

## THE NEW ECONOMICS OF INDUSTRIAL POLICY

Réka Juhász, Nathan Lane, and Dani Rodrik<sup>1</sup>  
August 2023

### ABSTRACT

We discuss the considerable literature that has developed in recent years providing rigorous evidence on how industrial policies work. This literature is a significant improvement over the earlier generation of empirical work, which was largely correlational and marred by interpretational problems. On the whole, the recent crop of papers offers a more positive take on industrial policy. We review the standard rationales and critiques of industrial policy and provide a broad overview of new empirical approaches to measurement. We discuss how the recent literature, paying close attention to measurement, causal inference, and economic structure, is offering a nuanced and contextual understanding of the effects of industrial policy. We re-evaluate the East Asian experience with industrial policy in light of recent results. Finally, we conclude by reviewing how industrial policy is being reshaped by a new understanding of governance, a richer set of policy instruments beyond subsidies, and the reality of de-industrialization.

---

<sup>1</sup> University of British Columbia, University of Oxford, and Harvard University, respectively. This is a draft of a paper prepared for the *Annual Review of Economics*. We thank Jessica De Simone for her expert editorial assistance.

## 1. Introduction

There are few economic policies that generate more kneejerk opposition from economists than industrial policy. This has not stopped governments from making abundant use of it, even when they seem ideologically hostile to it.<sup>2</sup> The salience of industrial policy has risen greatly in recent years, as governments have increasingly engaged in self-conscious industrial policies as they address a variety of problems – the green transition, resilience of supply chains, the challenge of good jobs, and geopolitical competition with China. Academic economists have often acted as by-standers (and often naysayers) as policies such as the CHIPS and IRA acts in the U.S. have been developed and implemented.

The good news is that there is much to be learned from the variety of industrial policies around the world. A considerable literature has developed in recent years providing rigorous evidence on how industrial policy really works and how it shapes economic activity. This literature is a significant improvement over the earlier generation of empirical work, which was largely correlational and marred by interpretational problems. The recent crop of papers offers in general a more positive take on industrial policy. More importantly, it provides a much more nuanced and contextual understanding of industrial policy. It enables economists to engage in the debates around industrial policy in a more productive manner, shedding light rather than heat.

We summarize the outline and main conclusions of this paper as follows. We begin in section 2 by providing our definition of industrial policy and presenting the standard arguments both in favor of and against the use of industrial policy. We argue that there is a generic and powerful economic case for industrial policy and that the usual critiques rely on practical rather than principled objections. In light of this, it is curious that the debate on industrial policy in economics has focused on the “whether” (“should governments carry out industrial policy?”) rather than on the “how” (“how should industrial policy be carried out?”).

In section 3, we turn to the actual practice of industrial policy. We discuss the difficulties of measuring industrial policy and then focus on recent systematic efforts to overcome these difficulties. We provide a summary characterization of current industrial policies. Several key conclusions emerge. First, it is no longer appropriate, if it ever was, to identify industrial policy with inward-looking, protectionist trade policies; contemporary industrial policies typically target outward-orientation and export promotion. Second, industrial policy has been ubiquitous, and its prevalence predates the recent rise in its use and prominence in public discussions. Third, it appears advanced economies are heavier users of industrial policies than developing countries; there is a steep income gradient in the reliance on industrial policies.

Section 4 focuses on the evidence of how industrial policy works. There is an inherent difficulty in ascertaining the causal effects of industrial policy since, *by design*, policy intervention is non-random and

---

<sup>2</sup> President Reagan, famous for labeling government as the problem rather than the solution, protected U.S. steel, auto, and motorcycle industries from import competition to encourage them to invest in technological upgrading. President Pinochet of Chile, as close to an avatar of free market liberalism as one can imagine, subsidized the forestry sector and promoted its exports. Margaret Thatcher actively courted Japanese auto makers and promoted their UK investments with financial incentives.

targets certain industries for a mix of economic, political, or administrative reasons that cannot be perfectly observed. We show that correlational work of the type that was the norm until recently is uninformative in a precise sense: such evidence cannot distinguish between the polar cases of rent-seeking and developmental governments – and anything in between – as these cases are observationally equivalent. We further show that standard identification techniques can be useful but will not resolve the debate on whether systematic use of industrial policy works or not.

The new generation of work we discuss has less ambitious aims and focuses on evaluating whether industrial policy elicited the desired behavioral response in specific settings. We focus on three types of cases: episodes of infant industry promotion (e.g., in textiles, shipbuilding and heavy industries); large-scale public R&D efforts (as in the “space race” between the U.S. and Soviet Union); and selective place-based policies targeting specific industries (as in the U.S. manufacturing drive during World War II and contemporary regional European subsidies). Interestingly, the most recent vintage of papers, paying serious attention to identification and observability difficulties, produces results that are much more favorable to industrial policy. These papers tend to find that industrial policy has typically shifted resources in the desired direction, often producing large long-term effects in the structure of economic activity.

Section 5 revisits the industrial policy experience in East Asia. The East Asian miracle constitutes not only one of the most important episodes of modern economic development, it also remains the focal point of debates surrounding the efficacy and desirability of industrial policy. Early literature on the East Asian experience was sharply divided. Many regional specialists, mostly sociologists and political scientists along with a few economists, ascribed at least part of the region’s success to the strong hand of the state in driving industrialization. Most mainstream economists, meanwhile, have held the view that industrial policies were at best ineffective and at worst harmful. A new generation of work on the region’s industrial transformation is beginning to paint a more nuanced picture. This work, often using disaggregated data, pays careful attention to the diversity of policies in question and the structure of linkages and production networks. It shows that certain types of industrial policy were powerful in driving structural change in countries such as Japan, South Korea, and China. But it also suggests that it is very difficult to derive broad generalizations for other countries and time periods from this experience without taking institutional differences in consideration.

Nevertheless, it would be a mistake to see the East Asian experience as idiosyncratic and exotic. As we emphasize, in section 5, there are some useful lessons to be drawn. The region’s experience calls for a broadly strategic and dynamic approach to the practice of industrial policy. The instruments used, the relative emphasis placed on domestic firms versus FDI, and the balance between import substitution and export promotion varied both across countries and over time, depending on local opportunities and constraints. Furthermore, while East Asian states have traditionally been characterized as “hard,” in contrast with the “soft states” prevailing elsewhere, East Asian models of industrial policy are a precursor of today’s successful practices and provide useful guidance on the design of future industrial policy.

We end the paper (in section 6) with a discussion of the new economic context, which is reshaping our traditional conceptions of industrial policy. We focus on three challenges in particular. First, we suggest that the actual practice of industrial policy often departs from the top-down model of regulation that economists typically adopt when they think of such policies. We use the concept of “embedded

autonomy,” borrowed from sociologist Peter Evans, to characterize an alternative model of regulation, based on iterative collaboration between government and firms. Second, we suggest that successful industrial policy often uses a broader range of policies that can be more effective than the classic instrument of subsidies (or of trade policy). These include, importantly, customized public services and inputs that are tailored to firms’ needs and target specific obstacles to productivity enhancing investments. Third, we argue that employment de-industrialization will necessarily reshape the practice of industrial policy, as manufacturing continues to employ fewer workers (relative to the rest of the economy). Policy makers will have to pay more attention to productivity in services, and come up with “industrial” policies that are more suited to services. We illustrate these arguments with specific examples.

## **2. The industrial policy debate: definitions, rationales, and critiques**

### **2.1. Definitions**

Industrial policy is much discussed but rarely explicitly defined. We define industrial policies as those government policies that explicitly target the transformation of the structure of economic activity in pursuit of some public goal. The goal is typically to stimulate innovation, productivity, and economic growth. But it could also be to promote climate transition, good jobs, lagging regions, exports, or import substitution. Since industrial policy targets structural change, a key characteristic is the exercise of choice and discretion by the public authorities: “we promote X but not Y,” though the latter part of this statement is typically left implicit.

Industrial policy has traditionally focused on promoting manufacturing industries such as steel, autos, shipbuilding, aircraft, or semiconductors – hence the name. But our definition is open-ended and includes support for services as well as particular types of R&D. Hence industrial policies overlap with what in other contexts might be called regional policies (Slattery & Zidar 2020), place-based policies (Neumark & Simpson 2015), or innovation policies (Mazzucato 2014). In developing countries, industrial policies are often called productive development policies (Fernández-Arias et al. 2016) or structural transformation policies – in part because of the negative connotation the term “industrial policy” has acquired, but also to reflect the reality that similar policies have to be deployed for a wider array of developmental challenges going beyond industrialization.

Industrial policies can take various forms but always create incentives for private-sector actors – firms, innovators, investors – to act in ways that are consistent with the intended direction of structural change. Subsidies (on specific types of exports, investment, R&D, etc.) are the most obvious types of industrial policy. But the gamut runs from import protection to exemptions from specific regulations to public provision of key inputs such as land or training. Since government attention is a scarce good, public-private collaboration focused on alleviating constraints faced by specific sectors or groups of firms, such as deliberation councils or business-government roundtables, also counts as industrial policy.

Since industrial policy, by design, favors certain types of economic behavior, it typically comes with some kind of conditionality. Conditionality can be of a limited kind, restricted to ex-ante eligibility criteria. For example, only firms with less than a certain number of employees in a particular sector or region might be able to receive the proffered subsidy. Alternatively, the incentives may be conditioned on ex-post behavioral changes, such as undertaking specific investment or eventually employing a target number of

workers. In the latter case, there might be explicit, quantitative performance criteria or a softer, iterative form of monitoring to ensure broad compliance.

## 2.2. Rationales

There is no shortage of well-grounded economic rationales for industrial policies. We can summarize those under three broad headings: externalities, co-ordination (or agglomeration) failures, and public input provision. The first two are related to well-known market failures, while the third has to do with the specificity of public inputs to particular economic activities. We say a few words about each rationale.

Externalities. Economic activity produces positive externalities when it generates benefits elsewhere in society that are not recouped in the revenues generated for those who carry out the activity. Learning externalities are common and widely recognized in the economics literature. The learning in question may be about how to produce a good or service more efficiently, as in the case of R&D or learning-by-doing spillovers across producers. It may also be about general cost and demand conditions for new goods in which such conditions are uncertain and initial entrants produce valuable information to subsequent entrants (e.g., cost-discovery externalities à la Hausmann & Rodrik 2003). But externalities can take other forms as well. There are national security externalities when reducing dependence on a foreign source of supply – for example, rare earths or semiconductor inputs – makes a country as a whole more secure. This is a social benefit that is not fully internalized in the input sourcing decisions of individual firms. Similarly, there are good-jobs externalities when creating middle-class jobs produces greater social cohesion and alleviates social ills such as crime and drug addiction (Rodrik & Sabel 2022).

Coordination (or agglomeration) failures. This category of market failure refers to situations in which the profitability of an individual producer depends on the level of related economic activities undertaken by others. The related activities may be goods and services that are complements in demand or production, or downstream and upstream activities. Coordination failures typically require scale economies of some sort. Such situations may produce multiple equilibria in the absence of government intervention. Consider a case where it would be profitable to produce good *A* when good *B* is also produced and not otherwise. Assume that the same situation holds true symmetrically for good *B* as well. In one equilibrium, neither good is produced. In the other, both goods are. If the social value of production of the two goods exceeds their opportunity costs, an economy can get stuck in what is a sub-optimal equilibrium. Government policy can help push producers to the superior equilibrium.

Activity-specific public inputs. Private production depends on the provision of public goods, such as law and order, appropriate regulation, education, and infrastructure. Economists generally think of such public goods as “horizontal” policies that do not prioritize certain activities or entail choice and discretion, and produce across-the-board benefits. Often public goods are indeed so. But in many real-world contexts, the needs of producers are highly specific to the nature or location of their activity. To take a simple example, infrastructure dollars can be spent on building/enlarging a port or on expanding the road transport network. Depending on the choice that is made, different kinds of producers reap the benefit. If the decision is to build a port, it can be located close to the copper mine, the steel complex, or a prospective green hydrogen facility. Similarly, worker skills are highly specific to the needs of different sectors, and the government has to decide what kind of professional training it should prioritize. In such situations, governments are essentially “doomed to choose” (Hausmann & Rodrik 2006) – to select

which activities are more deserving of public goods. When the choice is made consciously, the result is industrial policy as we have defined it.

These three rationales address different kinds of problems and call for diverse remedies. Externalities are best handled by Pigovian subsidies that are targeted at source and induce firms to internalize the value of the spillovers they produce. Coordination failures in principle do not need subsidies or financial incentives to be addressed: governments could simply bring the different groups of firms together and call on them to make simultaneous investment commitments. Alternatively, prospective investors could be provided with public guarantees (that will not need to be paid out if, as predicted by the logic of the problem at hand, the investments turn out to be profitable *ex post*). Finally, public inputs require the provision of specific investments by the government itself. These first-best remedies aside, there is an almost endless list of second-best instruments which could tackle these problems at least partially, even if not equally well.

### **2.3. The critique**

As this discussion makes clear, the theoretical case for industrial policy is broad and strong. The controversy over industrial policy revolves not around these rationales – which are generally well-accepted among economists – but around two practical objections. One of these objections is about information shortcomings, the other about political capture. The informational critique asserts that even if the market failures on which governments could act are widespread, real-world governments are unlikely to know enough about the location and magnitude of these failures to make the correct decisions. The political critique asserts that even if governments have (or could acquire) the relevant information, industrial policy opens the door to self-interested lobbying and political influence activities, diverting the government into activities that enrich private interests without enlarging the social pie. For either or both reasons, the argument goes, “governments cannot pick winners.”

This is the kind of debate which empirical analysis might have helped resolve. Indeed, proponents and opponents of industrial policy have relied on their favorite examples to make their case. For proponents, the economic miracle experienced by countries such as Japan, South Korea, Taiwan, China – where industrial policy has been rampant – clinches the case that industrial policy is an important component of a successful growth policy. Opponents point to the disappointments with industrial policy in Latin America and Africa under import-substitution and to particular instances of costly public initiatives (such as the Franco-British Concorde or the Malaysian Proton car). The proponents respond that the failures have more to do with implementation weaknesses and argue besides that many ISI experiences (such as those in Mexico, Brazil, Turkey) were successful for quite some time. The opponents counter by arguing that East Asian countries could have done even better without industrial policies, or that they are special cases that cannot be generalized (because East Asian countries have especially competent bureaucrats and “hard” states that could effectively discipline the private sector).

The task of providing more systematic evidence on whether industrial policy works has been stymied by a number of complications. First, there is a dearth of cross-nationally comparable quantitative measures on industrial policy interventions. Traditionally, empirical studies have tended to focus on straightforward measures such as import tariffs or credit subsidies, which may be only one component of industrial policies and may serve many other objectives besides structural transformation. This shortcoming of measurement is gradually being remedied in more recent studies, as we will discuss below.

Second, there is the difficulty of ascertaining success. Most studies focus on whether an intervention moved the needle on a quantity of interest, such as investment, exports, or TFP. This is a test of “effectiveness,” and at best captures only part of the picture when it comes to determining success. Successful industrial policy needs not only to have accomplished the targeted structural change, but do so (a) in a way that truly alleviates the underlying market failures, and (b) without causing too many distortions elsewhere in the economy. If, say, investment in steel is boosted but the market failures are in pharmaceuticals instead, industrial policy will be effective but inefficient. Performing the complete analysis is quite difficult in practice – even ex post – since market failures are rarely observable directly.

This would seem to strengthen the critics’ position and reinforce the argument that “governments cannot pick winners.” But ultimately what is required for industrial policy to work is far less than a consistent ability to pick “winners.” In the presence of uncertainty, both about the effectiveness of policies and the location/magnitude of externalities, the ultimate test is not whether governments can pick “winners,” but whether they have (or can develop) the ability to let “losers” go. As with any portfolio decision, it would be an indication of sub-optimal policy if the government did not back some ventures that end up as failures ex post. In the U.S., Department of Energy loan guarantees to Solyndra, a solar cell manufacturer, failed miserably, but a similar loan guarantee to Tesla enabled the company to avert failure and become the behemoth it is today. In Chile, successes in four projects supported by Fundación Chile – including most spectacularly salmon – is said to have paid the costs of all other ventures.

Letting losers go may still be a hard task, in light of political pressures that inevitably develop. Indeed, Solyndra, for example, was backed by the government long after it became clear that the company would not become financially viable. But it is far less demanding than governmental omniscience. Ensuring that governments can stop backing evident losers requires a set of institutional safeguards that include clear benchmarks, close monitoring, and explicit mechanisms for reversing course. We will return to the institutional underpinnings of successful industrial policy later in the paper.

Beyond practical critiques of industrial policy, there have also been some more technical critiques. Bartelme et al. (2021) reconsiders the textbook Pigovian case for intervention in a quantitative trade model with increasing returns to scale. Although the authors find sizeable external economies of scale across sectors, their framework suggests production subsidies may only promote modest aggregate welfare gains. The quantitatively small effects are driven by inelastic demand, which implies industrial policy produces little structural change. (Incorporating input-output linkages produces larger aggregate effects.) Related work by Laskharipour & Lugovsky (2023) uses a quantitative trade framework to highlight another potential trade-off: the allocative gains from unilateral interventions may be undone by terms-of-trade losses. Globally coordinated industrial policies, on the other hand, can potentially promote large gains.<sup>3</sup>

Together, this work illustrates how new theory-driven quantitative work can provide a richer analysis beyond the traditional either-or debate. The study of industrial policy requires a combination of both careful empirics and theory. We return to this point in Section 4.8.

---

<sup>3</sup> See also Itskhoki & Moll (2019) and Liu (2019), which formalize the classic case for industrial policies in distorted economies in a quantitative setting.

## 2.4. Where does this leave us?

Industrial policy is not that different from many other domains of policy where there exist clear theoretical justifications for government intervention, but evidence on what works is not that clearcut. One can list education policy (human capital externalities), health policy (moral hazard, adverse selection), social insurance and safety nets (incomplete risk markets, behavioral factors), infrastructure policy (natural monopoly), and stabilization policy (Keynesian “rigidities”). In all these areas, it is recognized that the market-failure arguments for intervention can be exploited by powerful insiders and overwhelmed by informational asymmetries. Similarly, the efficacy of different remedies remains contested, despite rich empirical literatures in each domain. But policy discussions in these areas typically focus not on *whether* the government should do it, but on *how*. The debates revolve around what works and under what conditions.

## 3. What do we know about the practice of industrial policy?

Until recently, we lacked a systematic understanding of industrial policy practice, and data played a key role in this deficit. This section discusses why measuring industrial policy is difficult and shows how new work has overcome many well-known challenges. We then characterize broad features of current industrial policy practice, globally.

### 3.1. Dilemmas of measurement

Industrial policies are complex and measuring them in the wild can be complicated. To illustrate these challenges, consider the complexity of a single policy: China’s recent push in the shipbuilding industry. China’s 11th National 5-year Economic Plan for 2006-2010 identified shipbuilding as a “strategic industry.”<sup>4</sup> With the goal of becoming the largest shipbuilding nation within a decade, China deployed a multitude of policy instruments, including production subsidies, investment subsidies, and entry subsidies. There were also changes along the way. In 2009, the government announced the Plan on Adjusting and Revitalizing the Shipbuilding Industry, which turned policy away from promoting entry and instead focused on industry consolidation. From 2013, the government periodically considered which firms met standards to receive priority incentives.

The example reveals why measuring industrial policy can be difficult, especially at scale. A single sectoral strategy, such as China’s shipbuilding push, often entails many tools, and the composition of these tools can change over time. Observing these different levers is challenging. Industrial policy may consist of flows of government funding in the form of subsidies and financial grants. Such flows may be tracked through fiscal expenditures subject to data availability. Alternatively, industrial policy may also be conducted through tools such as tariffs or export restrictions. These tools do not entail fiscal expenditure but instead try to shift relative prices in favor of some activity and require different types of data. Other industrial policies may consist of fundamentally different policies. For instance, development planners can set expectations and act to coordinate the behavior of firms. This type of “administrative guidance” is a fixture of industrial policy in many countries (Johnson 1982), yet these may be some of the most

---

<sup>4</sup> This description is based on Kalouptsi (2018) and Barwick et al. (2021).



difficult policies to monitor and quantify. In short, the toolset is vast, and some tools are easier to observe than others.

A related measurement challenge arises from the fact that we cannot assume that a specific policy tool is always used as an industrial policy. In fact, assuming a tool is necessarily an industrial policy can conflate various policy goals. The best-known example of this problem is the import tariff. While a tariff can and is used for industrial policy purposes, there are many other motivations for deploying tariffs, from raising fiscal revenue (Cagé & Gadenne 2018) to special interest politics (Goldberg & Maggi 1999) and terms of trade rationales (Broda et al. 2008). Thus, even if we could track all policy levers *potentially* used as industrial policy, only a subset will *actually* be used for industrial policy.

Given these nuances, the literature has struggled to quantify the practice of industrial policy, especially beyond individual episodes. For example, an early empirical literature (section 4) sought to explore the effects of industrial policy using tariffs and other simple measures of economic openness as proxies. Although problems with this literature have been thoroughly discussed (Rodriguez & Rodrik 2001, Harrison & Rodriguez-Clare 2010), focusing on easily observable measures, such as tariffs, has left a lasting confusion in the literature.

In particular, some scholars have mistakenly taken industrial policy to be synonymous with protectionist trade policy, where industrial policy implies the overt protection of domestic markets. This is not without some merit, of course. Industrial policy may take the form of protectionism. Indeed, this was certainly the case historically, with classic instances of import substituting industrialization (ISI). However, seeing industrial policy as tantamount to protectionism can omit contemporary forms of industrial policymaking in open economies, where interventions look different from the protectionism of the past. We show below that, more often than not, contemporary policy takes the form of promoting outward-oriented economic activity, say by promoting export activity.

In fact, it would be a mistake to think that industrial policy disappeared alongside globalization. Instead, practice has evolved alongside globalization – and perhaps survived despite it. The ascent of the East Asian Tiger economies belies a clean correspondence between protectionism and industrial policy. Notably, 1960s South Korean industrial promotion under Park Chung-hee targeted export activity (Westphal 1990), and measures of trade openness increased through its most conspicuously interventionist periods (Lane 2022). FDI has emerged as a tool for industrial policy targeting (Harrison & Rodriguez-Clare 2010, Harding et al. 2019).<sup>5</sup> Following liberalization, China adapted the strategies of South Korea and Japan to a more globalized world, famously wielding incentives and controls on FDI as tools of industrial policy in the post-WTO world (Thun 2004). In short, measures of protectionism cannot be taken as reliable proxies for industrial policy, especially when it comes to characterizing contemporary practice.

A distinct issue arises with other early studies which used government spending as a proxy for industrial policy interventions, especially in East Asia. Governments spend money on different sectors for a multitude of reasons, only one of which is industrial policy.

---

<sup>5</sup> Where FDI is usually associated with liberalization, strategic use of FDI is compatible with our definition of industrial policy. See Bai et al. (2019) and DiPippo et al. (2022) for the case of Chinese joint ventures.

Finally, globalization has likely shifted the way industrial policy is practiced towards policies that are more difficult to observe.<sup>6</sup> Many forms of government intervention may violate countries' commitments to multilateral institutions such as the WTO or supranational institutions such as the EU. Consequently, policymakers may choose measures that evade detection more easily, making measurement even harder.

### 3.2. Accounting for contemporary industrial policy practice

The return of industrial policy has brought attention to the paucity of systematic data, and scholarship has started to take stock of global practice. Recent work has clarified conceptual issues surrounding industrial policy and provides a glimpse into contemporary practice. A number of ambitious efforts have been launched to measure industrial policies through deep accounting of government activity (DiPippo et al. 2022, Criscuolo et al. 2022, Hanson & Rodrik 2023). A central mission of these projects is distinguishing government expenditures allocated for industrial policy aims. For example, an extensive effort by the Center for Strategic and International Studies (DiPippo et al. 2022) has produced quantitative descriptions of industrial policy for China and seven other economies (Brazil, France, Germany, Japan, South Korea, Taiwan, and the U.S.). DiPippo et al. estimate that in 2019 China spent about 1.5% of its GDP on industrial policy versus 0.3-0.7% of GDP in the other economies. These findings suggest that among these countries, industrial policy is an important part of the policymaker's toolkit.

In the same vein, large-scale, institutional efforts by the OECD, such as Criscuolo et al. (2022), have developed cross-country methodologies to quantify industrial policy using government expenditure data for a sample of OECD member countries. In addition to quantifying financial flows, Hanson & Rodrik (2023) provide comparative evidence across U.S. regions on the organizational landscape of place-based policies, tracking single-purpose entities such as workforce development agencies as well as coordinating agencies such as local economic development agencies. Together, these efforts go far beyond the early accounts of industrial policy, highlighting the depth and variation of industrial practice in major economies.

A different approach is taken by Juhász et al. (2022), who use natural language processing to classify industrial policies at a high resolution (country-industry-year level) using a publicly available policy inventory (the Global Trade Alert database or GTA, Evenett 2009).<sup>7</sup> The core idea of Juhász et al. is that textual descriptions of policy often convey information about the *objectives* of political actors, and thus allow researchers to identify whether a policy has industrial policy goals versus alternative objectives.

To illustrate the approach, consider the description of a Chinese subsidy scheme from their database: "In the PRC Ministry of Industry and Information Technology's policy released on the 1st of March 2017, a plan is laid out *to boost growth in the Chinese battery industry*, specifically, batteries for automobiles

---

<sup>6</sup> Although import tariffs have come down worldwide, there is evidence that trade policy has persisted through other (non-tariff) means (Ray & Marvel 1984, Bown & Tovar 2011, Limao & Tovar 2011, Bown & Crowley 2014, Garred 2018).

<sup>7</sup> Rather than taking a more macro accounting approach, the outputs of this method are similar to coverage ratios used by trade scholars. Thus, this shares commonalities to a literature measuring non-tariff measures.

[...]”. Here, the italicized text reveals the objective of the policymaker, which is textbook industrial policy designed to shape the composition of what is produced domestically.

Juhász et al. (2022) operationalize this idea, applying a supervised machine learning algorithm to classify policy descriptions from the GTA database. Applying this algorithm to the entire database allows them to construct a global dataset of industrial policies from 2010 to 2022. The authors train both simple logistic regression models and large language models (LLMs; specifically different versions of BERT) using training data from a random subset of GTA policies hand-classified by humans. Both types of models perform well. The simplest binary logistic regression model that classifies policies into two mutually exclusive categories – “industrial policy goals,” and “not industrial policy goals” – classifies 96% of unseen data correctly (the *accuracy* measure). The advantage of logistic regression is its transparency, yielding easily interpretable results. Table 3.1 shows the ten features (words) that are most predictive of industrial policy, based on the binary logistic classifier. Words such as “export,” “boost,” “research,” “development,” and “technology” are all reasonably associated with industrial policy. On the other hand, the advantage of BERT is its better performance, particularly when it comes to predicting multiple classes (distinguishing “industrial policy goals,” “not industrial policy goals,” and “not enough information”).<sup>8</sup> Both classifiers perform well, but the large language model receives a performance boost from its ability to pick up nuance and context.

This text-based approach overcomes a number of challenges associated with measuring industrial policy, reviewed in Section 3.1. First, given the broad range of policy measures covered by the Global Trade Alert (GTA), it encompasses a vast array of policy levers that can be used for industrial policy purposes. Notably, the GTA includes both trade policy measures that entail no fiscal spending, as well as those that do. There are also cases where coordination policies show up in the data. For example, the Chinese consolidation policy in shipbuilding that we described above (Plan on Adjusting and Revitalizing the Shipbuilding Industry) is a measure that is captured in the GTA and is classified by the model as industrial policy. Second, within policy measures, it is possible to distinguish industrial policy from other reasons for intervention; for example, a subsidy used for industrial policy versus one used to accommodate macro-economic shocks.<sup>9</sup> Third, the model is able to classify policies at scale using off-the-shelf textual data. However, an important limitation of this approach is that it yields count-based measures of industrial policy interventions. This means that a subsidy to a single firm and a major sectoral policy such as China’s consolidation of its shipbuilding sector can both count as one policy in the GTA. This is an important limitation to bear in mind, particularly when comparing these measures across countries. Nevertheless, a nontariff measure literature in trade policy has illustrated how count-based nontariff measures can be transformed into ad valorem equivalencies (AVEs) using various methods.<sup>10</sup>

---

<sup>8</sup> The simple binary logistic classifier collapses “not enough information” and “not industrial policy goals” into one class: “not industrial policy goals.” When extending the logistic model to multiple classes, the performance of the model worsens slightly, and this is where BERT’s superior performance becomes evident. The three-class logistic model’s accuracy is 88% vs. BERT’s 94%.

<sup>9</sup> In fact, the authors find that including variables for the type of policy (e.g., tariff, subsidy, etc.) as a predictor of industrial policy goals does not improve model performance, consistent with the notion that policy measures themselves reveal little about the objectives of the policymaker.

<sup>10</sup> See recent gravity-based methods proposed by Heid et al. (2021) and Herman (2022).

### 3.3. A first look at international industrial policy practice

The findings from Juhász et al. (2022) give us an initial glance at the global practice of industrial policy, with particularly good coverage of G20 nations.<sup>11</sup> First, industrial policy has indeed returned and is on the rise. Figure 3.1 shows an increase in the total count of industrial policy interventions through the 2010s, with major accelerations in 2018 and 2021. This trend does not seem to be driven by data coverage from the GTA improving over time, as the share of industrial policies among all policies in the GTA also increased. The increase in industrial policy is also reflected in the shares.

Second, higher income (e.g., G20) countries are major users of industrial policy. While industrial policy is practiced the world over, it may in fact be more prevalent in higher income countries. Figure 3.2 plots the breakdown of industrial policies by income quintile (based on GDP per capita in 2010) and region. The figure makes clear that advanced economies account for the overwhelming majority of industrial policy interventions. Outside this group of countries, industrial policy is more evenly distributed across regions. Given that most of DiPippo's 2019 sample measuring industrial policy spending overlaps with where Juhász et al. (2022) find the most intensive use of industrial policy interventions, we have a rough way of benchmarking how much fiscal expenditure industrial policy interventions may entail. Of the seven economies outside of China examined by DiPippo et al. (2022), all except South Korea and Taiwan are on the list of the top ten users of industrial policy based on Juhász et al.'s count-based measure. This suggests that the fiscal expenditure of industrial policy for some of the heaviest users is in the ballpark of 0.3-0.7% of GDP. The findings across the two studies paint a consistent picture: industrial policy is relatively important, particularly in advanced economies. Notably, these numbers (for 2019) predate the current surge in industrial policy spending (e.g., post CHIPS and IRA acts in the U.S.).

It is impossible to ignore the pattern in Figure 3.2, which suggests that the number of industrial policy interventions increases with income per capita. This finding is sensible in light of the fact that the authors also find that contemporary industrial policy is typically conducted through subsidies and export promotion measures (see below). In lower-income economies, fiscal constraints will bind fast if industrial policy is deployed through government spending. However, cross-country comparisons using a count-based measure cannot tell the entire story, nor is the GTA's sample complete. Thus, we interpret the patterns in Figure 3.2 as suggesting a potentially steep income gradient in industrial policy practice.

Juhász et al.'s findings support the idea that modern industrial policy is complex and often composed of many distinct policy levers that are outward-oriented; this finding is perhaps most relevant to the discussion in the preceding section (3.2). Different forms of subsidies and export-related measures, together, account for most industrial policy interventions across the income distribution, constituting almost 90% of all interventions. This underscores the point that modern industrial policy is expensive. Interestingly, countries across the income distribution tend to use many of the same policy levers.

Figure 3.3 shows the type of industrial policies used across different income groups and plots the 10 most frequently used policy tools in high-, middle-, and low-income countries, respectively. Trade financing, a policy measure that facilitates exporting, is ubiquitous across the income distribution. Interestingly, local value-added incentives (a trade related investment measure) are the second most used tool in low-income countries, accounting for 16% of industrial policy interventions. Developing

---

<sup>11</sup> The data presented here are from version 2 (v2) update of Juhász et al. (2022). Notes on v2 are available at <https://osf.io/tp6ak>.

countries seem to be trying to harness FDI in ways that would increase local linkages within the domestic country; this is compatible with the literature on FDI and industrial policy (see Harding et al. 2019). These patterns also highlight the fact the industrial policy can often be outward oriented.

Notably, import tariffs are not a prominent industrial policy tool, at least in the GTA and Juhász et al. data. Tariffs account for 1.3% of industrial policy interventions and 3% in the bottom two quintiles of the income distribution. Although tariffs are not irrelevant, modern industrial policy would be poorly approximated with tariffs.

What, then, do policies target? Juhász et al. (2022) find overlap across the income distribution. Figure 3.4 plots the top 10 sectors targeted by industrial policy in high-, middle-, and low-income economies, respectively. Sectors targeted by all three groups of countries include machinery and transport equipment (e.g., electrical machinery or vehicles). Given the vertical fragmentation of the value chain across international borders, it is quite possible that industrial policy is targeting different parts of the value chain in which countries already specialize. Notably, automobiles and capital equipment are clear examples where fabrication and assembly have been offshored to middle-income countries (Baldwin & Ito 2021). Likewise, industrial policy in some lower-income countries may also be attempting to move countries into different, higher value-added parts of the value chain, consistent with the prominent use of local value-added incentives in poor countries.

Figure 3.4 also shows important differences in the sectors targeted by vertical industrial policies. In the lowest quintiles of the income distribution, textiles and apparels receive a sizeable share of industrial policy, consistent with similar strategies that have been historically pursued at early stages of structural transformation. Interestingly, in middle-income countries, agriculture and livestock receive some industrial policies. Food and agriculture are also targeted by industrial policy in high-income economies, though relatively less.

In high-income countries, the starkest difference is the prominence of industrial policy targeted towards clean electricity generation. This may not jump out to the casual observer, as these policies fall under HS code 27, which, unfortunately, also includes production of energy using hydrocarbons. This 2-digit sector is more heavily targeted by high-income countries: 21% of policies target this sector in high income economies, compared to eleven to 12% in middle- and low-income countries. If we break down targeting within this sector and isolate clean electricity generation, this difference becomes even starker. 56% of policies in rich countries target clean electricity generation within the HS 27 sector, while only 42% (31%) target clean electricity generation in middle-income (low-income) countries. These patterns suggests that the green energy transition is being more heavily promoted by industrial policy (both in levels, and in relative terms) among high-income economies.

How do these findings inform research and thinking about industrial policy? First, despite skepticism and the lack of empirical study among academic economists, industrial policies are not uncommon – perhaps far from it. As of this study, Juhász et al. (2022) find that over a third of policies in the GTA commercial policy database qualify as industrial policies. In their studies of higher income countries and China, DiPippo et al. (2019) and Criscuolo et al. (2022) both estimate that a substantial share of GDP goes toward industrial policy. Figure 3.1 and recent evidence suggest that these trends may have been in motion before the 2020 revival of industrial policy. Even if the current embrace of large-scale industrial policy proves to be a temporary fad, it is unlikely to fade into obscurity anytime soon. We urgently need research that can inform policymakers of how to do industrial policy well, or less poorly.

Second, the picture of industrial policy that emerges brings systematic data to the assertions that industrial policy and outward oriented development strategies can go hand in hand. While in the past, industrial policy may have been more closely associated with import tariffs and trade protection, these results portray a considerably more variegated picture. Moreover, the evidence reinforces a point made repeatedly in the trade policy literature (Goldberg & Pavcnik 2016): that research needs to move away from policy instruments that can be readily measured (tariffs), to other measures that may be harder to capture but that reflect actual current practice. For industrial policy, this rings particularly true.

#### 4. Industrial policy meets causal identification

Until very recently, credible empirical evidence on the effects of industrial policy was virtually non-existent. However, the past few years have seen a proliferation of papers that use careful research designs to evaluate different forms of industrial policy. Though long overdue, the credibility revolution (Angrist & Pischke 2010) has finally arrived to research on industrial policy. In this section, we discuss the growing set of papers that use the toolkit of modern econometrics to evaluate industrial policy. We do not provide a comprehensive overview of current empirical research.<sup>12</sup> Rather, our goal is to evaluate where recent papers have made progress on empirical challenges and discuss where the largest gaps in our knowledge remain.

Empirical evaluation of industrial policy needs to contend with the inherent endogeneity of policy. By inherent, we mean that, *by design*, the policymaker is not acting *randomly* in targeting the promotion of certain activities. As a result, not-targeted units are generally unlikely to serve as credible counterfactuals. This issue presents a serious obstacle to studies of industrial policy that use observational data and also raises challenges for more modern, research design-based empirical methods.

In this section, we introduce these empirical challenges using a simple model of government behavior that embeds different stories about the nature of externalities as well as of government motivations and capabilities (the framework is adapted from Rodrik 2012). We then discuss the issues that arise with traditional correlational approaches, as well as modern causal identification techniques. We then turn to discussing how new empirical work has tackled these challenges and what insights they have yielded.

##### 4.1. The empirical challenge of evaluating industrial policy

We express the underlying level of economic performance ( $g$ ) as a negative function of a market failure parameter  $\theta$  (which lies between 0 and 1):

$$g(\theta) = (1 - \theta)A,$$

where  $A$  stands as some state variable that affects performance. For example, in a linear endogenous growth model (where  $g$  would stand for the rate of economic growth),  $A$  would be the level of productivity. Alternatively,  $g$  could be the rate of investment and  $A$  some variable linked to the (social) profitability of capital accumulation. The basic point here is that the greater the market failure  $\theta$ , the

---

<sup>12</sup> Some forms of industrial policy such as different forms of firm support policies (Slattery & Zidar 2020, McKenzie et al. 2021) or innovation policies (Bloom et al. 2019) have their own, excellent recent reviews. More generally, Harrison & Rodriguez-Clare (2010) remains an excellent resource.

larger the gap between social and private returns, and the lower the level of economic performance (“growth”) absent government intervention. The basic unit of observation could be a firm, group of firms, industry, or entire country. We will omit subscripts to denote different units (“industries”) to avoid excess notation.

Let the government’s policy be a subsidy  $s \in [0,1]$  which alleviates the distortion by closing the gap between private and social returns to  $(1 - \theta)(1 - s)$ . We assume the use of the subsidy comes with some fiscal or agency cost  $\varphi\alpha(s)$ , where  $\alpha(s)$  is a rising and convex function of  $s$  ( $\alpha(0) = 0$ ,  $\alpha'(s) > 0$ , and  $\alpha''(s) > 0$ ). The parameter  $\varphi$  is meant to capture the ability of the government to intervene effectively (“government capacity”), with higher values of  $\varphi$  indicating lower government capacity. With the policy in place, the modified expression for growth is:

$$g(s, \theta, \varphi) = (1 - \theta(1 - s))A - \varphi\alpha(s).$$

The socially-optimal level of the subsidy is denoted  $s^{soc}$  and is given by the value of  $s$  that solves  $g_s(s, \theta, \varphi) = \theta A - \varphi'\alpha(s) = 0$ . Note that it would not be socially optimal to set the subsidy at a level that would fully offset the market failure (which is given by  $s = 1$ ). The cost term  $\varphi\alpha(s)$  and its convexity ensures that  $s^{soc} < 1$ .

So far we have considered only the economic motives for the subsidy. The government can also have political motives. The subsidy can be used, for example, to channel public resources to favored or politically-connected groups. We use the function  $\pi(s)$  to represent the purely political benefits to the government of using the subsidy. This is a single-peaked, concave function with  $\pi(s)$  rising for small  $s$  and maximum-value defined by  $\pi(s^{pol})$ .

We model government behavior by assuming it maximizes a value function  $u(s; \theta, \varphi)$  where economic and political objectives both enter. Denoting by  $\lambda$  the relative weight placed on the economic motive, the decision problem is:

$$\max_s u(s; \theta, \varphi) = \lambda g(s, \theta, \varphi) + \pi(s).$$

The government chooses level of subsidy  $s^{gov}$  which is the solution to the first-order condition

$$\lambda g_s(s, \theta, \varphi) + \pi'(s) = \lambda[\theta A - \varphi\alpha'(s)] + \pi'(s) = 0.$$

With this theoretical background in place, let us now consider how observational data on government behavior and economic outcomes can inform us about how and whether industrial policy works. Ideally, empirical evidence ought to help us distinguish among three different, contending positions on industrial policy:

- the “developmentalist” view: governments can successfully identify and support growth/efficiency-enhancing firms/industries;
- the “inefficacy” view: governments seek growth/efficiency but do a poor job of supporting appropriate activities;
- the “rent-seeking” view: governments are beholden to special interests and do not seek desirable economic outcomes.

Remember that we do not directly observe the critical parameters  $\theta$ ,  $\varphi$ , and  $\lambda$  – the market failures, government capabilities, and relative importance of political motives, respectively. What we do observe is how economic performance ( $g$ ) varies with the level of government intervention ( $s$ ).

#### 4.2. Conceptual issues with early evidence on industrial policy

The first-generation of empirical studies on industrial policy used this correlation to infer industrial policies had been generally ineffective or counter-productive: the level of subsidies or protection, these studies found, were generally negatively correlated with measures of performance such as productivity across industries (e.g., Krueger & Tuncer 1982, Harrison 1994, World Bank 1993, Lee 1996, Beason & Weinstein 1996, and Lawrence & Weinstein 2001). Moreover, this was true even in countries like Japan and South Korea – countries where a wide range of qualitative evidence suggested industrial policies had been put to good use.

However, as our theoretical framework makes clear, the level of government intervention is endogenous, responding to a variety of economic, administrative, and political determinants. When this endogeneity is not explicitly taken into account, the inferences can be misleading. To isolate the effects of these determinants, let us consider each in turn.

Suppose first that the sole driver of subsidies is the varying importance that the government places on political motives across different industries ( $\lambda$ ). This would epitomize the case of the purely “rent-seeking government.” Using the model just sketched out, it can be checked that  $\frac{ds^{gov}}{d\lambda} < 0$  and  $\frac{dg}{d\lambda} > 0$ . In words, the greater the weight on politics (the lower the  $\lambda$ ), the larger the subsidy and the lower the rate of economic growth. An analyst who can observe only  $s^{gov}$  and  $g$  would note a negative correlation across industries between these two. This result would be consistent with the hypothesis that the government is operating in a rent-seeking mode.

Next consider the alternative hypothesis that the sole determinant of subsidies is the presence of market failures ( $\theta$ ). Under this hypothesis the government is a social-welfare maximizer, and differences in the government’s exercise of industrial policy would be driven by the variation in the importance of market failures across industries. The comparative statics with respect to  $\theta$  yields the following results:

$$\frac{ds^{gov}}{d\theta} = \frac{A}{\varphi\alpha''(s^{gov})} > 0,$$

$$\frac{dg}{d\theta} = -(1 - s^{gov})A < 0.$$

(Since the government is emulating the social maximizer in this case,  $s^{gov} = s^{soc} < 1$ .) The first of these results states that intervention levels are higher where the market failures are greater, as expected. The second result is that growth rates are lower where the market failures are larger. This follows from the fact that socially optimal policy does not fully offset the market failures, and performance still lags in firms/industries that are subject to larger market failures. Putting the two results together, we find that subsidies would be larger for those activities that are doing worse. Subsidies and performance are negatively correlated even though the government acts a social-welfare maximizer! This is exactly the same pattern of correlation as in the case of the purely rent-seeking government. (Note that the use of industrial policy is clearly welfare-enhancing in this case in the sense



that an ex-ante rule that forced the government to give up on subsidies and set  $s^{gov} = 0$  would leave the economy worse off.)

For completeness, consider the pure “efficacy” case where the sole driver of industrial policy is now government capacity. In this instance, differences in the use of subsidies would be due entirely to differences in the ability to implement them, captured by  $\varphi$ . It can be checked in similar fashion that  $\frac{ds^{gov}}{d\varphi} < 0$  and  $\frac{dg}{d\varphi} < 0$ . In words, the level of subsidies and growth are both decreasing in government capacity. This would produce a cross-sectional correlation between subsidies and growth in observational data that is positive. Discovering a negative correlation could refute the hypothesis that governments respond systematically and optimally to capacity constraints – but only under the extreme assumption that those constraints are the only variation in the data that could be driving government behavior and that other fundamental political and economic motives can be ignored.

To summarize, the two polar cases of rent-seeking and developmental governments – and anything in between – are observationally equivalent. It is not possible to say anything about the merits of industrial policy from the pattern of correlation between the extent of policy intervention and economic performance. We need explicit models of government behavior to evaluate industrial policy. Observational data can be quite misleading absent structural models that can be otherwise verified.

### 4.3. Challenges for causal identification

Since the problem here is the endogeneity of government policy, one alternative is to focus on identifying causality through exogenous or random sources of variation in government action. But this does not entirely resolve the difficulties in evaluating industrial policy. Consider the canonical empirical exercise whereby the researcher is able to extract an exogenous component of subsidies. In terms of the model above, we could express  $s^{gov}$  as  $s^{gov} = s^{gov*} + \varepsilon_s$ , where the first component is the systematic part that responds to unobservable economic, political, and administrative determinants that may also influence growth, while the second term is the orthogonal component. The analysis would now focus on correlations between  $\varepsilon_s$  and  $g$ . Since  $\varepsilon_s$  is by definition orthogonal to any factors that may simultaneously affect  $s^{gov}$  and  $g$ , this exercise would yield the causal effects of subsidy “shocks” on economic performance.

In essence, what we will be uncovering in this instance is the consequences of the government randomly sprinkling subsidies of varying sizes on different parts of the economy. This hardly resolves the question of whether subsidies are likely to work under real-world circumstances. Remember that proponents and opponents of industrial policy not only posit that governments respond to economic and political considerations, but that the nature of that response is crucial to their arguments. Proponents say market failures can be targeted sufficiently well and that administrative and political difficulties are surmountable. Opponents say the political considerations will drive interventions and/or administrative obstacles will overwhelm any good intentions.

As an illustration, consider a well-identified study which finds a positive result in the sense that (the exogenous component of) industrial policy produces desirable economic outcomes. An opponent of industrial policy could justifiably argue: “Yes, I can see that the results were favorable in this instance, but in practice the selection of projects/industries/regions by industrial policy will be hardly random; it will be driven by politics, lobbying, and rent-seeking, and these results do not speak at all to these

difficulties.” Conversely, suppose the study yields a negative finding, with (random) interventions producing adverse economic results. Now the proponent of industrial policy can legitimately argue: “Yes, but those of us who favor industrial policy never advocated that the subsidies should be deployed randomly! We always said selection of projects/industries/regions should be done after careful economic analysis and consultations with stakeholders, and these results say very little about the likely consequences when such processes are followed.” Statistically well-designed studies might not convince critics of either kind.<sup>13</sup>

#### 4.4. The empirical path taken

The preceding discussion risks presenting industrial policy as inherently unsuited to evaluation using the modern empiricist’s toolkit. That would be misleading. As we noted previously, industrial policy is not that different from many other domains of policy with clear justifications for policy intervention and ambiguous evidence. In those areas, both research and the policy debate revolve around what works and under what conditions. Consider, for example, innovation policy, a mature field that overlaps with industrial policy. A recent study by Bloom et al. (2019) is devoted exclusively to assessing the effectiveness of different innovation policy levers informed by rich empirical evidence.

In industrial policy, the emerging literature has taken a somewhat different path. Much of the new generation work we survey below is devoted to evaluating whether industrial policy can elicit the desired behavioral response in highly stylized settings. Consider Juhász (2018), which evaluates the famous infant industry argument using the disruption to trade resulting from a blockade against Britain in the 19<sup>th</sup> century. The paper is obviously of little use to a policymaker trying to understand how to effectively promote infant industry. However, what the paper shows is that infant industry can be a powerful economic mechanism in the real world. In the language of the model above, it suggests that the market failure,  $\theta$ , can be large. That finding is arguably of use in a debate centered on *whether* the government ought to do industrial policy.

Much of the new work we survey below has a similar flavor. Similar to Juhász (2018), many papers make an attempt to isolate different “layers” of treatment. We think of one layer being the question of whether the justification for intervention is valid. For example, is it really the case that external economies of scale are present and prevent the industry from developing under *laissez-faire* (i.e., is  $\theta > 0$ )? We think of this as evaluating the “economic mechanism,” which may or may not be at work. A second set of questions involves evaluating the efficacy of “implementation:” was the policymaker able to identify the right unit to treat; were the instrument(s) used to promote the desired activity effective; was implementation undermined by rent seeking or other political economy problems (which would speak to the size of  $\varphi$  and  $\lambda$  – government capabilities, and the relative importance of political motives, respectively)?

---

<sup>13</sup> In the Criscuolo et al. (2019) paper, for example, causal inference is premised on the exogeneity of the eligibility criteria for British regions (the criteria having been determined in Brussels). One set of critics might argue that those eligibility criteria are precisely what political logic suggests will be manipulated in other instances. Another might argue that the eligibility criteria will be set endogenously to target deserving regions/firms, producing better results than what the paper finds.

Most, but not all, new empirical work is focused on trying to evaluate the relevance of the economic mechanism as opposed to questions of implementation. However, this does mean that this research does not need to contend with the tension between the search for sources of exogenous variation and identifying outcomes of real-world interest. Continuing with the example of infant industries, evaluating the effect of a truly randomized import tariff on the *average* industry, or a *random* set of industries would not be a particularly useful exercise. We think  $\theta > 0$  applies only in specific settings: technological follower countries in industries subject to the types of market failures discussed in section 2. This is one reason why much new work examines very narrow settings. For example, Juhász (2018) evaluates the effect of trade protection not on general industrial development, but on one specific industry (mechanized cotton spinning) where historical evidence points to the possibility of an infant industry.

Another distinct dimension along which much of the new work has made substantial progress is the evaluation of episodes over the appropriate time horizon. Critiques of earlier studies have long argued that dynamic justifications of industrial policy required, at minimum, proper, long-run comparisons between treated and non-treated units (Lucas 1984, Harrison 1994). A long-time horizon is a feature of many recent studies and also one reason why much of the new work uses historical contexts where effects can be evaluated across decades. Indeed, many of the papers we review below have found a strikingly long time window over which treatment effects are increasing (e.g., Garin & Rothbaum 2022).

We now turn to discussing how credible empirical research designs have tackled these empirical issues and how these findings have informed the debate surrounding industrial policy. The reduced form methods we discuss here are well-suited to evaluating the question of whether industrial policy can elicit the desired behavioral response. That is, do the targeted sectors, activities, or places (henceforth, units) increase outcomes along the margins the policymaker wants? This is an important place to start, as so much of the controversy around industrial policy is about whether the policymaker can shape economic activity in this way. At the same time, it is important to bear in mind that most of the outcomes we discuss here are a necessary but in themselves insufficient building block of a full efficiency evaluation, which requires a model.

#### **4.5. New evidence on sectoral industrial policies**

An early empirical literature generally dismissed the effects of sectoral industrial policy based on a thin and at times only tangentially relevant literature. Much of this early empirical work suffered not just from the endogeneity problems we laid out above, but also from the observability problem we discussed in section 3 – a particularly large challenge when evaluating sectoral industrial policies (Rodriguez & Rodrik 2001, Harrison & Rodriguez-Clare 2010).

Specifically, much of the empirical evidence often evoked to argue against infant industry promotion relies on episodes where it is not at all obvious if policy (typically an import tariff) is being used for infant industry rationales (Harrison & Rodriguez-Clare 2010). This “observability problem” is relevant for cross-country tariff and growth, and trade and growth regressions that reached mixed findings (Lane 2020), as well as for cross-industry studies that typically estimate large positive effects of trade liberalization on productivity growth (c.f. Shu & Steinwender 2019).

The recent literature has advanced our understanding of the possibilities that sectoral industrial policy holds. Where these papers have made the most headway is in tackling challenges of identification and

observability. With one exception, they study contexts that are suited to isolating the economic mechanism but have little to say about implementation. By and large, the small number of papers in this emerging literature have found that in episodes that look a lot like textbook cases of infant industry (textiles, shipbuilding, and heavy industry promotion) industrial policy treatment elicits behavioral responses consistent with theory, though findings on the size of these responses is more mixed (Juhász 2018, Hanlon 2020, Lane 2022).

Interestingly, Juhász (2018) and Hanlon (2020) examine episodes in which the industry needed to adopt a new technology to remain competitive. As such, they speak to theoretical concerns that sectoral industrial policies may simply boost traditional, domestic production methods and in this way undermine the objective of the policymaker (Saure 2007).

Each paper tackles thorny issues of identification by leveraging the external and more or less exogenous (to the structure of the economy) reasons for deploying the “industrial policy.” In Juhász (2018), there is no industrial policy lever at all. Instead, the paper relies on the regionally varying level of natural protection the Napoleonic blockade (1806-1813) against Britain afforded French cotton spinners. As such, the policy mimics the canonical policy of infant industry promotion: a temporary import tariff. In Hanlon (2020), North American shipbuilders were differentially protected from British competition by natural barriers (the inaccessibility of the Great Lakes) and protective policies (that varied across the Atlantic U.S. and Canada – the latter being part of the British Empire and lacking access to independent policy levers such as tariffs).

In Lane (2022) the heavy chemicals and industry (HCI) push materialized because a changing military security environment convinced South-Korean leadership that they needed to strive for domestic capabilities in arms construction. Importantly, up to its implementation, no one was willing to fund HCI projects in Korea, including the World Bank, suggesting few people believed Korea could become competitive in this sector.

Each of these studies finds some support for behavioral responses consistent with the infant industry hypothesis. In Juhász (2018), French regions which became better protected from trade increased capacity in mechanized cotton spinning during the blockade, and the economic geography of the industry persisted, even after the blockade ended. This suggests that, consistent with the predictions of the infant industry mechanism, temporary protection changed the long-term profitability of production across different regions *within* France. Overall, France switched from being a net-importer of cotton textiles to a net-exporter, though these effects are not well-identified.

Hanlon (2020) finds more mixed results. On the positive side, the study finds that better protected areas of North American shipbuilding transitioned from wood to metal ships once “fundamentals” (in particular, metal input prices) equalized between North America and Britain – the global leader in metal shipbuilding. This suggests that protection played a role in technology adoption. However, there is little to suggest that protected North American producers became competitive in global markets, making it unlikely that American producers had a latent comparative advantage in metal ships.

Across a number of different specifications, Lane (2022) finds that targeted sectors increased their output, productivity and comparative advantage over time, while downstream sectors also benefited through network effects. It provides the clearest example of an economy drastically shifting its comparative advantage with the use of industrial policy tools.

Each study raises the question of what exactly the underlying mechanisms at work are. Put differently, what is the market failure? All studies appeal to external dynamic learning-by-doing economies and many provide extensive historical evidence consistent with this. Hanlon (2020) is able to pinpoint the development of pools of skilled local workers as one source of external economies of scale. In follow-up work, Juhász et al. (2023) produce empirical support for another channel in the context of the French mechanized cotton spinning industry: costly experimentation about complementary organizational practices. Using detailed plant level data on productivity and organizational practices such as the layout of the factory floor, the study finds that at initial stages of technology adoption, many plants were operating the new technology with extremely low efficiency, and a wide array of organizational practices. Over time the industry converged on best practice factory layout. This is consistent with Giorcelli & Li (2023) who find using plant level data that tacit knowledge was an important component of technology transfer from the Soviets to the Chinese in the mid-20th century.

The findings of Juhász et al. (2023) highlight a further important point when evaluating sectoral industrial policy. In particular, their results suggest that, unlike in most simple models, the external economies of scale used to justify infant industry promotion can be costly to attain, with gains and losses very unevenly distributed across firms. In that study, many first-generation cotton mills got certain aspects of experimentation such as the layout of the factory floor wrong. These early entrants paid the cost of industry-wide experimentation, while later entrants reaped all the benefits (as in the cost-discovery model of Hausmann & Rodrik 2003). This is an important point to keep in mind as the literature moves towards evaluating industrial policies at the micro (firm or plant) level. Depending on the context, evaluating the effect of sectoral industrial policy for individual treated units may make little sense if the costs and benefits of developing the industry are distributed heterogeneously. In these cases, it may look like the policymaker is targeting producers inefficiently, as some targeted firms do not produce with high productivity. This may miss the fact that through their experimentation or mistakes, these low productivity producers actually generate knowledge (positive externalities) for the entire industry.

These papers suggest that market failures used to justify infant industry promotion do seem empirically relevant in the real-life contexts in which infant industry promotion is often implemented. However, alongside the benefits that leveraging natural experiments entail, it is also important to highlight their limitations. In our view, by far their greatest limitation is that because of their highly synthetic nature, they are too abstract to inform policymaking in any realistic way. The only exception to this is the HCI drive in South-Korea where a vast array of modern industrial policy levers were deployed: preferential tax policy and financing for the targeted sectors, as well as trade policy that reduced input tariffs for targeted sectors.

While not losing sight of careful identification, we believe the literature needs to move towards questions of implementation. It seems clear that infant industry is relevant and can work in the sense of eliciting the desired behavioral responses in some cases. However, we have very little understanding about the second layer of treatment, implementation. The HCI drive in South-Korea, for example, was almost certainly exceptionally effective. In a cross-country study of steel industrial policies, Blonigen (2016) finds that the downstream sector's export performance worsens in developing countries, which is the opposite of the result in Lane (2022). We urgently need more work informing questions of what makes for effective sectoral industrial policies.

A second question the literature needs to grapple with is how to evaluate the efficiency of industrial policy beyond the scope of what reduced-form methods can speak to. While it is clear that such an evaluation needs to be conducted through the lens of a model, it is less clear what aspects of reality the model should focus on. Should we consider the linkage/network effects of a policy as Lane (2022) highlights? For a strategic industry like shipbuilding, how should we account for the national security benefits? If the U.S. had embarked on its historic shipbuilding effort during WW2 (decades after the study period) without relevant expertise in metal shipbuilding, could they have achieved such astounding levels of production and productivity? How do we think about evaluating industrial policy at the plant or firm level when external economies of scale may be costly to develop and heterogeneously distributed as in Juhász et al. (2023)? At the end of the day, a researcher will obviously need to trade off parsimony for realism, but the considerations of how to model the cost and benefits of an industrial policy seem very context specific, and careful, rich, empirical work should inform modeling choices in each case.

#### **4.6. New evidence on public R&D policies**

Much like the theory of infant industry, the idea that the market undersupplies innovation has a long tradition in the field. However, despite the fact that, much like sectoral industrial policy, public-led R&D initiatives are also subject to high-profile failures (Lerner 2009), there is little debate in the field about whether the government should pursue some form of innovation policy. As we mentioned above, the question in this literature is not, “Can innovation policy work?” but rather “What is the most effective form of innovation policy?” In our taxonomy, this means that studies focus on evaluating different forms of implementation, with a view to informing *how* policy should be implemented (Bloom et al. 2019).

These “micro-innovation” policies contrast to more ambitious and less well-understood public R&D policies that focus on particular technologies or places. While long out of fashion among most economists, these “moonshot” or “mission-oriented” policies have made a recent come-back. The productivity slowdown among advanced economies, increased regional inequalities in these same countries, and the climate crisis have led some to argue that the government needs to embrace innovation policies that do more than simply alter incentives on the margin (Mazzucato 2013, 2021; Gruber & Johnson 2019).

These large-scale public R&D policies are different from the “micro-innovation” policies discussed in Bloom, Van Reenen & Williams (2019), both in the scale of spending involved and in the extent of government activism they entail in choosing which activities to promote. Such programs are also more ambitious in targeting transformational outcomes. Economists have long been skeptical of these policies for many of the same reasons that they are wary of sectoral policy: critics worry about the government’s ability to pick projects and execute them efficiently as well as the effects that public R&D will have on crowding out private innovation.

Despite the skepticism, two recent, highly timely papers (Kantor & Whalley 2023, Gross & Sampat 2023) show that this type of public R&D may have fairly large effects locally and, more suggestively, also at the aggregate level. Both papers examine canonical episodes of applied public R&D “moonshots:” the U.S. government’s massive R&D effort during WW2 and the U.S. Apollo mission in the 1960s that culminated in the moon landing. In both cases the level of public investment was massive. For example, NASA received 0.7% of GDP in the mid-1960s (Kantor & Whalley 2023), and the government was actively

involved in picking which technologies would receive support (e.g., radar technology, mass-produced penicillin, or the guidance computer for the lunar landing module).

In terms of identification, both paper's research design builds on the fact that the public R&D "windfall" was exogenous to previous aggregate technology trends. Which technologies were prioritized were driven by factors related to winning the war or the Space Race. Kantor & Whalley (2022) deal nicely with the empirical challenge that NASA may have simply been choosing technologies ripe for harvesting (such as microchips) by using only space technologies as defined by the Soviets. These are more likely to isolate technologies related to mission requirements irrespective of U.S. technological superiority. The papers then compare local outcomes between places more and less treated by these public R&D windfalls within the U.S. using difference-in-differences type specifications.

Each paper finds sizeable and long-lived local effects of large-scale public R&D. Gross & Sampat (2023) find that government funded invention activity during WW2 shaped the geographic pattern of innovation in the U.S. for decades after the postwar period. The shock created broad-based growth that drew in entrants from different geographic areas and different technology categories. Suggestive of some form of aggregate effects, the authors also show that the long-term direction of U.S. innovation shifted towards electronics and communications.

Kantor & Whalley (2023) examine the local (county-industry) effects of NASA spending and find positive effects on their outcome of interest: local manufacturing value-added, both during and after the Apollo mission. The contemporaneous local fiscal multiplier is 1.6 both during the Space Race and after. These estimates are within range of the multiplier effects of other forms of government spending, which are around 1.8 (Chodorow-Reich 2019).

These papers raise the question of what drives the positive effects, and particularly the large positive effects in Gross & Sampat (2023). Partly, this could be due to displacement (as opposed to growth) effects,<sup>14</sup> though Kantor & Whalley (2023) account for many forms of possible spillovers and continue to find large effects. It may be that public R&D creates technological spillovers to the private sector. This is in line with other recent work that has found large and broad-based technology spillovers from (typically smaller scale) public R&D to the private sector (Azoulay et al. 2019, Moretti et al. 2019, Myers & Lanahan 2022).

However, another possible reason is that places were not targeted at random, but rather based on their perceived ability to succeed. Indeed, Kantor & Whalley (2023) show that NASA targeted spending towards locations that already specialized in pre-Sputnik space technology. Gross & Sampat (2023) find that the entire long-run effect is driven by 5% of clusters that were most innovative in 1930 (prior to receiving the government contracts). In this regard, it is interesting to contrast these findings to Schweiger et al. (2022) who study a similar R&D policy in Soviet Russia, but one in which a subset of the treated locations were built from scratch in sparsely populated areas. Similar to the previous papers, this study also finds effects on innovation, productivity, skill composition and wages, often lasting to the present day.

While more work is certainly needed in this area, the findings from this literature suggest that public R&D may have a place in the policymakers' toolkit. First, careful work trying to understand the aggregate

---

<sup>14</sup> See Redding & Turner (2015) for a discussion of displacement and growth effects in spatial settings.

effects of public R&D would be urgently needed. A lingering question with much of the current work is their limited ability to speak to economy-wide counterfactuals. While much of the evidence suggests that the economy would have looked different without these moonshots, it would be valuable to understand how.

Second, it is not clear if public R&D policies would have similar effects in normal times. Yet another explanation for the positive effects of the Apollo mission is that the Space Race engendered superhuman efforts from those involved (Mazzucato 2021) – an aspect that would be difficult to sustain consistently. Similarly, in the U.S. context, both papers targeted places that were likely to succeed, and even so, Gross & Sampat (2023) find long-lasting effects only for the most productive of these highly selected initial locations. This is a sobering insight for initiatives that envisage similar public R&D investment in less auspicious locations today (Gruber & Johnson 2019).

These open questions aside, it seems that careful, systematic work is casting doubt on many long-held assertions that have not previously been put to the empirical test. Rather than crowding out private R&D, a number of papers suggest the opposite: the potential for public R&D to crowd in private innovation. Similarly, in times of national crisis, the U.S. government seemed capable of picking technologies, places, and firms that could deliver the desired outcomes, often with long-lasting positive local effects. Moreover, in the Soviet/Russian context, public R&D achieved long-lasting local change even in areas that were built from scratch.

#### **4.7. New evidence on place-based industrial policy**

Increasing regional inequality and the persistence of economic distress in space have led to a resurgence of interest in place-based industrial policies. Broadly speaking, place-based policies target particular areas with the goal of improving economic conditions in that area along some margins. The justification for these policies relies either on localized market failures such as agglomeration economies or on equity considerations (Kline & Moretti 2014). There are two types of concerns with these policies. One is that they may prove ineffective if mobile workers and firms arbitrage away the benefits by relocating across targeted and untargeted areas. There may also be a trade-off between efficiency and equity if correcting an externality such as a knowledge spillover requires the targeting of places that are already better-off (Glaeser & Gottlieb 2008).

The literature on place-based policy in general is well-surveyed (Kline & Moretti 2014, Neumark & Simpson 2015, Duranton & Venables 2019) and beyond the scope of our discussion. Here, we focus on place-based industrial policies that are targeted at places, but in a selective way. Typically, this means trying to boost manufacturing activity in lagging regions or, similarly, to alleviate the decline of manufacturing in a distressed region. A number of recent papers use credible identification strategies to evaluate the effects of these place-based industrial policies.

Mitrunen (2021) and Garin & Rothbaum (2022) both use historical natural experiments to examine the effects of public industrial investment on local outcomes. The research design in Garin & Rothbaum (2022) is particularly intriguing. The authors exploit the extraordinary industrial mobilization undertaken in the U.S. during WW2, a period in which manufacturing output increased by an astonishing 300% between 1939-1942. In particular, they use the fact that where private firms could not be subsidized to undertake investment, the U.S. government built plants from scratch in places outside of established



manufacturing hubs. This presents a powerful identification strategy, as local outcomes in treated places can be compared to plausible untreated counterfactuals that had similar location fundamentals.

In contrast, Mitrunen (2021) studies the effects of Finnish war reparations paid in kind to the Soviet Union in the 20th century. The strength of this research design is that what Finland had to produce was dictated by the needs of the Soviet Union, forcing Finland to embark on a drastic switch into complex metalworking products (ships, locomotives, cables, and engines) – skill-intensive industries in which the country had little prior experience. The author then uses a shift-share research design to study the local effects of forced production on more affected Finnish municipalities.

Similar to the local effects of public R&D, both papers find sizable effects on long-run measures of industrial development. Strikingly, Garin & Rothbaum (2022) find that the local effect on manufacturing employment is not only large, but also highly persistent. Against a backdrop of declining manufacturing employment across the U.S., places treated by a publicly funded government plant during WW2 had about 20% higher manufacturing employment through the 20<sup>th</sup> century and into the 21<sup>st</sup>. In line with the findings from the public R&D literature, the long-lasting effects were driven by broad-based growth in manufacturing activity as private manufacturing establishments were drawn into the area after the war. Likewise, Mitrunen (2021) finds that treated municipalities experienced broad-based structural transformation. Not only did treated municipalities move out of agriculture and into industry at a faster rate, but the manufacturing base expanded to industries upstream and downstream of the initially “treated” sectors.

The historical nature of these episodes and the availability of individual-level earnings and income data allow the researchers to examine the intergenerational effects of treatment. Both find that children living in treated areas before treatment experienced upward economic mobility later in life. Garin & Rothbaum (2022) find that this is likely driven by the fact that children living in treated areas had access to high-paying manufacturing jobs in their own community later in life. In contrast, Mitrunen (2021) finds that the effects were primarily driven by increased educational attainment.

Similar to the other natural experiments we have examined, the power of these papers is their strong claim to credible identification and their ability to evaluate effects in the long run at the individual level. However, in recent years, this literature has also seen the emergence of papers that evaluate the second layer of treatment using detailed, firm level data: implementation, alongside a credible identification strategy. In particular, Criscuolo et al. (2019) evaluate a place-based industrial policy program in the UK (Regional Selective Assistance) whose goal was to create and safeguard employment in the manufacturing sector. Their research design builds on changing eligibility rules which led to quasi random variation in which distressed regions in the UK received more or less support in the form of an investment subsidy. The authors find that the policy was effective in increasing jobs and reducing unemployment, but it had no effect on TFP, and the effect operates solely through small firms. There is no evidence to suggest that job increases are due to displacement from neighboring ineligible areas.

Similarly, Cingano et al. (2022) also evaluate a public investment subsidy program in Italy designed to “stimulate fixed investment in underdeveloped areas of the country” (p. L488). Under this scheme, funds were allocated through calls targeting different sectors, primarily in industry. Like in Criscuolo et al. (2019), the main objective of this policy was job creation. The authors’ research design builds on the fact that projects were given numerical scores and funded in rank order until the funds were allocated. This institutional detail naturally lends itself to a regression discontinuity design. Consistent with the previous

findings, the authors find that the policy induced the desired behavioral response in terms of job creation: marginal firms near the cutoff increased investment by 39% and employment by 17% over a 6-year period. There is also no evidence of spillover effects on firms in the same local labor market.

It is interesting to contrast these findings to La Point & Sakabe (2021), who study a place-based industrial policy with slightly different objectives and a different policy lever. In particular, they examine Japan's Technopolis program in the 1980s that had the goal of promoting high tech manufacturing outside of the main metropolises. The program incentivized investment in these regions through a bonus depreciation, which allowed firms to deduct an additional fraction of physical capital costs in the first year of an asset's tax life. That is, the program lowered the effective price of capital in targeted regions relative to untargeted ones. Using a staggered difference-in-differences approach, the authors find that the policy was successful at generating both investment and employment growth in treated areas. However, the authors also find that for multiplant firms with a plant in a treated area, the employment response was over 6 times larger for plants in untreated regions relative to the treated plants suggesting some unintended consequences of the policy through leakages.

The findings from these recent papers highlight a few important points. Most significantly, there is evidence that place-based industrial policies can shift the composition of local economic outcomes both in lagging and declining areas consistent with the desires of the policymaker. Importantly, as the goal of these policies is to improve outcomes for those living in these areas, a particularly striking finding is the increase in upward mobility for children that grew up in treated areas. While a full general equilibrium analysis of the effects is outside of the scope of these papers, many show that there is little evidence of displacement effects in neighboring locations. In contrast however, some papers do show that one margin through which leakages occur is through multiplant firms. This suggests that looking for (positive or negative) spillovers in nearby places may be insufficient when testing for displacement effects.

#### **4.8. Taking stock**

There are strong theoretical and economic justifications for industrial policy, as well as many well-justified practical concerns about its use. Given the terms of the debate surrounding industrial policy, recent empirical studies have focused mostly on evaluating the empirical relevance of the theoretical justifications for its use. This is what we termed the economic layer of treatment. A striking finding that emerges across some (though not all) of these papers is just how large and long lasting the local effects of industrial policy can be. To us, a balanced reading of the emerging literature suggests that it is no longer possible to dismiss industrial policy as ineffective or counter-productive. In stylized environments where industrial policy comes about "by accident," we have seen the potential for long-lasting, transformational local effects. Put differently, the market failures that justify its use seem large. This should fill us with concern as the current literature is mostly far too removed from the real world to inform policymaking in any serious way.

One consequence of the recent stream of papers might be that this policy domain will be normalized, with research and discussion moving from the whether to the how. Reduced form methods will play a crucial but in themselves insufficient part in this new research agenda. Careful structural work, informed by a detailed understanding of the nuances of particular contexts, has an important complementary role to play.

## 5. Industrial policy and the East Asian miracle: new empirics

The East Asian miracle is one of the most important episodes of modern economic development and the focal point of debates surrounding industrial policy. The miracle economies, such as Hong Kong, Singapore, South Korea, Taiwan, and, earlier, Japan, experienced rapid structural transformation and joined the club of high-income economies. Their transformation was not only rapid, but improbable. Since WW2, few countries achieved the rapid economic growth seen across East Asian economies, especially outside resource booms (Krueger 1986, Wade 2019).<sup>15</sup>

Given the ubiquity of industrial policy across postwar Asia, what role did industrial policy play? Within economics, the episode remains controversial, and one with remarkably little empirical research that meets the standards of contemporary economics. New research, however, is changing this picture. This section reviews how contemporary empirical work is updating our view of industrial policy, considering the empirical issues highlighted in Section 4.

Competing schools of thought draw different lessons from East Asia. On the one hand, social science scholars see industrial policies as integral to the experience. There is a voluminous comparative literature on the role of industrial policy in East Asia, largely sympathetic to the role of these policies in the postwar period.<sup>16</sup> While this literature largely emanates from sociology and political science, some economists have been sympathetic to this view (see: Rodrik 1995, Stiglitz 1996, Stiglitz and Uy 1996, Lin 2012). Many mainstream economists, on the other hand, have been historically skeptical of the role of industrial policy, and this skepticism has produced influential challenges to the role of the state in industrial development (e.g., see Lal 1983, Ito & Krueger 1995, Lall 1996).

The new empirical explorations of the East Asian miracle paint a more varied landscape. This emerging picture is more precise, granular, and pragmatic than first-generation empirical work. Early regression studies largely considered the correlation between measures of industrial policy and industry-level performance, such as productivity, or externalities associated with justifications of industrial policy use (e.g., static scale economies).<sup>17</sup> For a discussion of early evidence, see Lane (2020). However, these early correlational studies are mired by the empirical issues reviewed above (section 4). In addressing these issues, new empirical work provides a richer analysis of a diverse episode—with potential insights for policy practice.

East Asia is a heterogeneous region, from city-polities, such as Singapore, to countries with robust domestic markets, like Japan. Strategies varied too, from the focus on macro-stability and entrepôt trade of Hong Kong to the dirigisme of South Korea under General Park Chung-hee. Even among the developmental states, industrial policies were varied (Vittas & Cho 1996). South Korea and Taiwan are often mentioned in the same breath, yet the favored policy instrument was credit subsidies in the first case and tax incentives in the second. Given such diversity in experiences, one quickly runs into difficulty when trying to generalize, unconditionally, the policy experience of the region. Recent empirical work

---

<sup>15</sup> Between 1960 to 2014, only 16 countries achieved high-income status (Cherif & Hasanov 2019).

<sup>16</sup> The qualitative literature in social science is immense, see notable contributions by Johnson (1982), Wade (1989), Haggard (1990), Amsden (1992), Chang (1993), Evans (1995), Kohli (2004), and many others. Also see the more contemporary popular treatment by Studwell (2013).

<sup>17</sup> See, for example, Beason & Weinstein (1996), Lee (1996), and Pack (2000).

has attempted to make sense of this experience by using higher-resolution data and making use of deep institutional details.

### **5.1 New empirical research on East Asian industrial policy**

Consider the case of South Korea's Heavy and Chemical Industry (HCI) drive, a landmark – and controversial – industrial policy pursued by President Park Chung-hee's military autocracy. New research has updated earlier, pessimistic perceptions of this push (Lee 1996, Yoo 1997). Lane (2022) was the first to use the HCI episode as a natural experiment to consider the impact of the bundle of HCI policies on South Korean industrial development. Using variation from the episode, he estimated the differences in the treated vs. non-treated industries and found that these policies increased short- and long-run growth in treated HCI industries – effects which seemed to persist. Likewise, Lane finds some positive, though weak, effects of HCI on plant total factor productivity after the policy period. Lane also finds cross-country support for the role of industrial policy in shifting the comparative advantage of HCI industry, an effect that took time to occur.

Lane (2022) has inspired follow-up work using the HCI natural experiment. Lee et al. (2022) use the HCI experiment with administrative data, finding similar patterns for plant growth and output but making a plausible argument that HCI contributed to misallocation. Lee et al. also find that HCI policy may have changed the network structure of the economy as well, making HCI sectors more central. Like Lee et al. (2022), Choi & Levchenko (2022) use spatial variation and microdata but take a slightly different approach, emphasizing the impact of foreign investment. They study the HCI episode by combining Lane's natural experiment with firm-level corporate finance data and a quantitative framework. Choi & Levchenko (2022) also find a positive impact of the HCI policy on industrial development and argue the HCI policy led to considerable long-run welfare gains. These studies all update the perceptions of the controversial episode, highlighting the importance of understanding the allocative and aggregate impacts of the drive.

The early economic literature on East Asia argued that industrial policies failed to target the appropriate sectors, as governments did not possess the information required to target market failures (Pack & Saggi 2006). Yet, recent work by Liu (2019) formally explores this issue, using a model network economy populated by imperfections. In such a setting, where would a planner without complete knowledge optimally target industrial policy interventions? What characteristics would these sectors have? Liu's analysis shows that optimal sectors to target are those where imperfections are most consequential given their impact through input-output linkages. Sectors may directly and indirectly impact the broader economy where market imperfections compound through linkages. Even if the sector targeted by policy is subject to weak or no externalities itself, if it is an upstream sector that produces inputs for many others downstream with market imperfections, the policy can be welfare-enhancing.

In his analysis, Liu provides a theoretically motivated sufficient statistic, "distortion centrality," which conveys the extent of misallocation associated with each sector and provides a potential guide for policy-making. Under certain conditions, this measure correlates to quantities calculated from an input-output table. Importantly, he shows that the policies used in China and South Korea's HCI drive (from Lane 2022) correspond to sectors with high distortion centrality.<sup>18</sup> Thus, Liu demonstrates that, far from having to

---

<sup>18</sup> These economies have "hierarchical" production networks, which make this calculation feasible.

possess perfect knowledge of market failures, distortion centrality corresponds to observable features of the input-output network. Moreover, those sectors with high social returns to policy may not necessarily be the largest or those with the features used by earlier studies. Similar insights may help consider key optimal targeting for R&D spillovers. Related work on R&D networks by Liu & Ma (2023) provides a measure of “innovation centrality” by considering the spillovers associated with innovation. Interestingly, Liu and Ma find Japan has the most efficient allocation of R&D, superior than in the U.S. Similar work, like König et al. 2019, has considered how subtle network features and competitive behavior could inform optimal R&D policy.

Such research demonstrates that we should be cautious of early empirical tests of how well policies aligned with market failures. In the work above, optimal targeting in networks may look a lot different than the criteria chosen by the econometrician (for example, the largest sectors or those with the highest number of links). These issues are closely related to early tests of optimal targeting by considering the correlation between policy and static economies of scale above. As well, many classic theories of optimal industrial policy also consider targeting *dynamic* economies of scale (Corden 1997).<sup>19</sup>

There are also important reasons why early studies uncovered a negative relationship between East Asian industrial policy and sectoral performance. For example, institutional features of the Japanese industrial policy system explain this negative relationship. Scholars have long noted that MITI and Japanese policies dedicated large resources to cushion and restructure languishing industry (for example, in response to oil shocks; Saxonhouse 1979). This was a well-documented dimension of Japanese policymaking (Uriu 1984, 1996); by some measures, declining (sunset) sectors received more resources than infant (sunrise) sectors (Teranishi 1986). This source of bias parallels our previous discussion in section 4.2. Using institutional narratives and updated measures of TFP, Diewert et al. (2011) revisits the correlation between Japanese industrial policy and productivity. Accordingly, the authors find different performance between the sunset and other sectors covered by earlier work (see Beason & Weinstein 1996).

In other words, the institutional context matters in how we evaluate industrial policies and interpret their results. Empirical research will likely have to consider the objectives of policies and deploy higher resolution data on policy levers. Institutional details not only help to interpret earlier episodes, but also help parse and contextualize recent ones, like China and Vietnam.

For instance, FDI is an essential feature of both Chinese and Vietnamese industrial policy, and one that is distinct from East Asian predecessors (Ye 2009, Huang 2000). While postwar Korea and, in particular, Japan carefully controlled foreign-directed credit, the specifics of the Chinese political economy and the forces of globalization made FDI an integral tool (Thun 2004). On one hand, the lingering prominence of state-owned enterprise in China and Vietnam may make it difficult to directly map the lesson from SOE-centered policies to other economies. The incentives issues, social objectives, and environment surrounding such policies may be quite different than most capitalist economies. On the other hand, the region’s use of FDI-as-industrial policy may offer more generalizable insights and policy lessons for developing economies (see: Harrison & Rodriguez-Clare 2010).

---

<sup>19</sup> Work by Pons-Benaiges (2017) considers evidence that dynamic learning economies in Japanese sectors may be different than sectors with static economies of scale during the miracle period.

The Chinese automobile sector embodies this centrality of FDI as an industrial policy (Thun 2006). Using rich data and a quantitative empirical framework, new work by Bai et al. (2022) explores the impact of Chinese quid-pro-quo style FDI, from 2001 to 2014. This controversial aspect of Chinese FDI requires foreign producers to enter joint ventures with Chinese automakers to produce and sell cars domestically. Spillovers from FDI to domestic markets are a critical rationale for these policies and have been a potent aspect of FDI for industrial development (Haddad & Harrison 1993, Javorcik 2004, Blalock & Gertler 2008). Using a comprehensive dataset on the Chinese automobile industry, Bai et al. study spillovers from foreign joint ventures to domestic affiliates, focusing on quality upgrading. Their results point to knowledge spillovers from foreign joint ventures to local affiliates, and these effects take time to manifest – evidence of learning. Moreover, using linkage data, they show that these spillovers are driven by worker flows (e.g., from joint ventures to affiliated domestic firms) and are embodied in high-quality, shared components. Relative to unrestricted FDI, they estimate that the quid-pro-quo FDI improved the quality of affiliated domestic models between 3.8-12.7% and raised their sales by 0.9-3.9% between 2007 and 2014.

Nevertheless, the literature on Chinese industrial policy is still scant relative to their scale and notoriety.<sup>20</sup> Aghion et al. (2015) use firm-level data to consider the relationship between Chinese industrial policies and the level of competition within industries. They find subsidies and tax holidays promote productivity when directed at more competitive industries; tariffs and loans do not. These results point to important potential lessons in policy design. Likewise, deep work on Chinese shipbuilding by Barwick et al. (2019), building off Kalouptsi (2018), points to the nuances of industrial policy design. Barwick et al. show that not all policy levers were efficacious in promoting the ascent of Chinese buildings. Production subsidies and investment subsidies were more effective, while entry subsidies promoted inefficient firms. The authors show that targeting policies toward productive firms would have likely been more beneficial.

Empirical research on East Asia is updating the policy details to be gleaned from the miracle episode. If the “devil is in the details” when getting industrial policy right, as shown above, many of these details varied across miracle economies and within them. The emerging research, while incomplete, is clarifying where and when policies worked and their mechanisms. Importantly, this work is highlighting the institutional contingencies of success and failure. Moreover, institutional context is necessary for interpreting such lessons. For instance, the role of state-owned enterprises in communist economies may embody a multitude of other political objectives (Malesky & London 2014), unlike the industrial policies of capitalist economies. Nevertheless, the specificity of the East Asian environment and history should not blinker us to the applicability of lessons.

## **5.2. Is this time different? East Asia and contemporary industrial policy**

Is the experience of East Asia relevant for today? The varieties of experience across Asia do yield contemporary lessons. The integral role of trade and the emphasis of export policy imbue current conversations of industrial development. Likewise, the pragmatism of East Asia shows us how countries adapted to the experiences of neighbors and predecessors, and how they negotiated industrial policy with changing global institutions (e.g., multilateral trade institutions). Rather than being idiosyncratic and exotic, the region’s experience calls for a broadly strategic and dynamic approach to the practice of

---

<sup>20</sup> See estimates of Chinese industrial policy for 2019 by DiPippo et al. 2022.

industrial policy, and, as we show in the next section, informs us about the design of future industrial policy.

In important ways, East Asian policies were a hint of development policy to come. Along with scholars of South Korea, Lane (2022) notes that throughout South Korea's miracle, even during its most interventionist period, formal measures of effective protection were decreasing. Likewise, Japan liberalized notably in the 1960s as they joined the GATT and further during their entry into the OECD in the 1970s. Industrial policies never (fully) went away; rather, they evolved with the globalizing economy. The importance of export and trade facilitation among East Asian tigers echoes the preponderance of export-oriented policies shown in Section 3. Juhász et al. (2022) find an important role of trade finance among modern users of industrial policies. Such instruments were ubiquitous across East Asian economies, through bodies such as the Export-Import Bank of Japan. The governance mechanisms behind East Asia's industrial policies were not so different from today's successful examples either, as we will discuss in the next section.

## **6. New thinking on (and context of) industrial policy**

In recent years, industrial policy has also been transformed by the demands of a new economy. Manufacturing still occupies center-stage in many government initiatives (as in creating local supply chains and promoting advanced manufacturing). But digitalization, the green transition, the middle-class squeeze (and the associated good-jobs challenge), and geopolitical imperatives have multiplied the objectives that industrial policy is tasked with. This creates inevitable tensions and trade-offs. For example, requiring domestic content in supply chains to spur local economic development makes green investments more costly. Subsidizing advanced semiconductor manufacturing – which is highly capital- and skill-intensive – may be what is required for geopolitical competition with China, but it is not an effective way of generating good jobs where they are most needed. Multiple goals require multiple instruments – a lesson that many governments have yet to internalize. They also require thinking of industrial policies in ways that are somewhat different from what economists are accustomed to.

### **6.1. Iterative public-private collaboration versus top-down regulation**

The image that economists have of industrial policy goes something like this: a group of bureaucrats (a) design some incentives that favored activities or sectors are to receive (e.g., export subsidies, import protection, cheap credit, etc.), and (b) select the sectors and activities that are to be incentivized in this fashion. They may then formulate some additional rules regarding what kind of firms qualify for the incentives, the specific firm actions or performance criteria on which the incentives are conditioned, and the consequences (or penalties) for non-performance. Ideally, the bureaucrats keep firms at arms' length throughout the process and thereafter, to provide them with insulation against political manipulation and rent-seeking. East Asian governments are supposed to have done an especially good job at disciplining private firms, and they serve as the example to emulate (though economists remain generally skeptical that this is possible outside the context of a few East Asian countries).

But this description of the hard, insulated state does not quite do justice to the reality of East Asian industrial policy. As the sociologist Peter Evans (1995) has argued, successful East Asian governments like South Korea combined autonomy from private interest groups with "embeddedness" in social ties that provided "institutionalized channels for the continual negotiation and re-negotiation of goals and policies." "A state that was only autonomous," Evans wrote "would lack both sources of intelligence and

the ability to rely on decentralized private implementation” (Evans 1995, p. 12). Evans called this system “embedded autonomy” and wrote that South Korea exemplified it. Economists might worry that such close relationships with private firms could have made the government more prone to capture. But Evans argued that these links were essential to ensure that governments had access to the information needed to design workable policies, could adjust to changing circumstances, and prod firms along new technological trajectories in the most effective ways possible. The difference with India and Brazil, the other, less successful cases Evans analyzed, was less in the formal aspects of the policies (the tariffs or subsidies), and more in the manner in which this cooperative relation was managed dynamically over time.

Evans’ discussion highlights that embeddedness can be as important as autonomy to successful industrial policy. Following Wright (1996), this argument can be summarized in the form of a 2x2 matrix where state characteristics can vary along both dimensions, as shown in Figure 6.1. The Weberian ideal of an autonomous, competent state engaged in top-down regulation is in the upper right cell, with high autonomy but low embeddedness. The mirror opposite is the clientelist state, with low autonomy but high embeddedness in the lower left cell. The predatory state has neither autonomy nor embeddedness (upper left cell) while the developmental state combines both attributes (lower right cell).

At first sight, embedded autonomy might appear to be a feature of East Asian states which other countries have been unable to emulate. But there are many examples from other settings that suggest this is not the case.

One of the most successful cases of industrial policy in advanced countries is the ARPA model of the U.S., which operates along lines quite similar to embedded autonomy. (In the ARPA language, the approach is called “active program management.”) The model originates with DARPA, the Defense Advanced Research Projects Agency, which was set up in the late 1950s to counter a perceived lag in missile technology vis-à-vis the Soviet Union. As the name implies, DARPA’s focus was on promoting defense-related frontier technologies, though this mission has been interpreted very broadly. The agency has played a key role in the development of some breakthrough technologies such as the Internet, GPS, and GUI which have transformed the economy. The model has been subsequently emulated at smaller scale in energy (ARPA-E, set up in 2009) and health sciences (ARPA-H, set up in 2022).

The central figure in the ARPA model is the program director, who is not a career government official, but a professional from academia or industry who is recruited for a three-year term. The program director works in a designated area and selects a portfolio of projects to support with grants. The selection process and the associated performance milestones are designed following intense consultations and workshops involving the program director and the relevant stakeholders of researchers and firms. The program director works closely with the grantees as the projects unfold, reviewing and revising targets as needed. Grantees are required to provide quarterly progress reports which the agency staff rates according to a traffic light system: projects that miss a critical milestone and are at risk of failing receive a “red;” projects that miss a milestone but can be expected to recover get a “yellow;” and projects that are on track to reach their goals get a “green.” Projects with red ratings call for more intense supervision and scrutiny on the part of program directors (in the form of site visits, conference calls, meetings, and written analyses of problems and possible solutions). If putting the project back on track seems infeasible, the program director can issue a letter warning of the possibility of termination. These performance criteria are coupled with considerable amount of flexibility. Specific



milestones, like much else in these programs, are provisional: they are added or deleted in nearly half of the projects. Similarly, budgets can be expanded or decreased in light of project developments. In short, the model is based on continuous iterative collaboration with the private sector and is quite different from the hands-off, arms'-length, strict ex-ante-rules regulatory approach favored by economists (Azoulay et al. 2018, Rodrik & Sabel 2022; see also Sabel & Victor 2022 for a general discussion of what the authors call "experimental governance" in the domain of green technologies).

For a different example of industrial policy from another part of the world but with many of the same features, consider the Mesas Ejecutivas of Peru. These were a system of sectoral roundtables established during the mid-2010s by Peru's then economy minister Piero Ghezzi to enhance productivity in selected sectors. The sectors included forestry, aquaculture, creative industries, textiles, logistics, and agricultural exports (Ghezzi 2017). The objective of the roundtables was to institutionalize public-private dialog aimed at addressing coordination failures among firms and between the firms and the government, and thereby encourage productivity-increasing investments. The roundtables started as open-ended conversations to share information on and uncover constraints. Instead of lengthy industry reports, the focus was on developing an initial list of blockages or obstacles to productivity and means of removing them, to be revised as needed as more knowledge was acquired in the process. The responsibilities for action were divided into separate categories of "my problems" and "your problems." The former category refers to government responsibilities (e.g., removing red tape for exports or establishing a national phytosanitary agency); the latter refers to firms' actions (e.g., making specific investments in quality upgrading). When removing identified constraints required action on the part of higher-level government bodies, the matter was bumped up to inter-ministerial or presidential levels.

Not all sectoral roundtables were effective or successful. But where they did work, as in forestry, they did lead to some policy action and private sector buy-in. Ghezzi quotes one business person as saying that "more was done [in the forestry roundtable] than was achieved in the previous 200 years of Peru's republican history" (p. 377). An important feature of the conversations was that the focus was on solving coordination failures and providing public inputs; subsidies and other types of financial incentives were generally off the table. This helped firms concentrate directly on productivity, rather than receiving handouts from the government.

## **6.2. Customized public services and inputs instead versus subsidies**

This focus on specific constraints and productivity-enhancing public inputs instead of subsidies is another important feature of modern industrial policies. In the U.S., this approach is best exemplified in the work of Tim Bartik who has carried out extensive analyses of local economic development policies geared towards job creation (Bartik 2019, 2020, 2022). State and local governments spend large amounts of resources annually in the form of tax credits, often competing with each other, to attract firms, establish industrial clusters, and create employment opportunities (Slattery & Zidar, 2020). While these subsidies generally lead to increased jobs, they do so at large fiscal cost. Bartik has argued that it would be far more effective to provide current and prospective investors with customized business services and inputs rather than subsidies. He estimates that public spending on infrastructure, manufacturing extension, specialized training, and brownfield development create new jobs at a fraction of the cost. For example, customized job training and manufacturing extension services cost \$34,000 per new job, compared to \$196,000 per job for tax incentives (Bartik 2020). Yet total spending on tax and other

financial subsidies are in the range of \$50 billion a year compared to roughly \$1 billion in total for manufacturing extension and customized training.

An additional advantage of customized public inputs is that they directly target shortcomings in the local business environment. Hence they are productivity-enhancing in addition to employment-creating. But this makes it imperative that these inputs be adequately geared to the real needs of existing businesses or likely investors. Otherwise the resources could be wasted, as it so often happens, for example, with generic training or infrastructure investments. This makes the kind of dynamic public-private dialog and information exchange we have discussed a critical component of this type of industrial policy.

### **6.3. Manufacturing versus services**

Industrial policy has traditionally focused on manufacturing, as the name indicates. But the economic rationales for industrial policy – externalities, coordination failures, specialized public inputs – are general and do not apply to just manufacturing industries per se. Moreover, manufacturing’s importance in the economy has generally shrunk in all advanced countries, when computed at current prices. Employment-deindustrialization has been even starker: in the U.S., the share of manufacturing in total employment now stands below 10%. A reduction in the share of manufacturing employment has been a common feature of all advanced economies, even those (such as German, South Korea, and Taiwan) that have maintained globally competitive manufacturing sectors. Moreover, de-industrialization is not something that is restricted to rich countries. As manufacturing technologies have become more capital- and skill-intensive and global competition has intensified, many lower-income countries have found it difficult to jump on the industrialization bandwagon, with “premature de-industrialization” setting in (Rodrik 2016).

One consequence is that governments are likely to look beyond manufacturing as they consider productivity-enhancing “industrial” policies in the future. This is especially the case when the focus is on “good” jobs – i.e., those that act as career ladders into the middle class (Rodrik 2022). It is almost a statistical certainty that the bulk of such jobs will have to be generated in services. So the question becomes whether the productive development policies typically applied to manufacturing can also be appropriate for sectors such as retail, hospitality, education, health care, or long-term care. There is very little experience and evidence on the benefits of sectoral policy in these areas. But “good-jobs” externalities are rampant in such services, and we know that these activities can benefit from complementary investments in new work practices, job-specific training, technologies that complement and empower workers, better tailored regulations, and improved organizational culture. Public-private initiatives that promote such investments can enhance labor productivity, enabling the provision of better jobs (see Rodrik 2022 for further discussion and examples).

In sum, the actual practice of industrial policy looks quite different from the way economists have traditionally conceptualized it. It entails dynamic, iterative collaboration between the government and firms in the pursuit of a more diffuse set of goals. We summarize in Table 6.1 the main differences between these traditional and “modern” conceptions of industrial policy.

## Tables

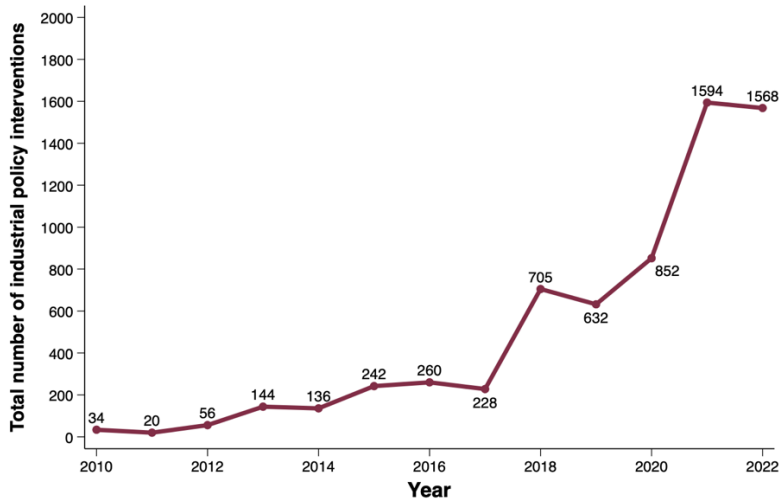
**Table 3.1. Words associated with industrial policies**

<b>Feature Names</b>	<b>Coefficient Size</b>
export	5.5
development	5.3
project	4.5
support	4.4
industry	4.0
million	3.6
energy	3.5
research	3.3
boost	3.3
technology	3.1

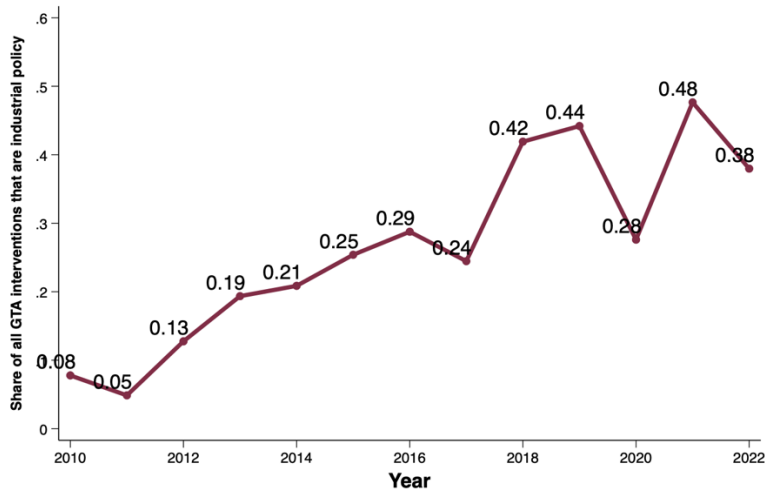
*Notes: This table lists the ten features most predictive of industrial policy from a text-based, binomial logistic regression, and that correspond to individual tokens. The text of these features are reported in the left column. The right column reports the estimated coefficients. Source: Juhász et al. (2022) (data update July, 2023).*

Figures

Figure 3.1. Time trend of industrial policies



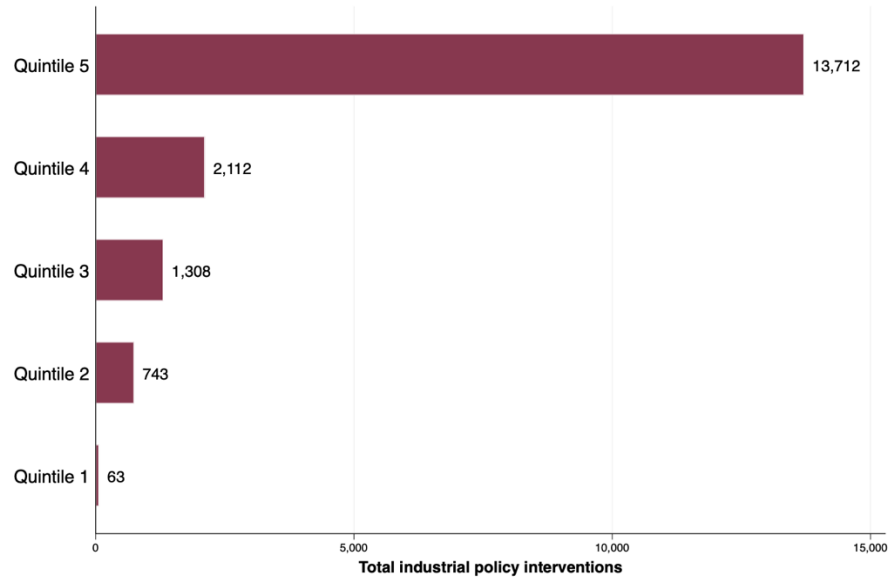
Panel A: Total number of industrial policy interventions



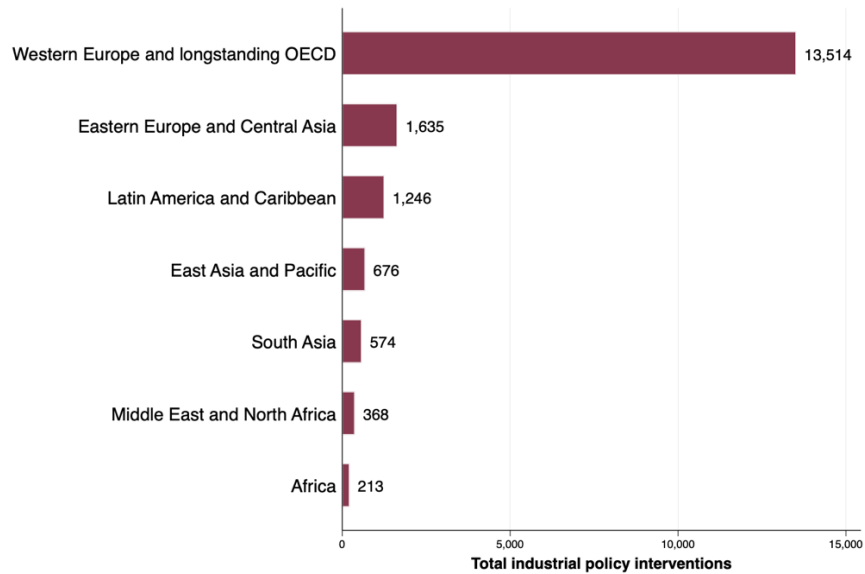
Panel B: Share of all interventions classified as industrial policy

Notes: Panel A plots the total number of industrial policy interventions by year globally; panel B plots the share of industrial policy interventions among all interventions in the GTA. Following guidance from GTA, only policies entered in the same calendar year are included to ensure comparability across time. See the July 2023 data update from Juhász et al. for details. Source: Juhász et al. (2022) (data update July 2023).

**Figure 3.2. Incidence of industrial policies by group, 2010-2022**



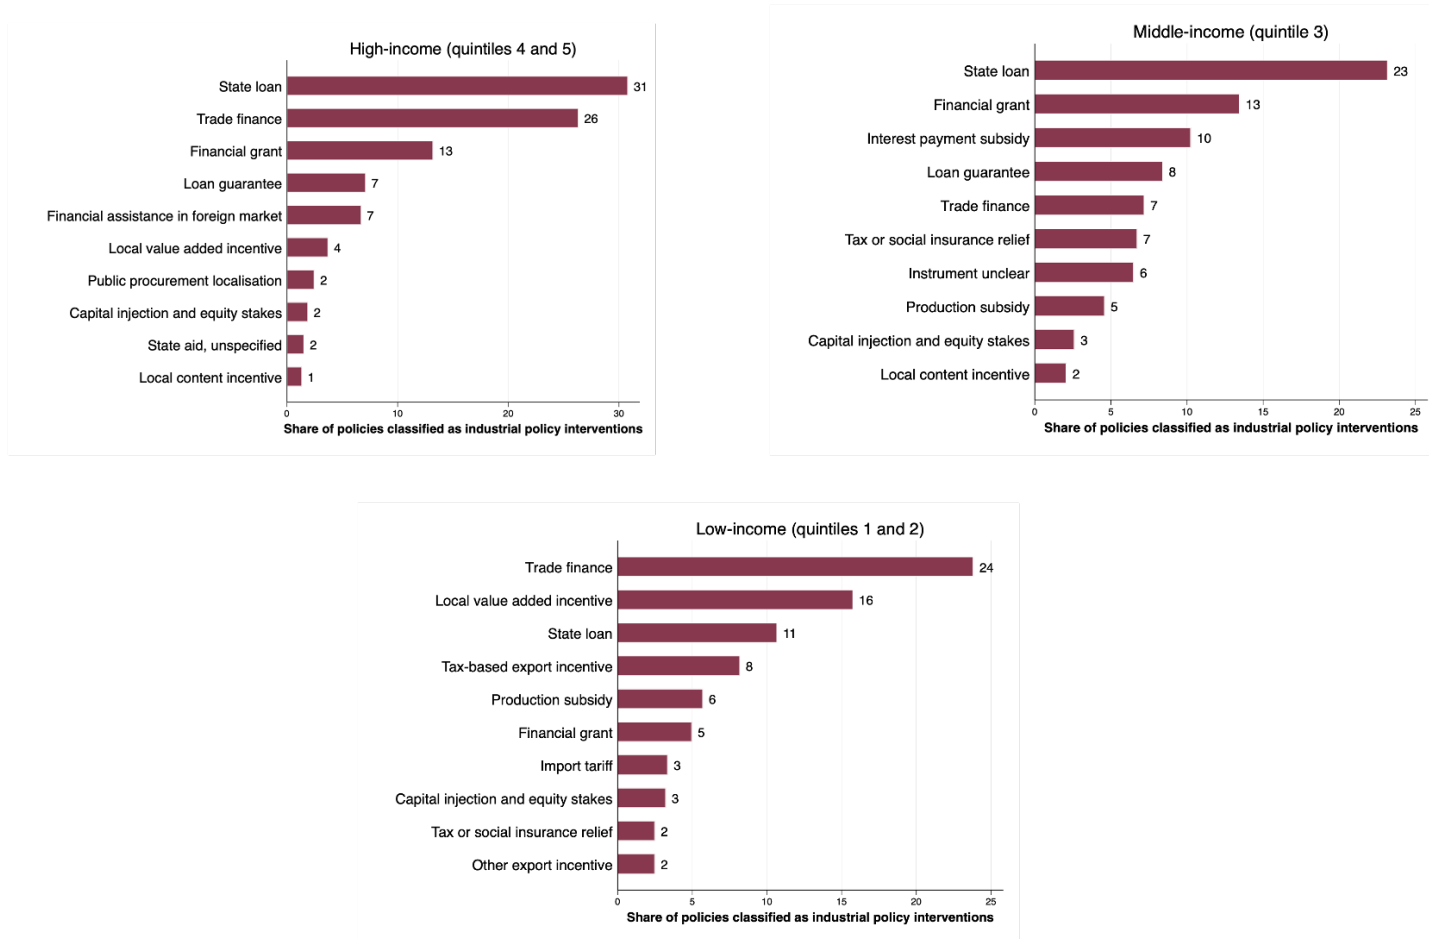
**Panel A: Total industrial policy interventions by income quintile**



**Panel B: Total industrial policy interventions by region**

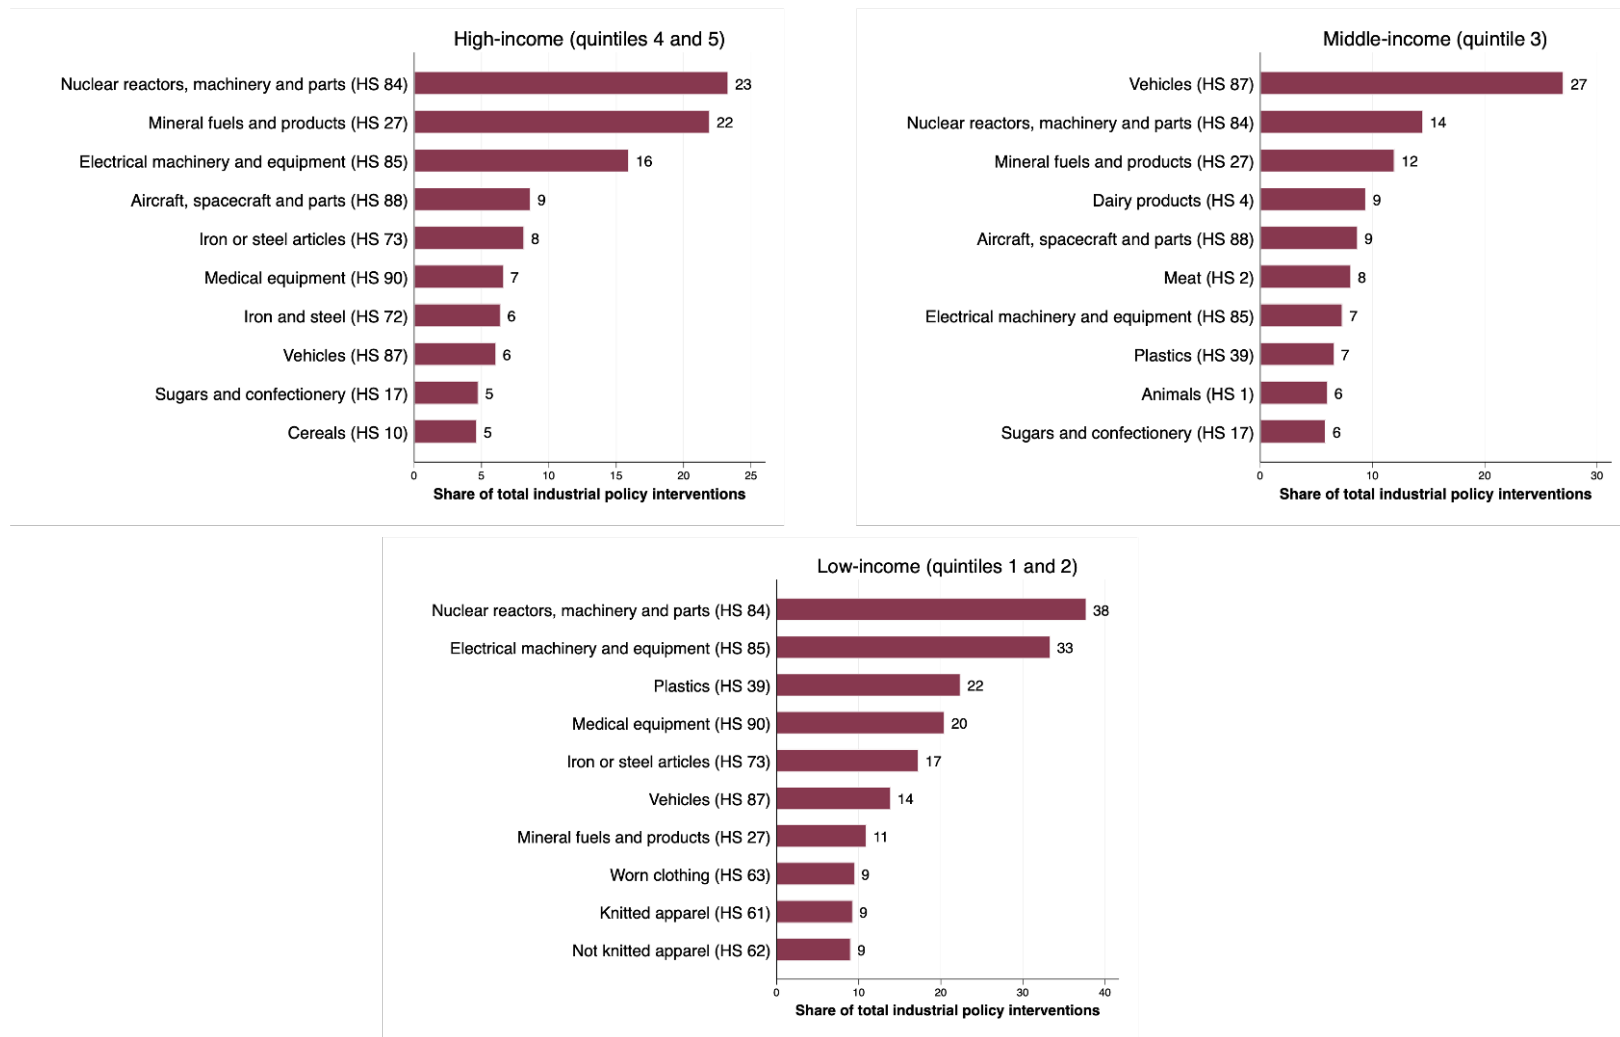
*Notes: Panel A plots the total number of industrial policy interventions (2010-2022) by income quintile based on GDP per capita in 2010. Quintile 5 is the highest income group; quintile 1 is the lowest income group. Panel B plots the share of industrial policy interventions (2010-2022) by region. See the July 2023 data update from Juhász et al. for details. Source: Juhász et al. (2022) (data update July 2023).*

**Figure 3.3. Share of industrial policies by measure type**



Notes: The figure plots the share of industrial policies accounted for by each measure type within a particular income group (top 10 measures reported). Income quintiles based on GDP per capita in 2010. Quintile 5 is the highest income group; quintile 1 is the lowest income group. See the July 2023 data update from Juhász et al. for details. Source: Juhász et al. (2022) (data update July 2023).

**Figure 3.4. Share of industrial policy interventions by sector**



Notes: Share of industrial policy interventions targeting a specific HS 2-digit sector by income quintiles (top 10 sectors reported; an industrial policy intervention can target multiple HS 2 sectors). Income quintiles based on GDP per capita in 2010. Quintile 5 is the highest income group; quintile 1 is the lowest income group. See the July 2023 data update from Juhász et al. for details. Source: Juhász et al. (2022) (data update July 2023).

**Figure 6.1: Embeddedness, autonomy, and the developmental state**

		Autonomy	
		low	high
Embeddedness	low	predatory state	Weberian regulatory state
	high	clientelist state	developmental state

Source: Authors, based on Wright (1996).



## Literature Cited

- Aghion P, Cai J, Dewatripont M, Du L, Harrison A, Legros P. 2015. Industrial policy and competition. *Am. Econ. J. Macroecon.* 7(4):1–32
- Amsden AH. 1992. *Asia's Next Giant: South Korea and Late Industrialization*. Oxford University Press. 2nd ed.
- Angrist JD, Pischke J-S. 2010. The credibility revolution in empirical economics: how better research design is taking the con out of econometrics. *J. Econ. Perspect.* 24(2):30–30
- Azoulay P, Fuchs E, Goldstein A, Kearney M. 2018. Funding breakthrough research: promises and challenges of the “ARPA model”. w24674, NBER. <https://www.nber.org/papers/w24674>
- Azoulay P, Graff Zivin JS, Li D, Sampat BN. 2019. Public R&D investments and private-sector patenting: evidence from NIH funding rules. *Rev. Econ. Stud.* 86(1):117–52
- Bai J, Barwick P, Cao S, Li S. 2022. Quid pro quo, knowledge spillover and industrial quality upgrading. w27644, NBER. <https://www.nber.org/papers/w27644>
- Baldwin R, Ito T. 2021. The smile curve: evolving sources of value added in manufacturing. *Can. J. Econ. Can. Déconomique.* 54(4):1842–80
- Bartelme D, Costinot A, Donaldson D, Rodríguez-Clare A. 2019. The textbook case for industrial policy: theory meets data. W26193, NBER. <https://www.nber.org/papers/w26193>
- Bartik T. 2019. Should place-based jobs policies be used to help distressed communities? W.E. Upjohn Institute for Employment Research
- Bartik T. 2020. Bringing jobs to people: improving local economic development policies. Aspen Institute
- Bartik T. 2022. How state governments can target job opportunities to distressed places. W.E. Upjohn Institute for Employment Research
- Barwick PJ, Kalouptsi M, Zahur NB. 2019. China's industrial policy: an empirical evaluation. w26075, NBER. <https://www.nber.org/papers/w26075>
- Beason R, Weinstein DE. 1996. Growth, economies of scale, and targeting in Japan (1955–1990). *Rev. Econ. Stat.* 78(2):286–95
- Blalock G, Gertler PJ. 2008. Welfare gains from foreign direct investment through technology transfer to local suppliers. *J. Int. Econ.* 74(2):402–21
- Blonigen BA. 2016. Industrial policy and downstream export performance. *Econ. J.* 126(595):1635–59
- Bloom N, Van Reenen J, Williams H. 2019. A toolkit of policies to promote innovation. *J. Econ. Perspect.* 33(3):163–84
- Bown CP, Crowley MA. 2014. Emerging economies, trade policy, and macroeconomic shocks. *J. Dev. Econ.* 111:261–73
- Bown CP, Tovar P. 2011. Trade liberalization, antidumping, and safeguards: evidence from India's tariff reform. *J. Dev. Econ.* 96(1):115–25
- Broda C, Limao N, Weinstein DE. 2008. Optimal tariffs and market power: the evidence. *Am. Econ. Rev.* 98(5):2032–65
- Cagé J, Gadenne L. 2018. Tax revenues and the fiscal cost of trade liberalization, 1792–2006. *Explor. Econ. Hist.* 70:1–24
- Chang H-J. 1993. The political economy of industrial policy in Korea. *Camb. J. Econ.* 17(2):131–57
- Cherif R, Hasanov F. 2019. The return of the policy that shall not be named: principles of industrial policy. IMF Working Papers. 2019/074. <https://www.imf.org/en/Publications/WP/Issues/2019/03/26/The-Return-of-the-Policy-That-Shell-Not-Be-Named-Principles-of-Industrial-Policy-46710>

- Chodorow-Reich G. 2019. Geographic cross-sectional fiscal spending multipliers: what have we learned? *Am. Econ. J. Econ. Policy*. 11(2):1–34
- Choi J, Levchenko AA. 2023. The long-term effects of industrial policy. w29263, NBER. <https://www.nber.org/papers/w29263>
- Cingano F, Palomba F, Pinotti P, Rettore E. 2022. Making subsidies work: rules vs. discretion. IZA Discussion Papers. 15172. <https://papers.ssrn.com/abstract=4114732>
- Corden WM. 1997. *Trade Policy and Economic Welfare*. Oxford University Press
- Criscuolo C, Lalanne G, Díaz L. 2022. Quantifying industrial strategies (QuIS): measuring industrial policy expenditures. OECD, Paris. [https://www.oecd-ilibrary.org/science-and-technology/quantifying-industrial-strategies-quis\\_ae351abf-en](https://www.oecd-ilibrary.org/science-and-technology/quantifying-industrial-strategies-quis_ae351abf-en)
- Criscuolo C, Martin R, Overman HG, Van Reenen J. 2019. Some causal effects of an industrial policy. *Am. Econ. Rev.* 109(1):48–85
- Diewert WE, Nakajima T, Nakamura A, Nakamura E, Nakamura M. 2011. Returns to scale: concept, estimation and analysis of Japan’s turbulent 1964–88 economy. *Can. J. Econ. Can. D’économique*. 44(2):451–85
- DiPippo G, Mazzocco I, Kennedy S, Goodman MP. 2022. Red ink: estimating Chinese industrial policy spending in comparative perspective. Center for Strategic and International Studies
- Dollar D, Sokoloff K. 1990. Patterns of productivity growth in South Korean manufacturing industries, 1963–1979. *J. Dev. Econ.* 33(2):309–27
- Duranton G, Venables AJ. 2019. Place-based policies: principles and developing country applications. Economics Series Working Papers. 893. <https://ideas.repec.org//p/oxf/wpaper/893.html>
- Evans PB. 1995. *Embedded Autonomy: States and Industrial Transformation*. Princeton, NJ: Princeton University Press
- Evenett SJ. 2009. Global trade alert: motivation and launch. *World Trade Rev.* 8(4):607–9
- Fernández-Arias E, Sabel C, Stein EH, Trejos A. 2016. *Two to Tango: Public-Private Collaboration for Productive Development Policies*. Washington, D.C.: Inter-American Development Bank
- Garin A, Rothbaum J. 2022. The long-run impacts of public industrial investment on regional development and economic mobility: evidence from World War II.
- Garred J. 2018. The persistence of trade policy in China after WTO accession. *J. Int. Econ.* 114:130–42
- Ghezzi P. 2017. Mesas ejecutivas in Peru: lessons for productive development policies. *Glob. Policy*. 8(3):369–80
- Giorcelli M, Li B. 2023. Technology transfer and early industrial development: evidence from the Sino-Soviet alliance. w29455, NBER. <https://www.nber.org/papers/w27503>
- Glaeser EL, Gottlieb JD. 2008. The economics of place-making policies. w14373, NBER. <https://www.nber.org/papers/w14373>
- Glick R, Moreno R. 1997. The East Asian miracle: growth because of government intervention and protectionism or in spite of it? *Bus. Econ.* 32(2):20–25
- Goldberg PK, Maggi G. 1999. Protection for sale: an empirical investigation. *Am. Econ. Rev.* 89(5):1135–55
- Goldberg PK, Pavcnik N. 2016. The effects of trade policy. In *Handbook of Commercial Policy*, Vol. 1, eds. K Bagwell, RW Staiger, pp. 161–206. Elsevier
- Gross DP, Sampat BN. 2023. America, jump-started: World War II R&D and the takeoff of the U.S. innovation system. w27375, NBER. <https://www.nber.org/papers/w27375>

- Gruber J, Johnson S. 2019. *Jump-Starting America: How Breakthrough Science Can Revive Economic Growth and the American Dream*. New York: Public Affairs Books
- Haddad M, Harrison A. 1993. Are there positive spillovers from direct foreign investment?: Evidence from panel data for Morocco. *J. Dev. Econ.* 42(1):51–74
- Haggard S. 1990. *Pathways from the Periphery: The Politics of Growth in Newly Industrializing Countries*. Cornell University Press
- Hanlon WW. 2020. The persistent effect of temporary input cost advantages in shipbuilding, 1850 to 1911. *J. Eur. Econ. Assoc.* 18(6):3173–3209
- Hanson G, Rodrik D. 2023. How does place based policy work? Reimagining the Economy Program. Harvard Kennedy School
- Harding T, Javorcik BS, Maggioni D. 2019. FDI promotion and comparative advantage
- Harrison A, Rodriguez-Clare A. 2010. Trade, foreign investment, and industrial policy for developing countries. In *Handbook of Development Economics*, Vol. 5, eds. D Rodrik, M Rosenzweig, pp. 4093–4214. Elsevier
- Harrison AAE. 1994. An empirical test of the infant industry argument: comment. *Am. Econ. Rev.* 84(4):1090–95
- Heid B, Larch M, Yotov YV. 2021. Estimating the effects of non-discriminatory trade policies within structural gravity models. *Can. J. Econ. Can. Déconomique.* 54(1):376–409
- Herman PR. 2022. A pragmatic approach to estimating nondiscriminatory non-tariff trade costs. *Rev. Int. Econ.* 30(4):1258–87
- Huang Y. 2000. An institutional and policy approach to the demand for foreign direct investment in China. *China Rev.* 249–84
- Imai M. 2009. Political determinants of government loans in Japan. *J. Law Econ.* 52(1):41–70
- Ito T, Krueger AO, eds. 1995. *Growth Theories in Light of the East Asian Experience*. Chicago, IL: University of Chicago Press
- Itskhoki O, Moll B. 2019. Optimal development policies with financial frictions. *Econometrica.* 87(1):139–73
- Javorcik SB. 2004. Does foreign direct investment increase the productivity of domestic firms? In search of spillovers through backward linkages. *Am. Econ. Rev.* 94(3):605–27
- Johnson C. 1982. *MITI and the Japanese Miracle: The Growth of Industrial Policy: 1925-1975*. Stanford University Press
- Juhász R. 2018. Temporary protection and technology adoption: evidence from the Napoleonic blockade. *Am. Econ. Rev.* 108(11):3339–76
- Juhász R, Lane N, Oehlsen E, Pérez VC. 2022. The who, what, when, and how of industrial policy: a text-based approach. <https://osf.io/preprints/socarxiv/uyxh9>
- Juhász R, Squicciarini M, Voigtländer N. 2023. Technology adoption and productivity growth: evidence from industrialization in France. w27503, NBER. <https://www.nber.org/papers/w27503>
- Kalouptsidi M. 2018. Detection and impact of industrial subsidies: the case of Chinese shipbuilding. *Rev. Econ. Stud.* 85(2):1111–58
- Kantor S, Whalley AT. 2023. Moonshot: public R&D and growth. w31471, NBER. <https://www.nber.org/papers/w31471>
- Kline P, Moretti E. 2014. People, places, and public policy: some simple welfare economics of local economic development programs. *Annu. Rev. Econ.* 6(1):629–62

- Kohli A. 2004. *State-Directed Development: Political Power and Industrialization in the Global Periphery*. Cambridge: Cambridge University Press
- König MD, Liu X, Zenou Y. 2019. R&D networks: theory, empirics, and policy implications. *Rev. Econ. Stat.* 101(3):476–91
- Krueger AO. 1986. The experience and lessons of Asia's super exporters. In *Export-Oriented Development Strategies*. Routledge
- Krueger AO. 1995. East Asian experience and endogenous growth theory. In *Growth Theories in Light of the East Asian Experience*, Vol. 4, eds. AO Krueger, T Ito, pp. 9–36. University of Chicago Press
- Krueger AO. 2019. The experience and lessons of Asia's super exporters. In *Export-Oriented Development Strategies*, ed. V Corbo, pp. 187–212. Routledge. Second ed.
- Krueger AO, Tuncer B. 1982. An empirical test of the infant industry argument. *Am. Econ. Rev.* 72(5):1142–52
- Krugman PR. 1994. The myth of Asia's miracle. *Foreign Aff.* 73(6):62–74
- Lal D. 1983. *The Poverty of Development Economics*. Institute of Economics Affairs
- Lall S. 1996. 'The East Asian miracle' study: does the bell toll for industrial strategy? In *Learning from the Asian Tigers: Studies in Technology and Industrial Policy*, ed. S Lall, pp. 107–23. London: Palgrave Macmillan UK
- Lane N. 2020. The new empirics of industrial policy. *J. Ind. Compet. Trade.* 20(2):209–34
- Lane N. 2022. Manufacturing revolutions: industrial policy and industrialization in South Korea. CSAE Working Paper Series. <https://doi.org/10.2139/ssrn.3890311>
- LaPoint C, Sakabe S. 2021. Place-based policies and the geography of corporate investment. <http://dx.doi.org/10.2139/ssrn.3950548>
- Lashkaripour A, Lugovskyy V. 2023. Profits, scale economies, and the gains from trade and industrial policy. [https://alashkar.pages.iu.edu/Lashkaripour\\_Lugovskyy\\_AER\\_2023.pdf](https://alashkar.pages.iu.edu/Lashkaripour_Lugovskyy_AER_2023.pdf)
- Lawrence RZ, Weinstein DE. 2001. Trade and growth: import led or export led? Evidence from Japan and Korea. In *Rethinking the East Asian Miracle*, eds. JE Stiglitz, S Yusuf. Oxford: Oxford University Press
- Lee J-W. 1996. Government interventions and economic growth. *J. Econ. Growth.* 1(3):
- Lee K, Hahn D, Lin J. 2002. Is China following the East Asian model? A "comparative institutional analysis" perspective. *China Rev.* 2(1):85–120
- Lerner J. 2009. *The Boulevard of Broken Dreams: Why Public Efforts to Boost Entrepreneurship and Venture Capital Have Failed—and What to Do about It*. Princeton, NJ: Princeton University Press
- Limão N, Tovar P. 2011. Policy choice: theory and evidence from commitment via international trade agreements. *J. Int. Econ.* 85(2):186–205
- Lin JY. 2012. *New Structural Economics: A Framework for Rethinking Development and Policy*. Washington, D.C.: The World Bank
- Liu E. 2019. Industrial policies in production networks. *Q. J. Econ.* 134(4):1883–1948
- Liu E, Ma S. 2023. Innovation networks and R&D allocation. w29607, NBER. <https://www.nber.org/papers/w29607>
- Lucas RE. 1984. An empirical test of the infant industry argument: comment. *Am. Econ. Rev.* 74(5):1110–11
- Malesky E, London J. 2014. The political economy of development in China and Vietnam. *Annu. Rev. Polit. Sci.* 17:395–419
- Maloney WF, Nayyar G. 2018. Industrial policy, information, and government capacity. *World Bank Res. Obs.* 33(2):189–217

- Mazzucato M. 2014. *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*. New York: Anthem Press. Rev. ed.
- Mazzucato M. 2021. *Mission Economy: A Moonshot Guide to Changing Capitalism*. Harper Business
- Mazzucato M, Kattel R, Ryan-Collins J. 2020. Challenge-driven innovation policy: towards a new policy toolkit. *J. Ind. Compet. Trade*. 20(2):421–37
- McKenzie D, Woodruff C, Bjorvatn K, Bruhn M, Cai J, et al. 2021. Training entrepreneurs. *VoxDevLit*. 1(2):3–27
- Mitrunen M. 2021. Industrial policy, structural change and intergenerational mobility: evidence from the Finnish war reparations
- Moretti E, Steinwender C, Van Reenen J. 2019. The intellectual spoils of war? Defense R&D, productivity and international technology spillovers
- Myers KR, Lanahan L. 2022. Estimating spillovers from publicly funded R&D: evidence from the US department of energy. *Am. Econ. Rev.* 112(7):2393–2423
- Nelson RR, Pack H. 1999. The Asian miracle and modern growth theory. *Econ. J.* 109(457):416–36
- Neumark D, Simpson H. 2015. Place-based policies. In *Handbook of Regional and Urban Economics*, Vol. 5, pp. 1197–1287. Elsevier
- Noland M. 2004. Selective intervention and growth: the case of Korea. In *Empirical Methods in International Trade: Essays in Honor of Mordechai Kreinin*, ed. MG Plummer, pp. 229–46. Cheltenham, UK; Northampton, MA: Edward Elgar
- Noland M, Pack H. 2003. *Industrial Policy in an Era of Globalization: Lessons from Asia*. Peterson Institute for International Economics
- Pack H. 2000. Industrial policy: growth elixir or poison? *World Bank Res. Obs.* 14(1):47–67
- Pack H, Saggi K. 2006. Is there a case for industrial policy? A critical survey. *World Bank Res. Obs.* 21(2):267–97
- Page J. 1994. The East Asian miracle: four lessons for development policy. *NBER Macroecon. Annu.* 9:219–69
- Pons-Benaiges O. 2017. Did government intervention target technological externalities? Industrial policy and economic growth in postwar Japan, 1964-1983. [https://delong.typepad.com/06\\_tech\\_externalities\\_japan.pdf](https://delong.typepad.com/06_tech_externalities_japan.pdf)
- Ray EJ, Marvel HP. 1984. The pattern of protection in the industrialized world. *Rev. Econ. Stat.* 452–58
- Redding SJ, Turner MA. 2015. Transportation costs and the spatial organization of economic activity. In *Handbook of Regional and Urban Economics*, Vol. 5, eds. G Duranton, JV Henderson, W Strange, pp. 1339–98. Elsevier
- Rodrik D. 1995. Getting interventions right: how South Korea and Taiwan grew rich. *Econ. Policy*. 10(20):53–107
- Rodrik D. 2012. Why we learn nothing from regressing economic growth on policies. *Seoul J. Econ.* 25(2):137–51
- Rodrik D. 2016. Premature deindustrialization. *J. Econ. Growth*. 21(1):1–33
- Rodrik D. 2022. An industrial policy for good jobs. Policy Proposals. Hamilton Project. <https://www.hamiltonproject.org/publication/policy-proposal/an-industrial-policy-for-good-jobs/>
- Rodrik D, Hausmann R. 2003. Economic development as self-discovery. *J. Dev. Econ.* 72:
- Rodrik D, Hausmann R. 2006. Doomed to choose: industrial policy as predicament. Harvard Kennedy School. <https://drodrik.scholar.harvard.edu/publications/doomed-choose-industrial-policy-predicament>

- Rodrik D, Rodriguez F. 2001. Trade policy and economic growth: a skeptic's guide to the cross-national evidence. *NBER Macroecon. Annu.* 15:261–325
- Rodrik D, Sabel CF. 2022. Building a good jobs economy. In *A Political Economy of Justice*, eds. D Allen, Y Benkler, L Downey, R Henderson, J Simons. University of Chicago Press
- Sabel CF, Victor DG. 2022. *Fixing the Climate: Strategies for an Uncertain World*. Princeton University Press
- Saure P. 2007. Revisiting the infant industry argument. *J. Dev. Econ.* 84(1):104–17
- Saxonhouse GR. 1979. Industrial restructuring in Japan. *J. Jpn. Stud.* 5(2):273–320
- Schweiger H, Stepanov A, Zacchia P. 2022. The long-run effects of R&D place-based policies: evidence from Russian science cities. *Am. Econ. J. Econ. Policy.* 14(3):322–51
- Shu P, Steinwender C. 2019. The impact of trade liberalization on firm productivity and innovation. *Innov. Policy Econ.* 19(1):39–68
- Slattery C, Zidar O. 2020. Evaluating state and local business incentives. *J. Econ. Perspect.* 34(2):90–118
- Stiglitz JE. 1996. Some lessons from the East Asian miracle. *World Bank Res. Obs.* 11(2):151–77
- Stiglitz JE, Uy M. 1996. Financial markets, public policy, and the East Asian miracle. *World Bank Res. Obs.* 11(2):249–76
- Studwell J. 2013. *How Asia Works: Success and Failure in the World's Most Dynamic Region*. Grove Press and Grove/Atlantic, Inc.
- Teranishi J. 1986. Economic growth and regulation of financial markets: Japanese experience during postwar high growth period. *Hitotsubashi J. Econ.* 27(2):149–65
- Thun E. 2004. Industrial policy, Chinese-style: FDI, regulation, and dreams of national champions in the auto sector. *J. East Asian Stud.* 4(3):453–89
- Thun E. 2006. *Changing Lanes in China: Foreign Direct Investment, Local Governments, and Auto Sector Development*. Cambridge University Press
- Trezise PH. 1983. Industrial policy is not the major reason for Japan's success. *Brookings Rev.* 1(3):13–18
- Uriu RM. 1984. The declining industries of Japan: adjustment and reallocation. *J. Int. Aff.* 38(1):99–111
- Uriu RM. 1996. *Troubled Industries: Confronting Economic Change in Japan*. Cornell University Press
- Vittas D, Cho YJ. 1996. Credit policies: lessons from Japan and Korea. *World Bank Res. Obs.* 11(2):277–98
- Vogel EF. 1991. *The Four Little Dragons: The Spread of Industrialization in East Asia*. Cambridge, MA: Harvard University Press
- Wade R. 1990. *Governing the Market: Economic Theory and the Role of Government in East Asian Industrialization*. Princeton University Press
- Wade RH. 2019. East Asia. In *Asian Transformations: An Inquiry into the Development of Nations*, ed. D Nayyar. Oxford University Press
- Westphal LE. 1990. Industrial policy in an export-propelled economy: lessons from South Korea's experience. *J. Econ. Perspect.* 4(3):41–59
- World Bank. 1993. *The East Asian Miracle: Economic Growth and Public Policy: A World Bank Policy Research Report*. Oxford University Press
- Wright EO. 1996. Review of Peter Evans, Embedded Autonomy. *Contemp. Sociol.* 25(2):176–79
- Ye M. 2009. Policy learning or diffusion: how China opened to foreign direct investment. *J. East Asian Stud.* 9(3):399–432
- Yoo J. 1990. The industrial policy of the 1970s and the evolution of the manufacturing sector in Korea. KDI Working Papers. 9017.  
[http://www.kdi.re.kr/kdi\\_eng/publications/publication\\_view.jsp?pub\\_no=863](http://www.kdi.re.kr/kdi_eng/publications/publication_view.jsp?pub_no=863)