Y_{jkt}$), in the set of instruments for p_{jkt}. In addition, we include the real manufacturing wage in the supplier country in the demand-side instruments. In the estimation of the upper-level demand equation (2), we construct similar instruments for the price index P_{kt} by averaging the aforementioned supplier-specific instruments across suppliers within a particular technology k.

On the supply-side, the instruments include – in addition to the demand-side instruments discussed above – the cumulative past world output in technology k, ($Y_{Wkt} - Y_{jkt}$); the one-period lag of a supplier's log quantity sold and market share in technology k (capturing AR(1) demand forecasts); the inverse-distance[26] of a particular chip technology k to all other chip technologies k' sold by the same supplier in the previous period, weighted by their market shares in the previous period; the number of quarters since the technology was introduced by any supplier times a dummy that takes the value one if the supplier is a follower and zero if the supplier was the first to introduce the technology; and the number of competitors in the previous period (i.e., number of products in the previous period in same technology), weighted by inverse of the number of quarters they have been in the market. These instruments capture the fact that a firm's stock of knowledge depends on the number of technologies it offers and the time each technology has been on the market as well as the order of entry; late entrants benefit more from a larger stock of accumulated knowledge. Note that we do not model entry in this paper; however, we use entry – the order of entry in particular – to identify learning spillovers. Finally, we set the interest rate outside the model as $r = 4\%$, which implies a risk-free rate of return consistent with US Government Treasuries.

The estimation algorithm is described in detail in Appendix B.2.

26. Distance is computed as the absolute value of the difference in chip diameter of all pairwise chip technology combinations offered by the supplier. Hence, the instrument takes the form: $\sum_{k' \in j} s_{jk'(t-1)} \frac{1}{|Current_k - Previous_{k'}|}$.

6.7 Data

The data used to estimate the model are from a proprietary database collected by the Global Semiconductor Alliance (GSA). The GSA is a nonprofit industry organization consisting of fabless firms. Each quarter the organization surveys its members to collect information on wafer fabrication prices, quantities, and characteristics of their orders from both domestic and foreign foundries. Responses are anonymous and firms which participate are granted access to the results. The dataset consists of 17,089 individual quarterly responses to the "Wafer Fabrication & Back-End Pricing Survey" covering years 2004-2019.[27] According to GSA the sample is representative of the industry and accounts for roughly one-fifth of all fabless semiconductor wafers produced worldwide. Our analysis focuses on the first three chip categories of Table 1 (i.e., microprocessors, "system of chip," and commodity ICs). We exclude memory chips because these are manufactured by large technological conglomerates such as IBM and Samsung, which also supply a large set of products outside the industry that are subject to different economic forces.

We note that because this data set covers only the 2004-2019 period, the overlap with the *GTA* data (2010-2021) is limited, so that the comparison of the model-based subsidy estimates with the *GTA*-based figures is only suggestive.

The data include nominal price paid, the number of wafers purchased, and the foundry's country location. We also observe wafer characteristics, including the line width, wafer size, and number of layers. We can therefore examine how foundry wafer prices vary by foundry location, after controlling for physical characteristics. Recall, we cannot identify the specific foundry which fulfilled each order, though the dominance of Taiwan Semiconductor Manufacturing Corporation (TSMC) in the Taiwanese market, Semiconductor Manufacturing International Corporation (SMIC) in Chinese market, and Chartered Semiconductor in Singapore market suggests that transactions which involve wafers fabricated in the Taiwan, Chinese, and Singapore markets where fulfilled by these firms, respectively. U.S. foundries produced on average 7% of the wafers during the sample. In comparison, fabrication plants in Taiwan, Singapore, and China accounted for 56%, 8%, and 10% of the market, respectively. Thus, while there are a variety of foundries producing semiconductor wafers for fabless semiconductor firms, the bulk of the wafers are produced overseas in Taiwan, presumably by TSMC.

6.7.1 Descriptive Analysis

The data reveals several interesting patterns, which we discuss before turning to model estimates.

27. Although data up to 2021 exist, we exclude observations from the COVID era.

Yields. We conduct an initial investigation of whether countries' market shares may be due to differences in yields across foundries. We do this via descriptive analysis of wafer "yields," defined as the percent of fabricated wafers which meet the specification required by the buyer; higher yield means higher quality.

Contract manufacturing service providers typically assume the risk of product defects. If some products are defective, the supplier must compensate by producing extra units to meet order requirements for functional products. This process increases the supplier's costs but does not generally affect the perceived service quality unless yields are so low that the supplier cannot meet the order timely. In such cases, low yields can cause missed delivery dates or unfulfilled orders, impacting service quality regardless of contractual terms.

However, in the semiconductor industry, both buyers and producers share the risk of low yields through explicit contractual terms. For example, contracts may define a minimum acceptable yield; if yields are below this threshold, the supplier must replace the wafer at no extra charge. Such contractual arrangements mean that while buyers assume some risk above this minimum yield, extremely low yields disrupt supply chain for the buyer which we interpret as the wafer being of low "quality." Byrne, Kovak, and Michaels (2017) report that a contract between Integrated Device Technology, Inc. and Taiwan Semiconductor Manufacturing Corporation required a minimum yield 65% for the first 300 wafers.

The data contain yield information for 2005 through 2008 where survey respondents (buyers) report whether their order fell into one of four buckets: 0-25%, 26-50%, 51-75%, and 76-100%. Table 8 presents these reported yields by source country. These results line-up with Byrne et al. (2017) and indicate Chinese yields are higher than yields reported in other countries. Interestingly, Taiwan has among the lowest yields. These patterns may be due to the fact that the reported yields refer to multiple chip generations. China produces predominantly older-generation chips, whereas Taiwan is at the technological frontier, producing the newest chip generation that has not yet benefited from learning-by-doing.

Pricing Trends. Figure 6 plots wafer prices for two different technologies, 90 nm and "mainstream" (50-89 nm),[28] where the number refers to chip diameter. Smaller diameters correspond to more advanced technologies. The plots suggest that prices start high and fall over time even when there is little competition. This downward trend is suggestive of learning-by-doing specific to each technology.

An alternative explanation for the decline in prices is increased competition rather than falling costs. As more suppliers enter production, markups, and therefore prices, may fall. Our empirical model allows for this possibility. On the supply side, both the quantity- and

28. The term "mainstream" corresponds to industry jargon

Table 8: Semiconductor Yields (2005-2008)

COUNTRY	0%-25%	26%-50%	51%-75%	76%-100%
CHINA	0	0	0	100
EUROPE	0	0	0.72	99.27
JAPAN	0	1.8	4.5	93.69
MALAYSIA	0	0	0	100
OTHER	0	0	16	84
SINGAPORE	0	0	3.87	96.12
SOUTH KOREA	0	0	3.9	96.09
TAIWAN	0.03	1.14	10.23	88.58
UNITED STATES	0	7.88	14.1	78

Notes: Authors' calculations based on quarterly GSA survey responses from 2005 through 2008.

price-competition specifications imply that markups decline as additional suppliers enter the market. The learning parameters we estimate are therefore net of this competitive effect.

Another interpretation, based on the demand side, is that in the early years after a new technology is introduced, demand may be especially strong because close substitutes are not yet available, which can lead to high prices. We capture this mechanism through time–technology interactions in the demand specification.[29] The learning parameters we estimate are again net of this demand-side effect.

The figure also helps clarify the various dimensions of learning-by-doing and our identification strategy.

Consider the price trajectories in Panel (a) for the 90 nm technology. For clarity, we focus on three suppliers: two early innovators (Taiwan and Japan) and a follower (China). A first observation is that China enters the market at a significantly lower price than the two leaders. A second observation is that prices for Taiwan and Japan fall rapidly in the early years after entry. Prices for China also decline over time, although the slope is much flatter (the scale of the figure makes this less visible).

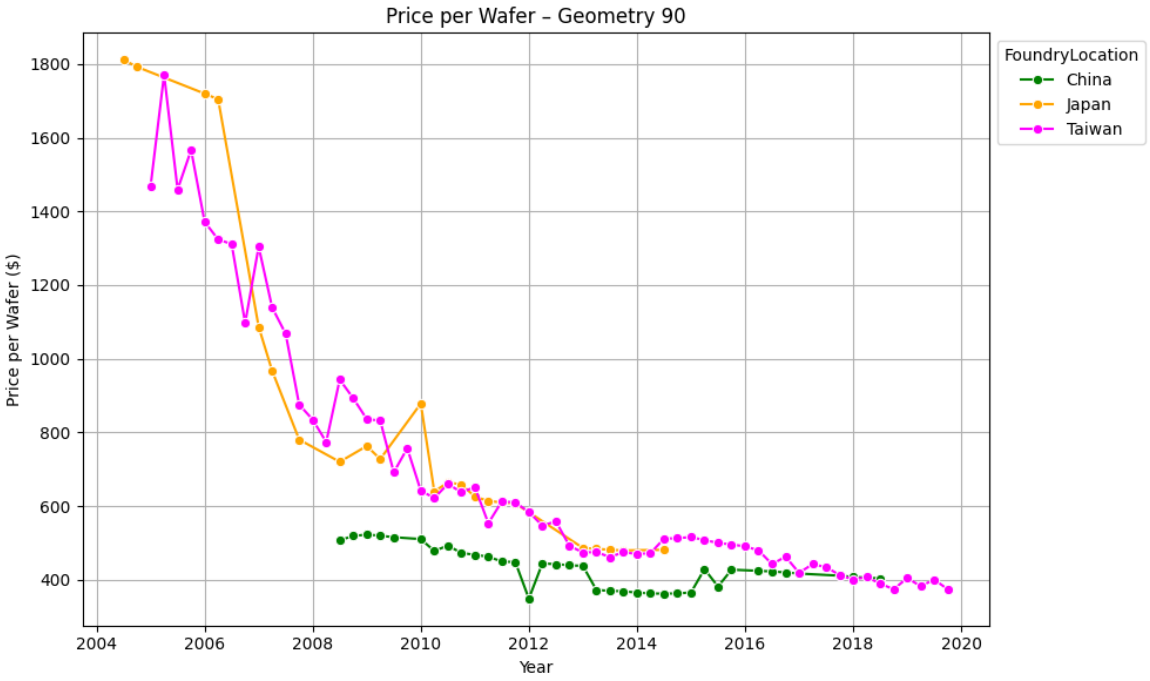
If we were to estimate the learning rate using the price paths of these three suppliers without differentiating between within-supplier versus across-supplier spillovers that followers enjoy, we would likely obtain a small learning effect. In contrast, if we allow late entrants (such as China) to benefit from learning generated by earlier entrants, we would attribute (at least part of) the lower entry price and the slower post-entry price decline to cross-supplier learning. We use these price differences at the time of entry to identify cross-border learning.

The lower prices offered by China could also reflect lower quality or a cost advantage. Given that the descriptive evidence shows higher yields for Chinese firms, the cost-advantage

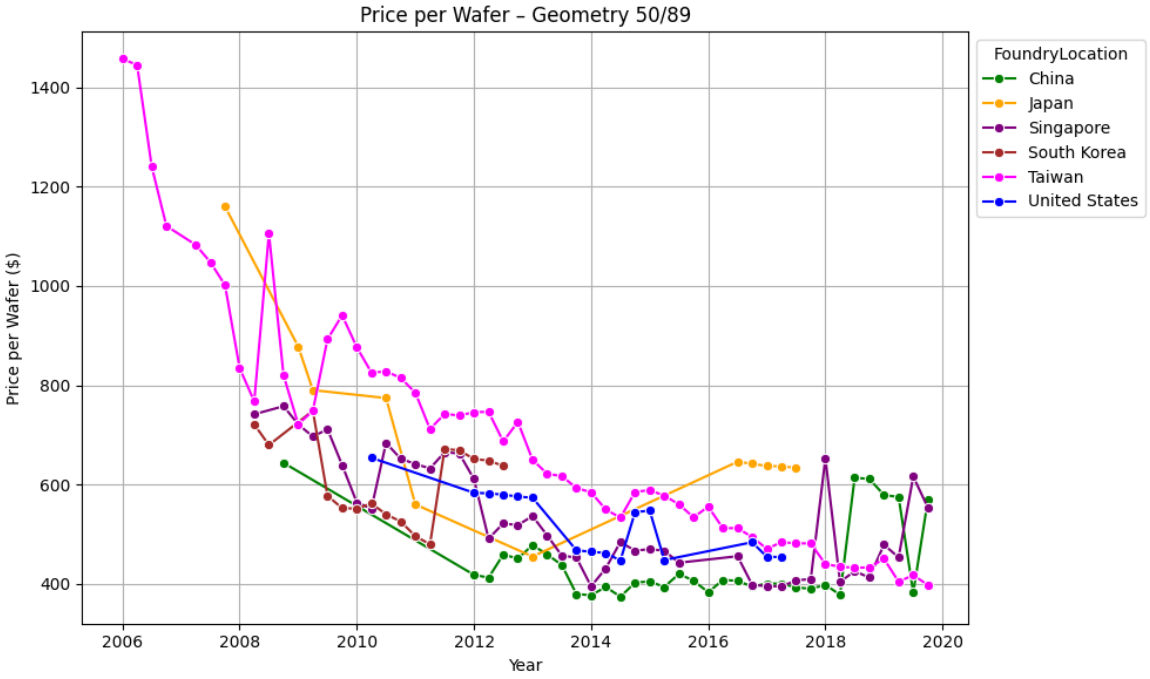
29. Note that the effect of lack of (or little) competition is already captured through the supply-side competition model that allows for markups to evolve with the number and market shares of competitors.

Figure 6: Pricing Trends

(a) Technology: 90nm



(b) Technology: “Mainstream” (50–89 nm)



explanation appears more plausible. Such an advantage could arise from lower wages, more efficient production practices, or government subsidies. The descriptive analysis alone cannot

distinguish among these possibilities. In our estimation, however, we control for cross-country wage differences, while time-invariant subsidies are absorbed by country fixed effects in the cost function.

A remaining possibility is that China subsidizes the early stage of production until learning-by-doing takes hold, allowing firms to become cost-competitive over time. While we cannot rule this out, the pattern we observe is not unique to China or to the 90 nm technology. Panel (b) shows similar price paths for the "mainstream" technology across several other supplier countries. In that panel as well, all follower countries enter at lower prices than the two leaders (Taiwan and Japan). This pattern forms the basis of our identification of cross-border learning.

6.8 Empirical Results

6.8.1 Results: Key model estimates

This subsection discusses some key estimates and their implications for the evaluation of industrial policy. The full set of estimation results are provided in Appendix B.3.

Demand-side Estimates The demand-side estimates are displayed in Table I in Appendix B.3. The two elasticities of substitution are large and precisely estimated, with magnitudes that conform to our expectations: The lower-nest elasticity of substitution $\sigma = -16.46$ ($s.e. = 1.76$) is larger in absolute value than the upper-next elasticity of substitution $\eta = -7.01$ ($s.e. = 2.71$). The parameter estimate for the "quality index" is positive and precisely estimated: 4.34 ($s.e. = 0.99$). As noted earlier, these estimates are not affected by the supply-side specification given that buyers are assumed to be static utility-maximizers.

Supply-side Estimates and Implications for Subsidies The supply-side estimates are reported in Table II in Appendix B.3. We estimate three specifications. In the first, learning-by-doing operates only at the supplier–technology level (1). In the second, we additionally allow for learning across technologies within a supplier (2). In the third, we also allow for cross-border learning spillovers across suppliers (3). Specification (3) is our baseline, and we base the discussion on this specification. We nonetheless report the first two specifications, as they help clarify the nature of learning and the sources of identification in our setting.

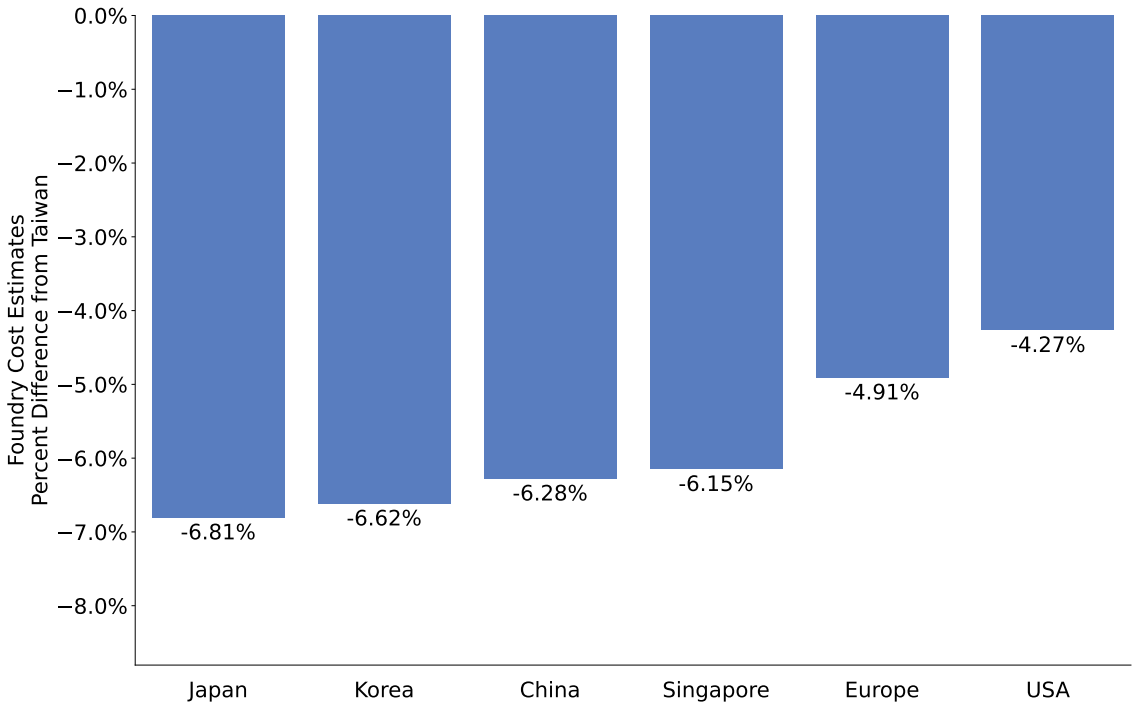
Wages and the "quality index" enter with positive coefficients, as expected, and are highly significant.

The production-side fixed effects are shown in Figure 7. For ease of presentation, we do not report standard errors in the figure, but all fixed effects are highly significant. Assuming that the marginal cost function adequately controls for all determinants of variable costs, these fixed effects can be interpreted as capturing average subsidies at the country level, averaged

over time and technologies. For example, the estimate of -6.62% for Korea implies that, over our sample period (2004–2019), production costs in Korea were on average 6.62% lower than in Taiwan, for reasons not captured in our model; this cost difference can then be attributed to subsidization. We emphasize that that our empirical specification controls for cost differences across technologies as well as sources of cost differences across countries. Labor cost is a significant cost component in semiconductor manufacturing, and our specification controls for wages. The second significant component of variable costs is the cost of inputs; these are internationally traded and absorbed by the time fixed effects. The empirical specification also controls for quality differences across suppliers that are associated with higher costs.

The bars indicate that on average, every country in our sample has subsidized foundries' manufacturing costs relative to Taiwan. This pattern aligns with results from the first part of the project based on *GTA* data, which indicate that Taiwan has not relied on production subsidies in recent years. The relatively small fixed effects for Europe and the United States, suggesting limited use of production subsidies, are consistent with evidence that these regions tend to subsidize research and development rather than production (OECD 2019). Consistent with the earlier GTA-based analysis, China does not appear to be an outlier, although we note that the two parts of the paper do not perfectly overlap in their sample periods. Finally, we stress that these fixed effects represent averages across time and technologies.

Figure 7: Estimated Country Fixed Effects
(Relative to Taiwan)



Learning-by-doing estimates. The estimated Spence learning rates under different assumptions about learning spillovers are reported in the last row of Table II in Appendix B.3. When we allow learning only at the supplier–technology level, and rule out both within-foundry spillovers across technologies and cross-border spillovers, the estimated learning rate is small, at 4.49%, with a 95% confidence interval of [3.61%–5.37%]. This estimate is well below industry lore and prior evidence, such as Irwin and Klenow (1994), who report learning rates in the range of 20–30%.[30]

The pricing patterns in Figure 6 help explain why this estimate is so low. As discussed in the descriptive analysis, prices for Taiwan and Japan—the two technology leaders—fall rapidly in the early years. In contrast, prices for the follower country, China, decline much more slowly over time. Because we identify learning-by-doing from price changes over time, the relatively flat price path for China in Figure 6 leads to a low estimated learning rate when spillovers are not allowed.

Once we allow for spillovers across technologies within a foundry, the estimated learning rate rises to 12.42%, with a 95% confidence interval of [6.32%–18.13%]. This estimate is consistent with anecdotal evidence pointing to economies of scope in learning. For example, semiconductor firms often foster a culture of experimentation and innovation (see Section 2.3), which is likely to benefit multiple product lines and facilitate the introduction of newer, more advanced technologies. In our framework, learning spillovers across technologies are identified using information on the number of technologies offered by each supplier, the timing of entry of each technology (with earlier entry generating larger spillovers not only for that technology but also for others produced by the same supplier), and the distance between technologies, where distance is measured by differences in chip diameter.

When we additionally allow for cross-border learning spillovers, the estimated learning rate increases further to 22.32%, with a 95% confidence interval of [16.35%–27.86%]. The intuition for this result is again illustrated in Figure 6. In the 90 nm technology, China enters at a lower price than the incumbents, Taiwan and Japan. In our framework, this lower entry price partly reflects learning from the earlier experience of other countries. As discussed above, this pattern is not specific to the 90 nm technology but also appears in other technologies and for all follower countries in our sample.

Focusing on the baseline specification in Column 3, which allows for all three types of learning spillovers, the estimates indicate that learning at the supplier–technology level is the strongest channel. The estimated values of γ, α, and μ are -0.36, 0.18, and 0.14, respectively. These estimates imply that a foundry learns about seven times more from an additional unit of its own cumulative output in a given technology than from an additional unit of cumulative output in that technology produced by foundries in other countries.

30. Irwin and Klenow (1994) study DRAM memory chips, which are much more homogeneous than the integrated circuits considered here.

At the same time, the cumulative past output of other foundries worldwide is large — often more than seven times the cumulative output of a single foundry in that technology. As a result, cross-border spillovers are substantial and, in aggregate, they likely outweigh the contribution of a foundry's own cumulative output. These spillovers are particularly important for follower countries, which – by definition –enter production with zero own cumulative output but can benefit from the accumulated experience of earlier entrants. In contrast, innovators (that is, countries that enter first, such as Taiwan) cannot draw on knowledge from earlier entrants in that technology, although they can still benefit from learning accumulated through the production of other technologies via economies of scope.

The sources of the large cross-border spillovers estimated in our framework are an important topic for future research. One possible explanation is learning along the global supply chain. In particular, buyers – the fabless firms in our data – as well as input suppliers may play a key role in spreading knowledge across borders. To leverage economies of scale in chip manufacturing, chip designers must coordinate closely with foundries and input suppliers (such as ASML) on the exact specifications of the chips and the manufacturing process. Once these specifications are established, other chip designers can adopt them to benefit from manufacturing economies of scale. This leads to the standardization of designs and processes and allows foundries to pool orders from multiple buyers. Once designs and processes are standardized, buyers and input suppliers can approach foundries in other countries with potential orders, sharing information gained from earlier interactions with different foundries. This process can create significant spillovers not only across buyers but also across foundries. For example, if TSMC works with fabless firms and ASML to improve a feature of the production process, that knowledge may diffuse through ASML's interactions with other foundries, generating spillovers throughout the industry.

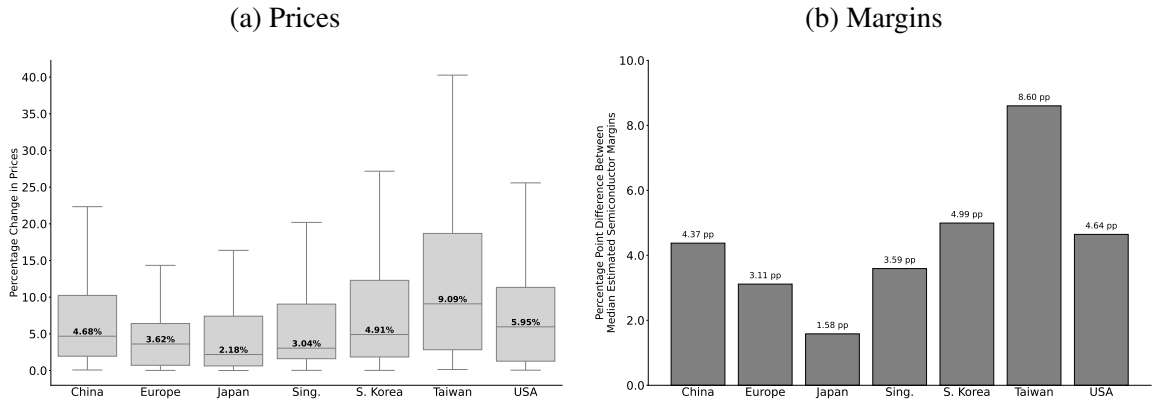
Other mechanisms that could generate cross-border spillovers are foreign direct investment (FDI), R&D collaborations, technology licensing and/or the recruitment of engineers or other experts from foundries in other countries. Our current framework does not allow us to distinguish between these hypotheses.

In sum, our results point to learning-by-doing with sizable cross-border spillovers. As discussed above, these spillovers may be linked to the fabless–foundry business model that has shaped the semiconductor industry over the past three decades and that facilitates cross-border knowledge sharing. Future work can help assess the importance of this channel relative to other sources of international spillovers.

Implications of learning-by-doing for prices and profit margins. We conclude the discussion of the results by examining how prices and profit margins are affected by learning-by-doing. Figure 8 shows by how much prices in each supplier country would have been higher

if suppliers did not internalize learning-by-doing (panel a), as well as the foregone profit margins (panel b)[31]

Figure 8: Impact of Internalizing Learning-by-Doing



If suppliers in each country did not internalize learning-by-doing, they would have charged higher prices, with the median percentage difference ranging between 2.18% for Japan to 9.09% for Taiwan. Regarding short-run profit margins, if firms did not internalize learning-by-doing, they would be higher, with the median (across technologies and time) difference ranging from 1.58 percentage points for Japan to 8.60 percentage points for Taiwan. The particularly large price and margin differences for Taiwan are consistent with anecdotal evidence suggesting that dynamic pricing induces TSMC to forego short-run profits. This pattern is also consistent with TSMC's position as the industry leader: TSMC/Taiwan benefits disproportionately from dynamic pricing in the early years following the introduction of a new technology. In contrast, follower firms in other countries can free-ride on cross-border learning spillovers, a benefit that Taiwan does not enjoy when it first enters a new technology.

6.8.2 Results: Implications for cross-border effects

The finding of large international learning spillovers implies large cross-border effects of subsidies. A subsidy provided by one country not only helps its domestic firms reduce costs but also leads to cost reductions in other countries. To the extent that some of the benefits of the subsidy are captured by foreign firms, one may question the wisdom of subsidization.

Note, however, that cross-border effects operate in both directions in this industry. As we noted earlier, learning may result from foreign technology transfer, which is – by defi-

31. To generate the counterfactual prices and margins shown in this figure, we solve the model under the hypothetical assumption that – despite the presence of learning-by-doing – firms maximize profits myopically based on static first-order-conditions. We maintain the assumption of quantity competition with product differentiation. The solution of the static profit maximization problem generates the (statically) optimal prices and quantities in each period. The quantities affect costs in subsequent periods due to learning, so they feed into the first-order-conditions of future periods, but firms ignore these intertemporal effects when they set prices myopically.

nition – a cross-border mechanism, or from firm-to-firm relationships between buyers and foundries located in different countries. Once learning has occurred, the same mechanisms that facilitated its creation will also help disseminate it across other foundries. The entire "fabless-foundry" model is based on these cross-border mechanisms. Therefore, one could argue that cross-border effects are not only present but are a defining feature of the current industry business model.

The realization of cross-border learning, however, is neither automatic nor inevitable. It depends on deliberate actions by market participants. Just as firms can share learning through business and research collaborations or technology licensing with other countries, they can also restrict access to frontier technology, cutting off certain countries. In such cases, international learning spillovers will be small, and the subsidy benefits will primarily accrue to domestic firms. However, positive cross-border effects may still occur if international buyers benefit from lower prices and innovations at the subsidized foundries.

It is also interesting to consider the effects when multiple countries subsidize their semiconductor industries. We are currently using model estimates to simulate such a "subsidy game" involving two countries. The results will provide insight into whether learning-by-doing makes subsidies global complements, suggesting that competition in subsidies could improve overall welfare. We aim to have these results available in the near future.

7 Conclusions

Our analysis highlights the important role of government support in the semiconductor industry's growth, especially during its nascent stages. Financial grants, tax incentives, loans, and equity injections have been fundamental in helping countries like Korea, Taiwan, China, and the United States enhance their semiconductor capabilities. Equally crucial is the role of cross-border technology transfer, which has historically enabled firms in follower countries to reach the technological frontier. Forces like foreign direct investment, research collaborations, and technology licensing not only promoted the industry's evolution but also fueled the rise of the "fabless" business model. These dynamics underscore the inherent nature of cross-border knowledge transfers in the semiconductor sector.

Quantifying subsidies, the primary form of government support, remains challenging. Our comparative analysis with the OECD study highlights the challenges of measurement, particularly in the case of China. In light of these challenges, our model-based approach offers a promising alternative for identifying and evaluating subsidies in this pivotal global sector. However, we emphasize that our model-based approach is not yet a definitive tool. Importantly, data constraints limit our ability to explore recent industry developments and to provide a more granular analysis of technology-specific effects.

Nevertheless, the subsidy findings implied by our model broadly align with the patterns of semiconductor policy documented using *GTA* data. These parallels are notable, given the distinct methods and the lack of overlap (in years) between the two datasets. Furthermore, our analyses indicate that subsidies are employed by all countries active in the semiconductor industry. While China is a significant user of subsidies in monetary terms, it is not an outlier when accounting for the size of its market.

Our model estimates suggest learning-by-doing at the firm-technology level that is consistent with industry lore as well as economies of scope within a firm and substantial cross-border learning spillovers. These findings are driven by pricing trends in our data rather than modeling assumptions. Although the exact source of international spillovers is beyond the scope of this study, we hypothesize that they result from cross-country technology transfers and the close relationships between fabless firms and foundries, the latter of which may facilitate the global dissemination of knowledge.

International spillovers imply potentially positive cross-border effects of subsidies. However, these effects depend on deliberate actions by market participants, meaning that firms in specific countries could be excluded. Future research will use counterfactual simulations to quantify these effects and explore the dynamics of a potential "subsidy race" among countries. Understanding whether subsidies are global complements or substitutes in the presence of learning-by-doing is essential, especially in understanding whether their use is compatible with the principles of multilateralism.

References

Amsden, A. H. (1989). *Asia's next giant: South Korea and late industrialization*. Oxford University Press, USA.

Bai, J., P. J. Barwick, S. Cao, and S. Li (2020). Quid pro quo, knowledge spillover, and industrial quality upgrading: Evidence from the chinese auto industry. Technical report, National Bureau of Economic Research.

Barwick, P. J., M. Kalouptsidi, and N. Zahur (2023). Industrial Policy Implementation: Empirical Evidence from China's Shipbuilding Industry. *The Review of Economic Studies, forthcoming*.

Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *American Economic Review 90*(4), 1034–1054.

Berger, A., H. Khan, A. Schrank, and E. R. Fuchs (2023). A new policy toolbox for semiconductor supply chains. *Issues in Science & Technology 39*(4).

Besanko, D., U. Doraszelski, and Y. Kryukov (2019). How efficient is dynamic competition? the case of price as investment. *American Economic Review 109*(9), 3339–3364.

Bown, C. P. (2020). How the United States marched the semiconductor industry into its trade war with China. *East Asian Economic Review 24*(4), 349–388.

Breznitz, D. (2007). *Innovation and the state: Political choice and strategies for growth in Israel, Taiwan, and Ireland*. Yale University Press.

Byrne, D. M., B. K. Kovak, and R. Michaels (2017). Quality-adjusted price measurement: A new approach with evidence from semiconductors. *Review of Economics and Statistics 99*(2), 330–342.

Chen, C.-F. and G. Sewell (1996). Strategies for technological development in south korea and taiwan: the case of semiconductors. *Research Policy 25*(5), 759–783.

Choi, H. (2008). Technology Importation, Corporate Strategies, and the Rise of the Japanese Semiconductor Industry in the 1950s. *Comparative Technology Transfer and Society 6*(2), 103–126.

Collard-Wexler, A., G. Gowrisankaran, and R. S. Lee (2019). Nash-in-nash bargaining: A microfoundation for applied work. *Journal of Political Economy 127*(1), 163–195.

Deng, S., D. Jia, M. Leccese, and A. Sweeting (2024, April). Bargaining and dynamic competition. Working Paper 32360, National Bureau of Economic Research.

Elliott, M., B. Golub, and M. V. Leduc (2022, August). Supply network formation and fragility. *American Economic Review 112*(8), 2701–47.

Evans, P. (1992). The state as problem and solution: Predation, embedded autonomy, and structural change. *The politics of economic adjustment 3*(3), 138–181.

Evenett, S. J. (2009). Global Trade Alert: Motivation and Launch. *World Trade Review 8*(4), 607–609.

Evenett, S. J. and J. Fritz (2022). The Global Trade Alert database handbook. Technical report.

Fajgelbaum, P., P. Goldberg, P. Kennedy, A. Khandelwal, and D. Taglioni (2024, June). The us-china trade war and global reallocations. *American Economic Review: Insights 6*(2), 295–312.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2020, 02). The return to protectionism*. *The Quarterly Journal of Economics 135*(1), 1–55.

Hallak, J. C. and P. K. Schott (2011). Estimating cross-country differences in product quality. *Quarterly Journal of Economics 126*, 417–474.

Hummels, D. and P. J. Klenow (2005). The variety and quality of a nation's exports. *American Economic Review 95*, 704–723.

Irwin, D. A. and P. J. Klenow (1994). Learning-by-doing spillovers in the semiconductor industry. *Journal of Political Economy 102*(6), 1200–1227.

Johnson, R. (2012). Trade and prices with heterogeneous firms. *Journal of International Economics 86*(1), 43–56.

Juhász, R., N. Lane, E. Oehlsen, and V. Perez (2023). The Who, What, When, and How of Industrial Policy: A Text-Based Approach. Technical report.

Juhász, R., N. Lane, and D. Rodrik (2024). The New Economics of Industrial Policy. *Annual Review of Economics*.

Kalouptsidi, M. (2018a). Detection and impact of industrial subsidies: The case of Chinese shipbuilding. *The Review of Economic Studies 85*(2), 1111–1158.

Kalouptsidi, M. (2018b). Detection and impact of industrial subsidies: The case of chinese shipbuilding. *The Review of Economic Studies 85*(2), 1111–1158.

Khandelwal, A. (2010). The long and short (of) quality ladders. *Review of Economic Studies 77*, 1450–1476.

Kugler, M. and E. Verhoogen (2011). Prices, plant size, and product quality. *Review of Economic Studies 79*, 307–339.

Liu, C.-Y. (1993). Government's role in developing a high-tech industry: the case of taiwan's semiconductor industry. *Technovation 13*(5), 299–309.

Lynn, L. H. (1998). Japan's technology-import policies in the 1950s and 1960s: did they increase industrial competitiveness? *International Journal of Technology Management 15*(6-7), 556–567.

Miller, C. (2022). *Chip war: the fight for the world's most critical technology*. Simon and Schuster.

Minnich, J. D. (2023). *Re-Innovation Nation: The Strategic and Political Logic of Technology Transfer Policy in Rising China*. Ph. D. thesis, Massachusetts Institute of Technology.

Miravete, E. J., M. J. Moral, and J. Thurk (2018). Fuel taxation, emissions policy, and competitive advantage in the diffusion of european diesel automobiles. *The RAND Journal of Economics 49*, 504–540.

OECD (2019). Measuring distortions in international markets: The semiconductor value chain. (234).

Sweeting, A., D. Jia, S. Hui, and X. Yao (2022, April). Dynamic price competition, learning-by-doing, and strategic buyers. *American Economic Review 112*(4), 1311–33.

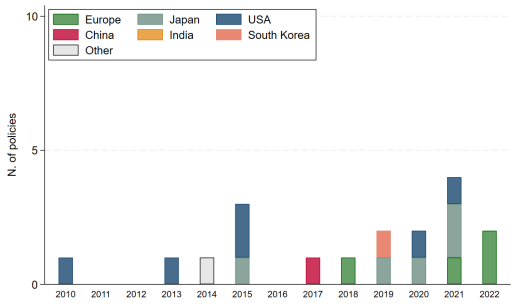
VerWey, J. (2019). Chinese semiconductor industrial policy: Past and present. *J. Int'l Com. & Econ.*, 1.

Online Appendix

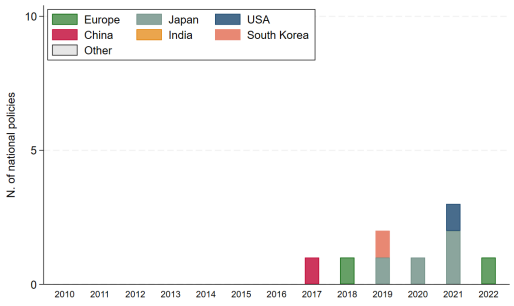
Industrial Policy in the Global Semiconductor Sector

Goldberg, Juhasz, Lane, Lo Forte, and Thurk

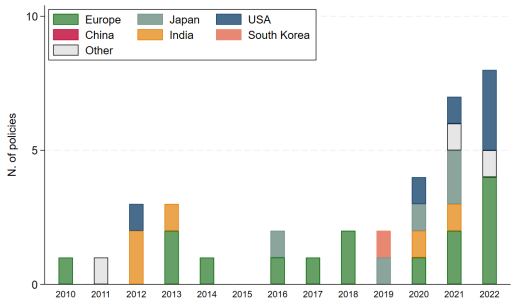
A Additional Figures



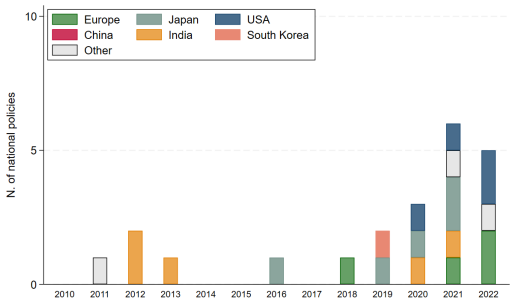
(a) Inputs (all)



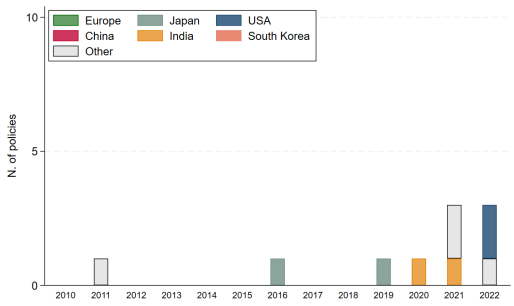
(b) Inputs (national policies only)



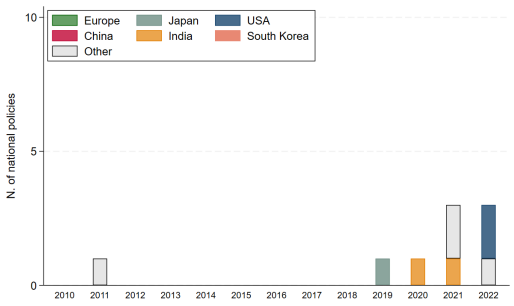
(c) DM (all)



(d) DM (national policies only)



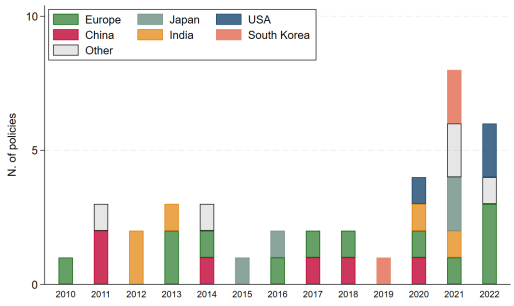
(e) APT (all)



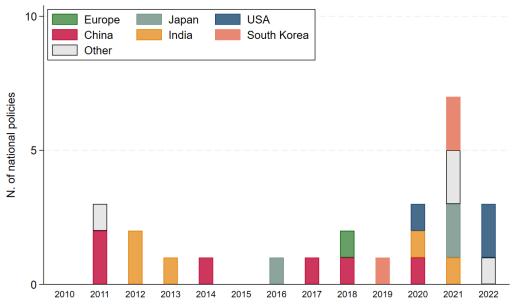
(f) APT (national policies only)

Figure A.1: Supply Chain targeted over time

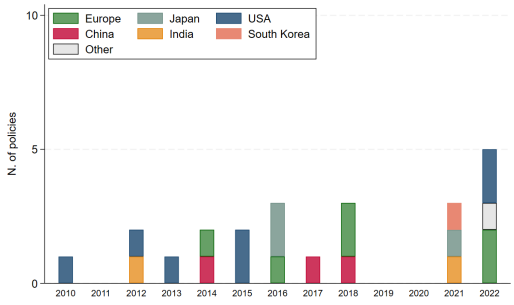
Notes: Panels A, C, and E include the count of all semiconductor policies. Panels B, D, and F exclude policies deploying funds to individual firms and only enumerate national-level policies. Panels A and B show policies targeting Inputs. Panels C and D show policies targeting Design and Manufacture (DM). Panels E and F show policies targeting Assemble, Package, and Test (APT). We group Europe together, it includes policies implemented by the following countries: Czechia, France, Germany, Italy, Netherlands, Malta, and the UK. "Other" comprises Brazil, Canada, Saudi Arabia, Thailand, and Russia.



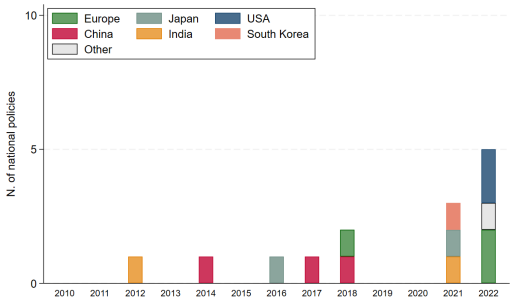
(a) Growth (all)



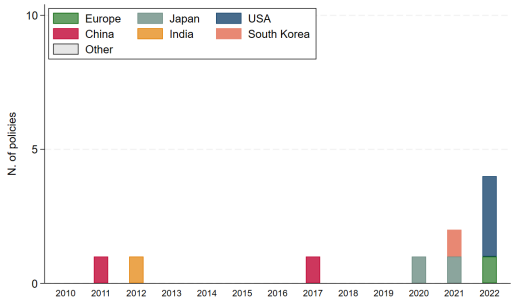
(b) Growth (national policies only)



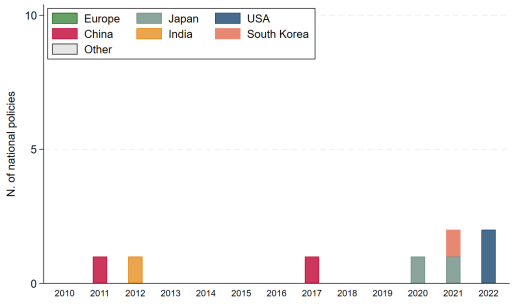
(c) Competitiveness (all)



(d) Competitiveness (national policies only)



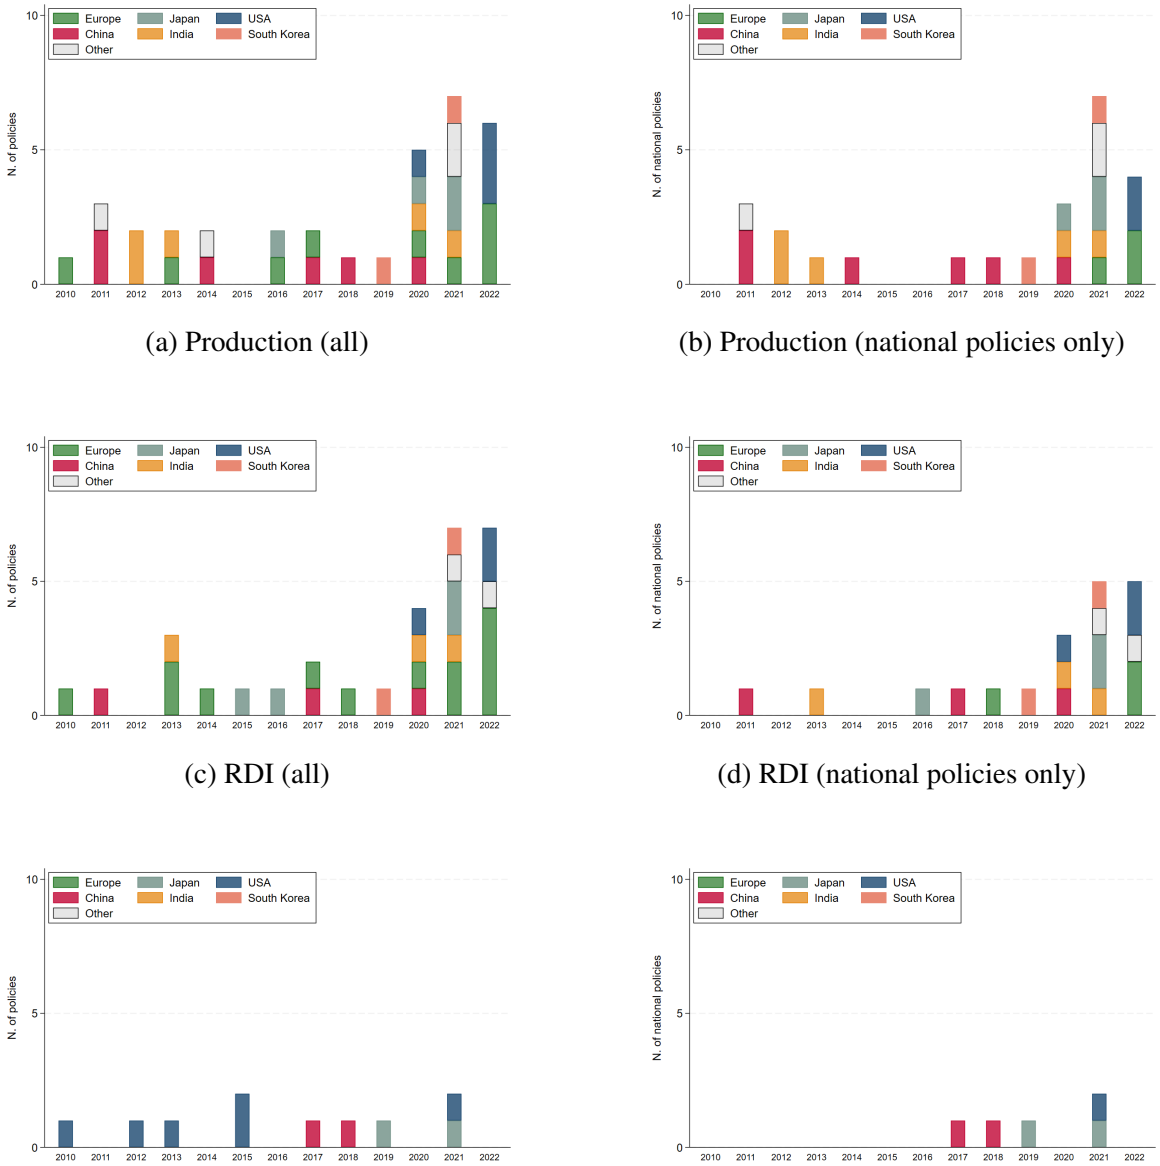
(e) Resilience (all)



(f) Resilience (national policies only)

Figure A.2: Goals over time

Notes: Panels A, C, and E include the count of all semiconductor policies. Panels B, D, and F exclude policies deploying funds to individual firms and only enumerate national-level policies. Panels A and B show policies with the goal of promoting economic growth and development. Panels C and D show policies with the goal of fostering international competitiveness. Panels E and F show policies aimed at improving resilience. We group Europe together, it includes policies implemented by the following countries: Czechia, France, Germany, Italy, Netherlands, Malta, and the UK. "Other" comprises Brazil, Canada, Saudi Arabia, Thailand, and Russia.



(a) Production (all)

(b) Production (national policies only)

(c) RDI (all)

(d) RDI (national policies only)

(e) ITI (all)

(f) ITI (national policies only)

Figure A.3: Means over time

Notes: Panels A, C, and E include the count of all semiconductor policies. Panels B, D, and F exclude policies deploying funds to individual firms and only enumerate national-level policies. Panels A and B show policies with the direct objective of improving production. Panels C and D show policies with the direct objective of supporting Research, Development, and Innovation (RDI). Panels E and F show policies with the direct objective of supporting International Trade and Investment (ITI). We group Europe together, it includes policies implemented by the following countries: Czechia, France, Germany, Italy, Netherlands, Malta, and the UK. "Other" comprises Brazil, Canada, Saudi Arabia, Thailand, and Russia.

B Additional Model Estimation Details

B.1 Construction of "Quality Index"

To construct the quality index $\bar{z}$ at the supplier-technology-quarter level, we make use of the *transactions-level* data. As noted in the data section, each observation in our data refers to a specific transaction, in which the buyer is anonymized; in the implementation of the empirical model, we aggregate the data to the product level jkt, but we use the transactions-level data when we construct our quality index. Let i denote the (unknown to us) buyer in a specific transaction "$ijkt$" involving the purchase of chips of technology k from supplier j by buyer i at time t, P_{ijkt} be the transactions-specific price, and Q_{ijkt} be the transactions-specific quantity of chips. The quality index for product jkt is constructed using three characteristics, number of masks, number of metal layers, and number of polysilicon layers, in the following five steps:

1. **Residual variation of metal layers and polysilicon layers**
   Because the number of masks ($\tilde{m}_{ijkt}$), number of metal layers ($metal_{ijkt}$), and number of polysilicon layers ($poly_{ijkt}$) are highly correlated, we first regress each of $metal_{ijkt}$ and $poly_{ijkt}$ on $\tilde{m}_{ijkt}$ to obtain the residual variation:

   $$metal_{ijkt},\ poly_{ijkt} \quad \Rightarrow \quad \hat{u}^L_{ijkt}, \hat{u}^P_{ijkt}$$

2. **"Hedonic" Price Regression**
   $$\log P_{ijkt} = \beta_P \log Q_{ijkt} + \gamma_m \tilde{m}_{ijkt} + \gamma_L \hat{u}^L_{ijkt} + \gamma_P \hat{u}^P_{ijkt} + \text{FE} + e_{ijkt}$$
   FE includes technology-time fixed effect interactions.
   We control for the quantity Q_{ijkt} of each transaction to account for the impact of potential quantity-based discounts on prices.

3. **Transaction-level quality**
   $$\bar{z}_{ijkt} = \hat{\gamma}_m \tilde{m}_{ijkt} + \hat{\gamma}_L \hat{u}^L_{ijkt} + \hat{\gamma}_P \hat{u}^P_{ijkt}$$

4. **Product-level Quality**
   $$\bar{z}_{jkt} = \sum_{i \in j} \log(Q_{ijkt}/q_{jkt}) \cdot \bar{z}_{ijkt}$$

5. We convert the log-quality estimate constructed in Step 4 into a strictly positive quality index, normalized so that the median quality in the first period equals 100.

B.2 Estimation Algorithm

We solve the model as follows conditional on a guess of the learning-by-doing parameters $\{\gamma, \alpha, \mu\}$.

1. Estimate $\{\hat{\sigma}, \hat{\delta}, \hat{\nu}_j^D\}$ via Equation 1 using 2SLS. The instruments used are described under "Estimation and Identification" in the main text. This estimation is unbiased provided unobserved quality $\Delta\xi_{jkt}$ is not serially correlated. As buyers maximize static utility, these estimates are fixed for any given $\{\gamma, \alpha, \mu\}$.

2. We recover $\{\hat{\eta}\}$ along with the technology and time fixed effects in Equation 2 as follows:

   (a) Recover $\hat{\xi}$ from estimates of Equation 1; i.e., $\hat{\xi}_{jkt} = \bar{z}_{jkt}^{\hat{\delta}} \times \widehat{\nu}_j^D \times \widehat{\Delta\xi_{jkt}}$

   (b) Construct price index $\widehat{P_{kt}} = \left[\sum_{j \in k} \left(\frac{p_{jkt}}{\hat{\xi}_{jkt}}\right)^{1-\hat{\sigma}}\right]^{\frac{1}{1-\hat{\sigma}}}$

   (c) Construct the technology-level quantity index $\widehat{Q}_{kt} = \left(\sum_{j \in k} \left(\hat{\xi}_{jkt} q_{jkt}\right)^{\frac{\hat{\sigma}-1}{\hat{\sigma}}}\right)^{\frac{\hat{\sigma}}{\hat{\sigma}-1}}$

   (d) Estimate $\{\hat{\eta}, \hat{\phi}_k^D, \hat{\rho}_t^D\}$ via Equation 2 using 2SLS. The instruments used are described under "Estimation and Identification" in the main text.

3. Given that suppliers have perfect foresight, we recover marginal costs $\{\hat{c}_{jkt}\}$ using Equation 6. We use the fact that learning-by-doing effects erode over time so that marginal costs for each supplier-technology pair (jk) are approximately fixed at the end of the sample period; i.e., $p_{jkT} = \frac{\varepsilon(s_{jkT})}{\varepsilon(s_{jkT})-1}\hat{c}_{jkT}$ at terminal value T. From this, backwards-induction generates the sequence of marginal costs $\{c_{jkt}\}_{t=1}^T$ using $\{q_{jkt}\}_{t=1}^T$, market shares in every period, $\hat{\sigma}$, $\hat{\eta}$, and Equation 6.

4. Estimate $\{\hat{\nu}_j^S, \hat{\phi}_k^S, \hat{\rho}_t^S, \hat{\beta}, \hat{\kappa}\}$ and recover supply-side structural errors $\{\hat{u}_{jgt}\}$ using a log-linear representation of Equation 3:

$$\log(\hat{c}_{jkt}) - \gamma \log(H_{jkt}) = \beta \log(\bar{z}_{jkt}) + \kappa \log(w_{jt}) + \log(\nu_j^S) + \log(\phi_k^S) + \log(\rho_t^S) + u_{jgt}$$

5. Construct orthogonality moments,$E[\hat{u}|Z^S] = 0$. The instruments for supply (Z^S) use the fact that firms make pricing decisions given their period t information set. At the optimal parameters, the structural errors are orthogonal to the information set. The supply-side instruments are described under "Estimation and Identification" in the main text.

B.3 Estimation Results

In this appendix we report select parameter estimates. Both the demand and supply-side specifications include a large set of fixed effects. We do not report those in the tables, but plot the relevant ones in the main text. The notes under each table indicate which fixed effects are included in each case. Table I presents the estimation results for the demand side. Table II presents GMM estimation results for three different specifications of learning-by-doing.[1]

Table I: Demand Estimation Results

PARAMETER	ESTIMATE	SE
LOWER NEST		
PRICE (σ)	-16.458	(1.762)
QUALITY INDEX	4.340	(0.999)
N. OBS	1,847	
UPPER NEST		
PRICE INDEX (η)	-7.008	(2.707)
N. OBS	443	

Notes: The table presents 2SLS demand estimates and their standard errors based on the GSA Semiconductor Wafer Pricing Survey (2004-2019). The specification for the lower nest includes technology-time (technology-quarter) fixed effects, as well as country fixed effects. The specification for the upper nest includes technology and time (quarter) fixed effects. Prices in the lower nest are instrumented using wages, the experience of the firm in the same technology, and the experience of the firm in any other technology. The price index in the upper nest is instrumented using the average of the instruments used in the lower nest at the technology-quarter level. Wages are quarterly average hourly manufacturing wages collected from each country's local statistical bureau and converted to US dollars using nominal exchange rates from Feenstra, Inklaar, and Timmer (2015). Prices and wages are in 1982 dollars using the consumer price index from FRED.

1. An implication of our model is that demand estimates are not affected by how we specify the learning process. This is because the buyers are assumed to be static utility maximizers.

Table II: GMM Estimation Results

	(1)		(2)		(3)	
PARAMETER	ESTIMATE	SE	ESTIMATE	SE	ESTIMATE	SE
WAGES	0.110	(0.026)	0.085	(0.025)	0.107	(0.026)
QUALITY INDEX	0.399	(0.050)	0.269	(0.049)	0.238	(0.050)
TECHNOLOGY LEARNING (γ)	-0.066	(0.007)	-0.191	(0.050)	-0.364	(0.054)
FIRM LEARNING (α)			0.174	(0.075)	0.179	(0.063)
WORLD LEARNING (μ)					0.138	(0.058)
N. OBS	1,847		1,847		1,847	
LEARNING RATE	4.5%		12.4%		22.3 %	
[95% C.I.]	[3.6%, 5.4%]		[6.3%, 18.1%]		[16.4%, 27.9%]	

Notes: The table presents the supply-side GMM estimates based on the GSA Semiconductor Wafer Pricing Survey (2004-2019). All specifications include time (quarter), technology, and country fixed effects. In specification (1), foundries learn only from past experience in the same technology. Specification (2) allows learning from experience in other technologies, and specification (3) – our baseline specification – additionally allows learning from world experience in the same technology. The instrument set in specification (1) includes lagged quantities sold and experience in the same technology. Specification (2) augments this set with experience in other technologies. Specification (3) further adds world experience in the same technology for technologies smaller than 180 nm (the reason for the focus on technologies smaller than 180 nm is that – as noted earlier – we exploit the order of entry, and in particular the lower prices of follower countries at entry; to identify the impact of international spillovers on costs; however, given that our sample starts in the 2000s, we do not observe entry for technologies that are older than 180 nm, as such technologies are available by all suppliers in our sample); the number of quarters since technology introduction interacted with a follower-foundry indicator; lagged market share; the number of competitors in the previous quarter weighted by their tenure; and the distance to technologies previously produced by the supplier weighted by market share. Wages are quarterly average hourly manufacturing wages collected from each country's local statistical bureau and converted to US dollars using nominal exchange rates from Feenstra et al. (2015). Wages are in 1982 dollars using the consumer price index from FRED.

B.4 Robustness

The formulation of the supply side can accommodate several models of competition. Specifically, depending on the assumptions about firm competition, the elasticity facing the firm $\varepsilon(s_{jkt})$ will be given by:

- **Quantity Competition**: $\varepsilon(s_{jkt}) = \left[\frac{1}{\sigma}(1 - s_{jkt}) + \frac{1}{\eta}s_{jkt}\right]^{-1}$
- **Price Competition**: $\varepsilon(s_{jkt}) = \sigma(1 - s_{jkt}) + \eta s_{jkt}$

- **Monopolistic Competition**: $\sigma = \eta; \varepsilon = \sigma; Markup = \frac{\sigma}{\sigma - 1}$

- **Homogeneous Cournot** (homogeneous products within technology k): $\varepsilon(s_{jkt}) = \frac{\eta}{s_{jkt}}$

Our results regarding learning rates, and in particular cross-border knowledge spillovers, are robust to these alternative assumptions. Below we report results for the first three cases in Table III.

B.4.1 Results for alternative competition models

In column 1, we report the results for our baseline specification: quantity competition with differentiated products. In column 2, we report results for price competition, and in column 3, for monopolistic competition. We did not estimate the specification corresponding to homogeneous product quantity competition, as the assumption of homogeneity is clearly violated in our data.

Table III: GMM Estimation Results for Other Market Structures

	(1)		(2)		(3)	
PARAMETER	ESTIMATE	SE	ESTIMATE	SE	ESTIMATE	SE
WAGES	0.107	(0.026)	0.116	(0.026)	0.109	(0.026)
QUALITY INDEX	0.238	(0.050)	0.239	(0.050)	0.242	(0.050)
TECHNOLOGY LEARNING (γ)	-0.364	(0.054)	-0.375	(0.055)	-0.360	(0.055)
FIRM LEARNING (α)	0.179	(0.063)	0.209	(0.070)	0.183	(0.066)
WORLD LEARNING (μ)	0.138	(0.058)	0.166	(0.067)	0.143	(0.060)
N. OBS	1,847		1,847		1,847	
LEARNING RATE	22.3%		22.9%		22.1%	
[95% C.I.]	[16.4%, 27.9%]		[16.9%, 28.4%]		[16.0%, 27.7%]	

Notes: The table presents the supply-side GMM estimates based on the GSA Semiconductor Wafer Pricing Survey (2004-2019). All specifications include time (quarter), technology, and country fixed effects. In specification (1), foundries learn only from past experience in the same technology. Specification (2) allows learning from experience in other technologies, and specification (3) additionally allows learning from world experience in the same technology. The instrument set in specification (1) includes lagged quantities sold and experience in the same technology. Specification (2) augments this set with experience in other technologies. Specification (3) further adds world experience in the same technology for technologies smaller than 180 nm; the number of quarters since technology introduction interacted with a follower-foundry indicator; lagged market share; the number of competitors in the previous quarter weighted by their tenure; and the distance to technologies previously produced by the supplier weighted by market share. Wages are quarterly average hourly manufacturing wages collected from each country's local statistical bureau and converted to US dollars using nominal exchange rates from Feenstra et al. (2015). Wages are in 1982 dollars using the consumer price index from FRED.

References

Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). The next generation of the penn world table. https://www.rug.nl/ggdc/productivity/pwt/. Version 10.01, retrieved from https://www.rug.nl/ggdc/productivity/pwt/.