



STEG WORKING PAPER

THE WHO, WHAT, WHEN, AND HOW OF INDUSTRIAL POLICY: A TEXT-BASED APPROACH

Réka Juhász, Nathan Lane,
Emily Oehlsen, and Verónica C. Pérez

JANUARY 2023
STEG WP050
SRG 78

The Who, What, When, and How of Industrial Policy: A Text-Based Approach*

Réka Juhász Nathan Lane Emily Oehlsen Verónica C. Pérez

(UBC, NBER & CEPR)

(Oxford)

(Oxford)

(Columbia)

December, 2022

Abstract

Since the 18th century, policymakers have debated the merits of industrial policy (IP). Yet, economists lack measures and data on its use. We provide a new approach to measuring industrial policy from text and study its global patterns. We create an automated classification algorithm and categorize policies from a global database of commercial policy descriptions, 2009–2020. By quantifying policy at the country, industry, and year levels, we provide a first disaggregate analysis of international industrial policies. We highlight four findings. First, IP is common (25% of policies in our database) and has expanded since 2010. Second, instead of blunt tariffs, IP is granular and technocratic. Countries tend to use subsidies and export promotion measures, often targeted at individual firms. Third, the countries engaged most in IP tend to be wealthier (top income quintile) liberal democracies. In our data, IP is rarer among the poorest nations (bottom quintile). Fourth, IP is targeted toward a subset of industries and is highly correlated with an industry's revealed comparative advantage. Our approach to measuring industrial policy shows that contemporary practice is likely much different from the past.

JEL: L52, C38

Keywords: Industrial policy, Supervised Machine Learning, Natural Language Processing

*We thank Elliott Ash, David Byrne, Chiara Criscuolo, John Andrew Fitzgerald, Amy Handlan, Gordon Hanson, Tarek Hassan, Myrto Kalouptsi, Guy Lalanne, Colin McAuliffe, Phillip McCalman, Suresh Naidu, Ashwin Narayan, Gaurav Nayyar, Dani Rodrik, Francesco Squintani, Chris Woodruff, and seminar audiences at Brown, Columbia, ETH Zurich, Harvard, the IGC-Stanford Firms, Development and Trade conference, Melbourne University, Michigan, Pakistan Institute of Development Economics, and Warwick for helpful comments and suggestions. We benefited from conversations with Gerard DiPippo, Ilaria Mazzocco, and scholars at the CSIS; the OECD Productivity, Innovation and Entrepreneurship team; and the UNIDO Research and Industrial Policy Advice Division. We especially thank Johannes Fritz for helpful discussions about the *Global Trade Alert* project. Ida Brzezinska, Caran Niousha Etemad, Max Marczinek, Sarah Wappel, Megan Yamoah, and a team of Columbia and Oxford annotators provided outstanding research support. Funding from ISERP and STEG/CEPR is gratefully acknowledged.

1 Introduction

A major question across the social sciences is whether the state should shape the composition of economic activity. This state action, known as industrial policy, is controversial. The debate about the state's role in the process of development is as old as modern capitalism (List, 1856; Taussig, 1914; Chang, 2002). Many scholars view the state as an essential, active promoter of development, emphasizing the importance of industrial policy in economic transformation (Wade, 1990; Amsden, 1992; Evans, 1995). Others have argued that industrial policies are harmful (Krueger, 1990; Pack, 2000), exemplified by Gary Becker's quip "The best industrial policy is none at all."¹ Although these questions surrounding industrial policy are fundamental, the scarcity of basic facts about its practice is remarkable. We lack measures and comparative data on industrial policy, and consequently, do not have a systematic picture of industrial policy use.

This study addresses an empirical deficit surrounding industrial policy. We propose a new text-based approach to measuring industrial policy. Our method uses machine learning to automatically classify industrial policies from policy text. By applying our classification algorithm to a comprehensive global policy database, we quantify industrial policy use at the country, industry, and year levels. We use our disaggregated measures to study the contours of current international industrial policy practice. Specifically, we provide a first descriptive analysis of industrial policy and consider the who, what, when, and how of policy use, from 2009–2020.

The first part of this paper presents our text-based approach for measuring the industrial policy goals of economic interventions. Unlike many policies, industrial policies are often not directly observable. Consider tariffs, a canonical tool of industrial policy. Although historically used to protect infant industries, tariffs are also used to raise government revenue (Cagé and Gadenne, 2018) and may be used for other objectives, like terms of trade rationales (Broda and Weinstein, 2004). Alone, a tariff rate cannot tell us much about industrial policy. This challenge is compounded by the fact that industrial policies take the form of non-tariff measures (NTMs). Attributing entire classes of policy measures, *ex-ante*, as industrial policy, can lead to confusion about the relationship between trade policy, industrial policy, and development (Rodriguez and Rodrik, 2001; Harrison and Rodríguez-Clare, 2010; Rodrik, 2012). These comparative questions require precise measures of IP. In short, researchers require more information to discern whether a policy is truly an industrial policy.

1. Business Week, 26 August 1985.

To measure industrial policy, our approach turns to the language and objectives of policies. More precisely, we consider the extent to which the descriptions of policy reflect industrial policy goals—the intent to shape the composition of economic activity. These transformative goals take many forms. They can be a policymaker targeting an infant industry to change domestic production. They can be a policymaker incentivizing R&D to expand innovative activity. They can be a policymaker subsidizing green technology to move away from fossil fuels. Or they can be a policymaker reviving a sunset industry facing economic headwinds. In each case, policymakers intend to alter the structure of the economy.

Industrial policy objectives often permeate the way political actors characterize policy. Consider the following description of a Chinese subsidy from our paper’s dataset:

“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. [...]”

Here, the goal (italicized) is clearly stated in the opening sentence. The policymaker wants to shape the composition of the economy by boosting a particular sector. In this case, vehicle batteries. This example is textbook industrial policy. This language is not anomalous—indeed, we show that policy descriptions often express the goal of the policy. Our method draws from this information to distinguish measures with industrial policy goals from those with different goals (e.g., public health-oriented interventions).

Our approach uses machine learning to automatically classify industrial policies using text. We draw from an emergent social science literature, which uses supervised machine learning to categorize complex concepts (e.g., ideology, populism, and inequality) using textual documents (Nelson, Burk, Knudsen and McCall, 2018; Grimmer, Roberts and Stewart, 2022). We develop a classification algorithm and make novel use of an exhaustive international corpus of commercial policy, *The Global Trade Alert* database (Evenett, 2019), which tracks and collects English-language descriptions of new government policy announcements (Section 3 describes the project’s scope). We apply our classification algorithm to this rich data to quantify industrial policy practice at the country, sector, and year levels.

We can summarize our automatic classification process in four steps. First, we develop a parsimonious definition of industrial policy goals. Annotators use this definition to manually categorize, or “label,” a subset of policy descriptions in our database (~ 7% or approximately 2000 observations). Next, we use this data to train a simple classification model and assess its performance. For transparency, this

paper focuses on a simple binary classifier, logistic regression. Third, we use this classification model to automatically categorize, or predict, the remaining policies in our database (~28,000 observations). In effect, this supervised classification algorithm replicates the manual categorization of policies to classify policy text at scale.

Our classification process delivers two important results. First, annotators successfully use our formal definition to categorize industrial policies. They largely agree on what industrial policy is, with minimal guidance. Second, we test our baseline classifier on labeled data not seen (“held out”) by the model during training and show it performs well across metrics. In terms of accuracy, our model correctly classifies about 93 percent of all policies. In terms of precision, 87 percent of measures identified as industrial policy by our model are indeed industrial policy. In terms of recall, 88 percent of the actual industrial policy measures in our data are correctly identified as industrial policy by the baseline model.²

We demonstrate the validity of our classification model in multiple ways. We establish ‘face validity’ (Grimmer et al., 2022, p. 31) for our baseline logistic classifier and show that words, or “tokens,” most predictive of industrial policy are reasonable.³ For example, predictive tokens include ‘technology’, ‘development’, and ‘green’. We establish ‘hypothesis validity’ and show that our findings adhere to anticipated, real-world patterns (ibid). For example, we consider the policy dynamics during the COVID-19 pandemic; here, we expect non-industrial policies to rise due to the enormous quantity of emergency health and social welfare measures (both are not industrial policies under our formal definition). Indeed, our data detects a sharp rise in *non-IP* crisis measures at the height of the 2020 pandemic, while IP use declines slightly. Our data also confirms widely reported accounts of the rise in industrial policy over the full sample period. The trends in our data match qualitative reports (e.g., the IMF’s (Cherif and Hasanov, 2019) and recent national accounting studies (e.g., DiPippo, Mazzocco and Kennedy (2022), who study eight countries) on the proliferation of industrial policy.

The second part of our paper uses our measures of industrial policy to establish basic facts surrounding policies. We highlight four findings. First, we show that industrial policy is fairly common (25% of policies in our database) and is on the rise. International industrial policies doubled through the 2010s. Second, we find that industrial policy is technocratic and granular. By this, we mean that the majority of industrial policies identified in our data take the form of subsidies and export-related

2. Accuracy is the number of total true positives and true negatives divided by the total number of observations. Precision is the number of true positives divided by the total sum of true positives and false positives. Recall is the number of true positives divided by the total sum of true positives and false negatives.

3. Face validity requires that, at a minimum, the measures pass inspection by a domain expert.

measures. Importantly, these measures tend to be targeted at individual firms. These patterns suggest that current industrial policy practice differs substantially from the blunt import tariffs used in the past.

Third, we look at who uses industrial policies. We find that high-income countries extensively use industrial policies. Our analysis shows a robust positive correlation between income and industrial policy. This pattern is robust to controlling for different reporting biases and is robust to alternative measures of industry policy use. The wide use of industrial policies among wealthier countries is striking, given these policies are popularly discussed in the context of developing economies.

Fourth, we show countries are selective in the sectors they target with industrial policies. That is, industrial policies are usually directed toward a subset of sectors. We also document a striking systematic pattern in which industries receive these policies. We find that within countries, industrial policy is targeted at sectors that have higher revealed comparative advantage in international trade. We show this correlation holds at different levels of aggregation (Harmonized System 2 and 6-digit level data), and controlling for a variety of time-varying fixed effects. In other words, on average, countries direct new industrial policies toward sectors in which they have an established international presence.

Thus, our new descriptive analysis provides substantive insights into the practice of industrial policy. First, the positive relationship between income and industrial policy is consistent with a multitude of qualitative work about industrial policy from bureaucratic quality and autonomy (Evans, 1995), state capacity (Maloney and Nayyar, 2018), and binding international constraints on low-income countries (Chang, 2002; Wade, 2009). In light of our facts, if industrial policy practice requires a state to use technocratic tools targeted at individual firms, the poorest countries will be limited by fiscal and administrative capacity to do the same. Worryingly, if the world economy is characterized by increasing industrial policy, as our findings suggest, it may be difficult for the poorest countries to compete in international markets.

Second, autocracies and late-industrializers have long been associated with industrial policy (Gerring, Gjerløw and Knutsen, 2022). While there are authoritarian and weak democracies among the largest users of industrial policy in our data (e.g., China, Russia and Saudi Arabia), wealthier democracies are highly represented. In fact, one important finding of our study is that IP is used heavily across rich countries, which runs counter to conventional wisdom. IP is not the purview of middle-income autocracies, but, instead, is prevalent in rich, liberal democracies.

Our paper relates to the literature on industrial policy practice in multiple ways. First, our study contributes to the notoriously thin literature on the empirics of industrial policy (Harrison and Rodríguez-Clare, 2010). Using our new, text-based

approach, we are able to provide a comprehensive analysis of IP for the first time. Recently, individual case studies have emerged to refine our understanding of IP in precise domains. These include well-identified historical episodes (Juhasz, 2018; Hanlon, 2020; Mitrunen, 2019; Lane, 2021), and contemporary case studies (Kalouptsidi, 2018; Barwick, Kalouptsidi and Zahur, 2019). Similarly, quantitative work has emerged to address and test specific theoretical motivations for policy (Liu, 2019; Bartelme, Costinot, Donaldson and Rodriguez-Clare, 2019). We complement this emerging literature, by providing what is, to the best of our knowledge, the first descriptive cross-country, cross-industry analysis of the practice of industrial policy. Importantly, this dataset will be continuously updated, providing a “live” database on industrial policy.

Second, we contribute to the literature on trade policy and development by confronting key issues: i) a need to move beyond import tariffs when measuring trade policy; and ii) better measurement of trade policy (Goldberg and Pavcnik, 2016). For the case of industrial policies, we show that both points have practical relevance. Our results suggest that import tariffs are not the main policy lever used for industrial policy. Moreover, a narrow focus on import tariffs as a proxy for industrial policy gives a misleading picture. Although the use of tariffs has fallen in developed countries, these same economies are the ones that use sophisticated, well-targeted NTMs extensively for industrial policy purposes. Our methods illustrate the potential of using text and alternative data to confront issues of measurement in the trade policy literature.

Third, we contribute to the literature on the developmental state by building data for comparative study of industrial policy use. Seminal qualitative work has long considered the role of the state in promoting development (Evans, 1995; Kohli, 2004; Haggard, 2015), of which industrial policy plays a critical role (Johnson, 1982; Amsden, 2001). We view our work as a quantitative complement to this qualitative work in comparative social science. By providing systematic data on real-world industrial policy practice, we hope to better understand what makes for effective—and ineffective—industrial policies. The success of these policies likely depends on institutional characteristics and important factors, like state capacity. Our study provides inputs into further empirical work on the relationship between state capacity and long-run development (Besley and Persson, 2010; Gennaioli and Voth, 2015).

Last, we contribute to text-as-data literature in the social sciences, which has used insights from natural language processing (NLP) and machine learning to transform unstructured data into inputs for empirical analysis (D’Orazio, Landis, Palmer and Schrod, 2014; Gentzkow, Kelly and Taddy, 2017; Grimmer et al., 2022). We join Baker, Bloom and Davis (2016) and Hassan, Hollander, Van Lent and Tahoun (2019), who

combine text mining methods with mass textual corpus to capture hard-to-measure quantities (economic policy uncertainty and firm-level risk, respectively) impacting the economy. We depart from dictionary-based methods of measurement and show the potential for supervised learning methods to build policy-related indices at scale.

We organize our study in the following way. In Section 2, we describe the measurement issues surrounding industrial policy and detail our approach to measurement. In Section 3, we describe the data we use in this study, and focus on the text and meta data of our database. We then detail our methodology in Section 4 and describe the details of our four-step workflow. We use the outputs of this workflow for our descriptive analysis. In Section 5, we our descriptive analysis and provide stylized facts of industrial policy. Finally, we conclude in Section 6.

2 Measurement

We now turn to the issues of measurement and industrial policy. First, we start with practical issues of observability and show how policy language is useful for capturing key aspects of industrial policy. Second, these issues inform our conceptual approach to measuring police. We describe our concept of industrial policy and provide a formal definition of IP which can be taken to data. Finally, we show the implications of classifying descriptions of policy using this definition.

2.1 OBSERVABILITY

Researchers face the following dilemma: we rarely observe industrial policies directly in the real world. This issue of observability hexes the study of industrial policy. These interventions can take many forms, and, importantly, the same measures (tariffs, subsidies, etc.) can be used for multiple ends—not merely the goals associated with industrial policy. Thus, observing a policy alone cannot tell us whether a policy is an industrial policy or one used for another end.

Consider tariffs, a measure historically associated with industrial policy. Although a classic means of infant industry promotion (Baldwin, 1969), tariffs have also been widely used to raise government revenue (Johnson, 1951) and may be applied for terms of trade rationales (Broda and Weinstein, 2004). Thus, a tariff rate alone cannot tell us whether this tariff embodies industrial policy or other objectives.

This ambiguous nature of tariffs matters. For instance, developing economies have historically relied on tariffs as a source of fiscal revenue (Balassa, 1989). Accordingly, the poorest and least fiscally developed countries may be the same countries most

likely to use tariffs for revenue (Cagé and Gadenne, 2018). Data on these interventions are easily collected and thus readily observed the econometrician, while more technocratic ones—requiring both more fiscal and administrative capacity—are not (Kalouptsidi (2018)). As a consequence, approximating for industrial policy with off-the-shelf tariffs may present a distorted picture of industrial policy use. Our results suggest that this is the case.

The industrial policy toolbox is also vast, compounding issues of observability. Like tariffs, various non-tariff measures used as industrial policy can serve other goals. As tariffs became increasingly limited by international agreements, many suspect countries have substituted away from tariffs (Goldberg and Pavcnik, 2016). Through the post-war period, an increasingly complicated menu of policy levers have emerged, many of which are more complicated to observe and categorize in cross-country datasets (Anderson and Wincoop, 2004).

This inability to distinguish the goals of policy measures has fed debates about the efficacy of IP (Harrison and Rodríguez-Clare, 2010), and the relationship between trade policy and development, more broadly (Rodriguez and Rodrik, 2001; Rodrik, 2012). Lehmann and O'Rourke (2011) show that the relationship between tariffs and sectoral growth is sensitive to whether tariffs are used for revenue or developmental ends. Early empirical work on IP fail to distinguish the ends of policy (see: Rodriguez and Rodrik (2001); Lane (2020)). Blunt measures of IP make it difficult to disentangle the impacts of IP from the multitude of objectives, such as optimal sunset policies (correlated with industry decline) or revenue extraction during export booms (correlated with industry ascent).

Simply put, industrial policy scholars face a conundrum: discerning which policies are, indeed, industrial policies. Rather than the policies quantities themselves, we consider the contextual language of policies. To do so, we turn to the tools of supervised machine learning, to categorize, *en masse*, policies as industrial policy. More precisely, we appeal to supervised machine learning to classify text into categories informed by labeled input data. Increasingly, social scientists use these methods to categorize complex, socially constructed, theoretical concepts, such as populism or inequality (Nelson et al., 2018, p. 204). Industrial policy is a similar, theoretical concept. As we have argued above, it is difficult to readily observe industrial policy. However, as we discuss below, the concept of industrial policy and its goals are well-defined.

2.2 THE CONCEPT OF INDUSTRIAL POLICY

To identify industrial policies, we must be clear about this concept. We refer to industrial policy as strategic state action directed at changing the structure of the domestic economy. Industrial policy is activism motivated by a long-run goal—a vision for what the economy should look like. Best summarized by Chalmers Johnson, “[t]he very existence of industrial policy implies a strategic, or goal-oriented, approach to the economy” (Johnson, 1982, p. 19). Broadly, this conception of industrial policy is useful in that it encompasses two policy views: those of neoclassical economists and those of scholars of the developmental state. Both see industrial policy as state action executed in the service of broader goals.

First, mainstream economists conceptualize industrial policy largely using the concepts of neoclassical economics. Here, justifications for state intervention rest on the existence of market failures, such as externalities and credit market imperfections (Corden, 1997). In this view, the state acts with a goal in mind: intervening when the market, left to its own devices, cannot achieve an efficient allocation of resources. To overcome market failures, industrial policies act to change the relative prices, shaping the composition of economic activity in the service of an ultimate end goal. Classic examples in this framework are subsidies to an infant industry to realize external economies of scale, or to innovative activity in an industry where such activity may be undersupplied.

Second, our conception of industrial policy encompasses another tradition, which we refer to as the developmental state view. This view takes a more expansive view of industrial policy. Where economists start with market failures, developmentalists see the state as necessary, or at least a critical actor, in achieving larger economic ends. To promote industrial development, the state may organize capital markets (e.g. Gerschenkron (1962)) or direct entrepreneurs to certain activities (e.g. Hirschman (1958)) to name but a few examples (see Evans (1995, pp. 28-32) for an overview). For these scholars, examples of the developmental state at work include post-war Japan (Johnson, 1982), Korea and Taiwan (Wade, 1990; Amsden, 1992). The developmental state takes a more expansive stand on the desirable structure of the economy and sees a more expansive role for the state in achieving that structure.

This idea of a goal-oriented, activist state has a long history. Before the developmental state tradition, earlier “Neomercantilist” ideas about the desirability of industrial policy have roots that date back to the early days of the Industrial Revolution. To emulate British industrialization and dominance, thinkers like Alexander Hamilton, Friedrich List, and Henry Carey also advocated for IP. To achieve commercial wealth and power, neomercantilists envisaged an activist state using strategic

trade protection and other levers to change the composition of economic activity in service of this goal.

Our concept of industrial policy is broad enough to encompass different views guiding state activism. This is important because our goal is to define industrial policy in a way that can be used to classify (i.e., determine the 'IP-ness') of actual, real-world policies. Our concept takes the common components of these traditions: the goal-oriented nature of IP and its intent of changing the composition of economic activity.

2.3 DEFINING INDUSTRIAL POLICY

We now codify our concept of industrial policy using a formal definition, which focuses on the goals of policy. This formal definition allows us to be categorizing descriptions of industrial policy in practice. Importantly, this formal definition is not new; we build on decades of work from scholars and practitioners of industrial policy (United States International Trade Commission, 1983; Warwick, 2013; DiPippo et al., 2022; Criscuolo, Gonne, Kitazawa and Lalanne, 2022).⁴ Our formal definition is as follows,

Formal Definition of Industrial Policy

- A) Stated goal - Industrial policy is goal oriented state action. The purpose is to shape the composition of economic activity. Specifically: industrial policy seeks to change the relative prices across sectors or direct resources towards certain selectively targeted activities (e.g., exporting, R&D), with (ii) the purpose of shifting the long-run composition of economic activity.
- B) National state implementation - Industrial policy is aimed at the stated goals at the level of the national economy. Specifically: industrial policy action is taken by a national, or extranational, state. These actions are sanctioned and financed by national governments, supranational bodies, or amalgamations of these units.

Part (A) captures the fact that industrial policy is goal-oriented, with a vision of shaping forms of economic activity within an economy. This follows directly from our concept of IP. The goal could be the correction of a market failure in the service of growth (the neoclassical view), achieving structural transformation (the

4. See important contemporary thinking about industrial policy from DiPippo et al. (2022) and Criscuolo et al. (2022).

developmentalist view), geopolitical power (the neo-mercantilist view). By pursuing these goals, the state takes a stand on the desirable structure of the economy and plays an active role in achieving those goals. Note that in emphasizing the goals of policy, part (A) of our definition encompasses different intellectual traditions.

Part (A) implies that industrial policy has specificity in how it shapes the economy. It can shape different dimensions of economic activity. First, industrial policies may be sectoral—or “vertical”—and these policies often constitute more traditional forms of industrial policy (e.g. infant industry policy aimed at steel or textiles). This type of sectoral policy has intellectual roots in all theories of IP discussed above and is also typically included in measurement exercises. Second, policies can shape *forms* of economic activity, targeting types of activity (e.g., exporting, innovation or de-carbonization) rather than sectors. Take the case of innovation, a textbook case of a market failure described above, and also a type of IP aimed at shaping the quantity and type of knowledge production an economy undertakes. Thus, our definition is agnostic to whether a policy is sectoral or not.

Our definition is also agnostic about which measures are used as industrial policy. Part (A) reflect states acting according to economic goals, rather than the ends to these goals. We view this as a key strength of our definition and more generally, our approach to measurement.

Part (B) is about administrative scope of policy. Nearly all definitions of industrial policy involve states as actors Warwick (2013). Industrial policy is state action—industrial policies transform a national economy. Thus, the second part of the definition means that industrial policy is aimed at transforming a national or supra-national economy, by the political actors. Thus, in our study, we are interested in economies (e.g. Guatemala) or aggregations of national economies (e.g. the European Union), where state actors act consciously to transform the economies of their constituencies. ⁵

Thus, we exclude sub-national policies, which can include many “place-based” policies, as we exclude policies implemented by sub-national administrative units that are disconnected from national economic goals. Our rationale is that sub-national interests and objectives may conflict with national interests, or compete with other sub-national jurisdictions. Thus, such policies may not reflect political actors shaping the structure of *national* economic activity, explicitly.

5. We follow the literature and consider the state as an aggregation beyond merely the government or the bureaucracy.

2.4 IMPLICATIONS AND EXAMPLES

This formal definition is used to categorize policies using their text. Let us briefly describe what types of policies are included and excluded by our definition. We believe most of these implications are standard and they comport with many current definitions and practices DiPippo et al. (2022). We justify cases where the implications are more nuanced.

Implications

Implications we consider are,

The Definition Excludes Policies Without a Bias - Part (A) of our formal definition entails a state changing the composition of economic activity and doing so with specificity. First and foremost, this excludes policies that do not discriminate between sectors or forms of activity (e.g., a program intended to promote youth employment). These policies typically fail to shape the composition of economic activity in an intentional way.

The Definition Excludes Regional Policies - Our definition involves the goals and actions of the state at the national level. As such, we exclude the activity of local governments, constituencies, and other sub-national units.

The Definition Excludes Policies in Response to Shocks and Social Programs - Industrial policy excludes policy support in response to shocks (e.g., business cycle fluctuations) and measures that act as social welfare programs. Part (A) of our definition precludes either type of policy.

The Definition Does Not Exclude Agriculture and Primary Industries - Nothing about the formal definition excludes agriculture, or other sectors not traditionally considered the domain of industrial policy (commodity production, natural resources, forestry, etc.).

Indeed, industrial policy has historically been the purview of manufacturing. Nevertheless, conceptually, IP is agnostic about the sectors targeted thus we do not exclude entire sectors, ex-ante. Moreover, including non-manufacturing sectors is also important, in light of industrial policy practice in the developing world, where these activities are potentially substantial.

Examples

Let us consider two descriptions of policy and examine them in light of the definition. First, consider again the description of a Chinese policy from 2017, discussed in the introduction:

“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. One of the instruments mentioned in the plan to be used to achieve this is the use of specific government ‘funds’ to invest into relevant firms. The release contemplates the establishment of an industry-specific development fund as well as utilizing existing government funds designated for the IT industry more generally.”

This policy is the type of intervention many have in mind when discussing Chinese industrial policy. Using our definition, the policy satisfies both part (A) and part (B) of our definition. For part (A), the goal of this policy is clearly to change relative prices in the economy in order to make battery production a more attractive and bigger sector. For part (B), the policy is executed by a national-level authority (The Ministry of Industry and Information Technology).

In contrast, consider a Thai policy introduced in 2016, summarized in the following text:

“On 21 April 2016, the Thai Ministry of Commerce announced a reduction in rice-growing zones to stabilise rice prices for the upcoming harvest season. The available area for rice harvest has been reduced from a total of 61.7 million rai to 55.8 million rai (9 hectares). The amendment comes in the wake of an oversupply of rice, with affected areas reserved for other types of crops. In addition, the Ministry has encouraged rice mills to purchase paddy directly from the farmers in an effort to supplement the farmers’ income and incentivise them to stop from further growing rice. The amendment is in effect from May 2016.”

This description makes clear that the government is responding to a shock—an oversupply of rice—with the goal of stabilizing rice prices. The text also references the fact that the goal of this policy is to increase farmers’ incomes, or has social welfare dimensions. Taken together, we conclude that this description does not satisfy part (A) of our definition. In this instance, the policymaker is not motivated by a goal to shape the structure of the economy.

These examples clarify how our definition makes it possible for a human to distinguish industrial policy in our data. We now turn to the database of policy descriptions and the data we use for this study.

3 Data

To create measures of industrial policy practice, our methodology requires comprehensive textual data describing economic policies. We make use of the Global Trade

Alert (GTA) database, which contains detailed information on commercial policies, starting from 2008 to the present day. We use this database to classify industrial policy use at the country, industry, and year level. We now describe the GTA database, the core textual corpus for this project, and the economic data used for our descriptive analysis.

3.1 GLOBAL TRADE ALERTS AND POLICY TEXT CORPUS

The Global Trade Alert initiative is ambitious in scope and coverage. It deploys an international network of policy experts to identify state policy measures and credible announcements *that discriminate against foreign commercial interests* (Evenett and Fritz, 2020a).⁶ Thus, since its inception in 2008, the GTA project strives to capture an expansive set of measures, including textual data on these policies.

Given the remit of the GTA project, the database has become arguably the most comprehensive compilation of non-tariff measures available (Evenett, 2019). Their coverage rivals longer-run projects from multi-lateral institutions, such as the United Nations Conference on Trade and Development (UNCTAD) database and the World Trade Organization's (WTO) own surveillance projects. One key advantage of GTA is that it is independent, meaning it does not require compliance by reporting countries.

These data are well-suited to our application for three reasons. First, the policies covered by the GTA are a superset of industrial policies. Industrial policies will typically impact foreign commercial interests as, by definition, an industrial policy make some targeted activity relatively more attractive. Industrial policies fall squarely under the scope of GTA surveillance. In theory, any industrial policy *should* be included in the GTA. Of course, not all policies included in the GTA will be industrial policies. Second, the GTA provides English-language summaries of commercial policies. Third, the GTA strives for comprehensive, international coverage of policies.

Our version of the database (August 2020) contains approximately 28,000 observations over 175 countries. Basic descriptive statistics are reported in Table 1. For each observation, the GTA provides two types of data:

The first piece of data is critical: the GTA provides *English-language* policy summaries for each state act observation. Thus, each observation in the data contains paragraph descriptions summarizing policy, usually between 69-178 words in length. Though the GTA project is multilingual, all descriptions are in English. These textual descriptions of policy are what will be used for our classification process. In the

6. The GTA verifies measures and documents them through official statements of administrative institutions (Evenett and Fritz, 2020b, p. 1). Foreign commercial interests, for the GTA, include imports, exports, foreign investments, foreign properties, and foreign employees (Evenett, 2009a, p. 608).

language of text analysis, these textual summaries comprise our “corpus” and the policy variables comprise the “metadata.” Section 2.2, our methodology sections describe how we process this textual data for our classification pipeline.

The second piece of data are key policy variables, which we refer to as “meta-data.” Of these, we use following covariates: policy date; the type of intervention (e.g., a tariff, state loans, etc.); level of implementation; implementing jurisdiction; HS6 digit code of affected sectors; and whether there was firm-level scope tied to the intervention. The GTA has a multi-stage taxonomy for categorizing commercial policies into one of over 60 policy categories (from import bans to FDI incentives). The GTA’s taxonomy further disaggregates the UNCTAD’s Multi-Agency Support Team (MAST) code nomenclature for policy measures. One important caveat to note with respect to the current draft is that information for affected sectors (the HS6 code) is missing for 32% of the observations. We are currently predicting missing sectors codes, and this may affect results about the sectoral coverage of industrial policy.

Importantly, the GTA records credible policy changes (flows) as opposed to stocks of existing policies at a given point in time. This means we capture new industrial policies since 2009. Given that our aim is to analyze the current patterns of industrial policy-making around the globe flow data is relevant. However, it should not be confused with stocks, which we do not observe.

We use the data above to classify national industrial policy use at the country, sector, and year level. We use these binary classifications to create count-based indices or coverage ratios, which we then merge with economic data for analysis.

3.2 ECONOMIC DATA

We merge disaggregated measures of industrial policy use with two types of economic data. First, we merge our data with trade flow values from the United Nation’s COMTRADE. We use the BACI version of the UN COMTRADE dataset (Gaulier and Zignago, 2010), produced by the Centre d’Etudes Prospectives et d’Informations Internationales (CEPII). CEPII’s BACI database is a processed version of the UN COMTRADE data, which further cleans the original UN COMTRADE data and accounts for irregularities. We consider trade flows at the Harmonized System (HS) 6 and 2-digit levels; for the latter, we aggregate our IP index. Trade flow values are reported in USD.

In addition to bilateral trade flow data, we use of country-level economic statistics. We turn to Princeton’s World Economics and Politics Dataverse for cross-country economic statistics and compare industrial policy practices across the income distribution using GDP per capita in 2010 (at current USD).

4 Methodology

Our methodology uses the machine learning to classify instances of industrial policy at scale. Intuitively, we use our definition to automatically categorize policies as industrial policy or not. We classify observations in our database using English-language summaries of economic policies from the GTA database described above. Our approach to classification is “supervised.” We train our classification algorithm with data manually labeled using our formal definition of industrial policy. We then use our trained classification model to predict instances of industrial policy use in our database. This means classifying each of the approximately 28,000 policies in our database.

Our classification algorithm has four steps. Before detailing each part, we can summarize the workflow and core results:

Step 1 - Labelling Data With Formal Definition - Annotators hand-labeled approximately 2,000 policy descriptions (~ 7% of the data). In each case, the annotators determined whether the policy description satisfied the definition of industrial policy (specifically, part A). This process was based on a codebook that was produced by the research team.

Result - Our concept and definition of industrial policy are valid and reliable. We find humans agree on industrial policies.

Step 2 - Text Processing and Numerical Representation - We transformed policy descriptions (labeled and unlabelled) into numerical vectors and reduced the dimensions of the data. Text is pre-processed using a standard pre-processing workflow. For brevity, our study considers simple representation of text-as-data.

Step 3 - Training and Cross-Validation - We train a classification model using only policy descriptions as inputs. Specifically, we trained the models on one subset of the labeled observations and evaluate performance on a held-out sample. For exposition and simplicity, we focus on a logistic binary classifier using vectorized text.

Result i) - Our baseline models show strong performance in predicting industrial policy on unseen labeled data.

Result ii) - For our baseline logistic classifier, the language most predictive of industrial policy is intuitive. The model uses features of language associated with industrial policy, ex-ante.

Hyperparameters for our model are chosen using grid search and cross-validation.

Step 4 - Prediction We select the best-performing model in step three and use this model to predict industrial policy for the remaining unlabeled observations.

We now consider each step in detail.

4.1 STEP 1 - LABELLING DATA WITH FORMAL DEFINITION

First, we hand-coded policies that correspond formal definition to construct our training and testing data. To do so, we developed a codebook with instructions for determining whether a policy satisfies part A of our definition. Henceforth, we refer to satisfying part (A) of the definition as “being an industrial policy” (though in reality, here, we only focus on part A). The codebook is short, eight pages in total, with three of those pages dedicated to working examples. The definition was the focus of the coding process, and our aim was to minimize the guidance provided to annotators. We wanted the definition itself to be sufficient for categorizing policy descriptions.

In our codebook, annotators were asked to look for “the why” of the policy, or its goal. We note that many policies state their goals explicitly (e.g., “in order to boost domestic industry by making Egyptian cars more competitive”), or implicitly (e.g., “China’s ‘Major Technical Equipment’ policy grants tax-free imports to firms in certain sectors involved in the production of said equipment.”). Both types of evidence are sufficient to satisfy the criteria for classification. Our codebook also describes policies that are implicitly, industrial policy, see Section 2.

Using the codebook, a group of annotators (undergraduate and graduate student research assistants at Columbia University or the University of Oxford) hand-labeled 2,081 policies, or seven percent of the full dataset. The policy observations were randomly drawn from the full dataset, and stratified by measure type.

Each observation was independently annotated by four RAs and assigned one of three possible labels: “industrial policy,” “not industrial policy,” or “not enough information.” We assign the label “not enough information” in cases where the description did not contain sufficient information to determine whether part A of the definition was satisfied.

We used majority voting across annotator votes to assign a label to each observation. In cases with an even split across annotators, we labeled the observation as “nan”. Figure A.1 shows the initial breakdown of annotations. For 38.49% of the descriptions that were annotated, the goal of the policy could be determined (ie, the policy was

assigned an “industrial policy” or “not industrial policy” label). This is important. To the best of our knowledge, these summaries were not designed to capture the goal of the policymaker. Yet it turns out that in generic policy descriptions, this is information that is often stated. This is encouraging for the broader applicability of our methodological approach beyond this dataset. Moreover, many of our hand-labeled policies (23.11%) do indeed satisfy part (A) of our definition.

For simplicity, we collapse the three labels to a binary categorization. We assign each observation a label of “not industrial policy” (combining the first and third categories; “not enough information” and “not industrial policy”) or “industrial policy” (the second category, “industrial policy”). Thus, policies are either industrial policy or not.

In rare cases (179 observations), annotators were evenly split on categorization. There are numerous to deal with these (“nan”) edge cases. For our baseline workflow presented in this paper, these ties are dropped.⁷ In future versions of the paper, we will pursue alternative approaches to better incorporate these edge cases into the workflow.

The first annotation step of our workflow produces 2,081 labeled observations, which we use to train our binary classifier. Before turning to the details of this classification model, however, we first ask whether our definition of industrial policy is reliable.

Step 1 Result - Humans agree on IP

For hand-labeled data, it is important to demonstrate that annotations are “reliable” (Artstein and Poesio, 2008; Krippendorff, 2004). Ideally, we want different annotators to produce the same results. This would allow us to conclude that they have internalized a similar understanding of industrial policy based on our formal definition and codebook. Demonstrating reliability indicates that our definition of industrial policy is a valid theoretical concept.

We use two standard diagnostics to assess consistency across annotations, Krippendorff’s alpha and Conger’s Kappa, and refer to these measures as intercoder reliability.⁸ The measures for the four rounds of annotations in our workflow take

7. We have also experimented with another approach: we had an ‘expert’ (one author of the study, Reka) annotated these descriptions and definitively assigned them to a category. When these expert annotations are included, model performance degrades slightly. If we compare predictions with and without these observations—holding all other aspects of the pipeline constant—the two models disagree on less than five percent of the total 28,087 in the full dataset. That said, our view is that the majority of these annotations are true edge cases (as opposed to noisy labels), and “forcing” a label on them may not be the best approach. Thus, we currently exclude them in the baseline model.

8. Both metrics take values between 0 and 1, with 0 meaning perfect disagreement and 1 meaning perfect agreement. Both metrics are suited for instances of over two coders.

values between 0.64 and 0.84 with an increasing trend over subsequent rounds (suggestive of learning).

What do we make of these measures? For traditional content analysis applications of intercoder reliability measures, a Krippendorff's alpha between 0.67-0.8 is regarded as tolerable quality, and values above 0.8 ensure high-quality (Krippendorff, 2004). However, recently, scholars have shown this is not necessarily the case for machine learning applications where intercoder reliability measures can be misleading to the point that these thresholds should not be used (Reidsma and Carletta, 2008).⁹ Statistical models have been shown to correctly recover labels from noisy data (Artstein, 2017; Passonneau and Carpenter, 2014). In light of this, we use these metrics as general guidance and take our measures as reliable (particularly in annotation rounds 2-4, where Krippendorff's alpha is just below 0.8).

4.2 STEP 2 - TEXT PROCESSING AND NUMERICAL REPRESENTATION

In step two, we transform each textual policy description into a numerical array. In this section, we focus on a particularly simple way of representing text as numerical data. More complex representations of text deliver similar classification results.

We start by pre-processing our textual data and follow standard conventions. First, we remove common stop words, including conjunctions, articles, and prepositions.¹⁰ Stop words, while critical for proper sentence structure and grammar, did not carry critical information about whether a policy is industrial policy. Second, we remove all punctuation and numbers. Like stop words, we expect punctuation to be uninformative about our quantity of interest. We remove numbers for simplicity.¹¹ Last, we lemmatize words and replace words with their roots, which is also standard practice.

After pre-processing, we transform vectors of words into vectors of numbers using term frequency-inverse document frequency (*tf-idf*). Using *tf-idf* emphasizes words that are distinctly present or absent in each document. With this method, very uncommon words in the document and very common words across all measure descriptions have low *tf-idf* scores.

9. This is because if the source of disagreement is due to random noise, machine learning can tolerate data with lower agreement (Passonneau and Carpenter, 2014). However, if the disagreement is systematic, even reliability measures with values 0.80 and above will provide an unwanted pattern for the machine to detect (Reidsma and Carletta, 2008).

10. We follow conventional practice and use baseline stop word removal included in the spaCy Python library, an open-source package for natural language processing.

11. One may expect that numbers within policy descriptions likely do contain information about industrial policy. For instance, it might be the case that a particular date ("May 2017") is predictive of industrial policy.

We run versions of our models using unigrams (as in the example above) and unigrams plus bigrams, together. Bigrams use the methods described above and define tokens over two-word phrases. An advantage of bigrams, is that phrases may capture important information that is lost when considering the composite words in isolation, ignoring word order. A disadvantage is that phrases increase the vocabulary size and make computation more cumbersome.

4.3 STEP 3 - TRAINING, CROSS-VALIDATION, AND GRID-SEARCH

In step three, we create a mapping from a document to a prediction of industrial policy.¹² For clarity, we use a simple and interpretable model to create a proof of concept for our general approach. More complex models and textual representations can be used, but we illustrate our workflow using simple methods.

We use a simple binomial logistic regression, which has been shown to perform well in applications such as ours.¹³ We use standard L2 regularization and include the lambda parameter in the set of hyperparameters selected through grid search, described below.

In the simplest version of our model training, we randomly split our labeled observations into two subsets. Two-thirds of our labeled observations are assigned to the “training dataset” ($n = 1268$ observations total) and one-third are assigned to the “testing data” (635 observations total). We stratify on industrial policy in the training data and retain the original balance of industrial policy. Our two classes, however, are unbalanced in our training data, with only 25 percent of the observations in the labeled data are industrial policy. We deal with this imbalance by oversampling or duplicating industrial policy observations until we have a balanced sample. Figure A.2 illustrates this process. However, model performance is not sensitive to this procedure. For robustness, we used two alternative procedures, under-sampling, or dropping the dominant class, and running our model on unbalanced training data. Neither approach changed the results materially and we over-sample for our baseline approach.

We perform cross-validation on our model and select optimal model parameters using a grid search algorithm (GridSearchCV). The goal of cross-validation is to prevent overfitting, where we evaluate our model’s performance on different held-out

12. We use an industry-standard library in Python, scikit-learn. Unless otherwise specified, we use scikit-learn’s default options.

13. According to Gentzkow et al. (2017), “for simple text-regression tasks with input dimension on the same order as the sample size, penalized linear models typically perform close to the frontier in terms of out-of-sample prediction.”

sub-samples of the data.¹⁴ K-fold cross-validation divides the training data into k bins of equal size. Each iteration leaves out a different partition for testing, and the remaining $k - 1$ partitions are used for training the model.

We perform a grid search to select the set of hyperparameters that maximizes our preferred scoring metric. See Online Appendix for the full set of hyperparameters used in this process. Using grid search, we iterate through different combinations of hyperparameters. Then, using k-fold cross-validation, we evaluate the model's predictions and chose the set of hyperparameters that maximize our model's performance.

Step 2 Results - Our baseline models perform well

Our model performs well in classifying industrial policy. Table 2 summarizes the performance of our baseline model across common metrics. In terms of accuracy, the model correctly predicts 93% of policies. In terms of precision, 87% of policies identified as industrial policy by our model are, indeed, industrial policy; 96% of the policies identified as not industrial policy are, indeed, not industrial policy. In terms of model recall, 88% of the policies that are industrial policy the model identifies as such, and 95% of the policies that are not industrial policy the model identifies as such. See Table 2 for formal definitions. Table 2 also shows the F1 score, a weighted average of precision and recall, which provides a holistic measure of performance. The results above show that our basic model performs well in correctly classifying industrial policy from our textual database.

Step 2 Results - Our baseline models appear valid

We consider the validity of our model and consider the coefficients in our logistic classifier Gentzkow et al. (2017). Table 3 shows the largest coefficients from the baseline logistic regression, where positive coefficients are those that are most predictive of IP.¹⁵ The table contains words that we expect to see—such as “technology”, “green”, “development” and “export”—commonly associated with attempts to selectively shift economic activity.

4.4 STEP 4 - PREDICTION

In the last step of our pipeline, we take the model above and predict IP for the unlabeled observations in our data. Table 4 shows that 26% of policy observations

14. There are a number of different out-of-sample error terms we can use to evaluate our model. We use $F1$ score, a weighted average as precision and recall, as our preferred metric.

15. We follow the advice from Gentzkow et al. (2017, p. 27): “As in text regression, it is usually worthwhile to look at the largest coefficients for validation but not take the smaller values too seriously.”

in the full dataset satisfy part (A) of our definition of industrial policy. Column B in the same table shows that part (B) of our definition holds for 97% of the policies. Conditioning on both (A) and (B), 25% of the policies are industrial policies. These predictions are the basis of the next section on the contours of industrial policy.

Before we analyze this data, we perform simple validation exercises. Step three discussed a basic validation check based on the language elements most predictive of industrial policy. Here, we examine the extent to which the data we produce adheres to a set of expected patterns (known as ‘hypothesis validity’). One approach suggested in the literature is to test a hypothesis in the data that is so obvious, it would be difficult to explain if it were not true [Grimmer et al. \(2022, p. 32\)](#). To this end, we examine how our model labeled data fare in the face of the Covid-19 pandemic. Our hypothesis is that while many countries enacted a huge number of commercial policies to mitigate the effects of the pandemic, these were not typically industrial policies. Instead, they were crisis mitigation responses that had the aim of trying to alleviate various negative effects of the pandemic. That is, they violate part (A) of our definition.

Our hypothesis is that the pandemic should be evident, but only for the trend of non-industrial policies. In Figure 1, we plot the monthly policies implemented from January 2019 through July 2020 distinguishing those that our model classifies as industrial policy. Panel A plots the number of policies, while Panel B plots the shares. The results are reassuring. The effect of the pandemic is evident, *but only for the ‘not industrial policy’ category*. There is no similar movement for policies classified as industrial policy; if anything, they decline slightly, which would make sense if the state were using its scarce resources to fight the pandemic. We see a clear spike in ‘not IP’ interventions, globally, increasing from around 100 policies a month to over 350 in March, 2020. These policies then level off through the summer and quickly revert back to pre-pandemic levels. The shares tell a consistent story.

Finally, in terms of model validation, it is also important to note that some of the findings we discuss in the next section confirm widely held hypotheses about current industrial policy practice. Most notably, as we discuss below, our finding that industrial policy is on the rise validates many expert views. In addition, our finding that industrial policy is prevalent in major rich countries is consistent with other recent work ([DiPippo et al., 2022](#)). Taking all pieces of evidence together, we conclude that the model passes multiple, simple, intuitive validation tests. In short, our model seems to identify industrial policy well.

5 Descriptive Analysis: The Contours of Industrial Policy

We now take the model output and use it to examine key questions surrounding industrial policy use across the globe today.

Fact 1 - Industrial policy is common and it is on the rise.

A key finding of our study is that 25 percent of policies in the GTA data were classified as industrial policy by our model. This suggests that industrial policy is indeed quite common. According to our data, industrial policy is also on the rise. Figure 2 examines the trend for IP between 2009 and 2019.¹⁶ Panel A plots a clear upward trend for instances of industrial policy over our study period, expanding from 462 in 2009 to more than 1000 in 2018.¹⁷

Additionally, Panel B of Figure 2 also shows a marked upward shift in the *proportion* of policies classified as industrial policy, moving from only 20 percent of total policies in the early 2010s to nearly 50 percent in 2019. This suggests that the upward trend we measure in industrial policy is unlikely driven by improve coverage by the GTA over time. It also unlikely that we are simply picking up a general turn towards protective policy since 2008 (Evenett, 2009b). The results in Panel B validates, for the first time, the widely held hypothesis that industrial policy is currently on the rise (e.g. Stiglitz, Joseph E., Lin, Justin Yifu, Monga (2013); Cherif and Hasanov (2019)).

Fact 2 - Industrial policy is technocratic and it is granular.

What policy levers are used to conduct industrial policy? Industrial policies identified in our data tend to take the form of subsidies and export-related measures. These are the most prominent types of policy seen in Figure 3, which breaks down IP policies using UNCTAD's MAST code policy taxonomy (Panel A). An even more disaggregated breakdown of policy measures is possible using the GTA's in-house classification system (Panel B).¹⁸ The most prominent forms of industrial policy are (in order of importance) trade financing, state loans, financial grants, financial assistance in foreign markets, local sourcing, loan guarantees and import tariffs.

Furthermore, we find that industrial policies are remarkably firm-specific. By this we mean that they are aimed at *specific firms*. We know this, since this information is provided by our source database, which codes the extent to which policies are aimed

16. We drop observations from 2020 here as it is a partial year—we obtained the data in August, 2020.

17. We measure a slight decrease in 2019 in the total number of policies. We suspect that there may be some backfilling of policy announcements over time in the GTA. This is supported by the fact that the *share* of industrial policy continues to rise through 2019 (see below). The dataset will be updated for the next version of the draft, allowing us to better understand what is happening in 2019.

18. The GTA handbook provides a mapping between the MAST chapter codes and their own classification.

at particular firms. We find that over 60 percent of industrial policies are firm-specific, compared to only 20 percent for non-industrial policies. This is of course consistent with the policy measures that are being used above.

Taking these two findings together, the forms of industrial policy that we capture are often a far cry from the blunt import tariffs of decades past. Implementing many of the most common forms of IP will almost surely require high levels of fiscal and administrative capacity. That is, a state that has fiscal revenue to spend subsidizing firms and promoting exports; and a state that has the administrative capacity to identify which firms to support. These dimensions of state capacity provide important context for thinking through our stylized facts.

Fact 3: Industrial policy is unevenly used and skews heavily towards rich countries.

Though industrial policy is common, it is not evenly distributed across countries. Figure 4 plots the use of industrial policy, by quantile, for countries in our data. Strikingly, we pick up next to no industrial policy in the majority of countries. Instead, as the quantile plot makes clear, a handful of countries account for the lion's share of industrial policy.

It is instructive to examine the top countries that use industrial policy. Panel A in Figure 5 lists the top 20 countries engaged in industrial policy. Germany tops the list and dwarfs lower-ranked countries. It has about double the number of industrial policies as the countries that come next: Japan, Brazil and Canada. Notably, 12 of the countries in the top 20 list are in the richest income quintile, based on GDP per capita in 2010 (in current USD). This pattern suggests one potential correlate of what type of countries engage in IP: income.

We now systematically examine whether industrial policy use is correlated with income. We regress a country's (log) total number of industrial policies on a set of binary indicators denoting a country's income quintile.¹⁹ Panel A in Figure 6 plots the coefficients for each quintile. The excluded category is the poorest income quintile, so each coefficient measures the effect of a being in a particular income group relative to that income group. The pattern is striking: higher income quintiles are associated with more industrial policy. The difference is statistically significant for the third, fourth, and fifth income quintiles—middle to upper-income quintiles, respectively.

One important concern is that forms of reporting biases in the GTA data may correlate with income. First, it is plausible that higher income economies receive disproportionate attention from GTA enumerators. However, we have a direct way to account for this by controlling for the total number of policies that are observed (i.e.,

19. By splitting countries into income groups, we more formally examine where industrial policy is concentrated along the global income distribution.

the sum of policies classified as industrial policy and those classified as not industrial policy). Panel B of Figure 6 plots the coefficients when a control for the (log of) total policies is added. The findings are instructive. Conditional on the number of total policies measured in the GTA, the difference in industrial policy use between the poorest countries and middle-income countries disappears. The coefficients are close to zero, though the confidence intervals increase.

In contrast, the difference between industrial policy use at the top and bottom of the income distribution survives. Income quintiles four and five have statistically and economically more industrial policy, even conditional on the total number of observed policies (though point estimates shrink in magnitude across the board). There are two ways to interpret these findings. First, if one assumes there is no systematic reporting bias in the GTA, these results suggest that industrial policy skews towards rich countries even more than general commercial policies that affect foreign interests. However, if the GTA is subject to some reporting bias, our results suggest that this bias, in and of itself, cannot explain the entire correlation between IP and income. If measurement error accounts for our results, it must be the case that the GTA captures industrial policy differently, above and beyond other policies. This could be the case for example, if the language used to describe policies in poorer countries contains less information about the goals of the policy. We aim to tackle this issue in the next version of the draft.

Second, it is possible that high-income countries have more granular reporting standards. For example, perhaps richer countries release detailed data on which firms receive state subsidies, while poorer countries only release more aggregated data. In this case, we may systematically underestimate industrial policy because our measure counts the total number of policies. To tackle this, we calculate *coverage ratios*. Instead of using the total number of industrial policies in a country, we consider the *share of sectors* covered by industrial policy in a given country.

Figure 7 plots the coefficients for each income quintile using the same specifications as before and with coverage ratios (at Harmonized System 2-digit level) as our dependent variable. The results are qualitatively similar. Countries in higher income quintiles have substantially more HS 2-digit sectors where industrial policy is present. Once again, the differences are statistically significant for quintiles three, four and five. Moreover, the magnitudes seem large. Moving from the poorest to the middle quintile implies that about 10 more sectors are covered by at least one industrial policy.²⁰ Similarly, moving from the poorest to the richest income group implies that there are almost 20 more sectors covered by some form of industrial policy.

20. There are 97 2-digit HS codes, so a 10 percentage point increase in the coverage ratio implies about 10 more sectors covered.

Fact 4 - Industrial policy is sectorally selective and correlated with comparative advantage.

Another key question surrounding industrial policy is what sectors are targeted? All of them? Some of them? If so, are they systematically chosen in some way? Panel B in Figure 5 plots the top 20 users of IP once again, but now examining the *share* of HS 2-digit sectors covered by industrial policy, as opposed to the total number of policies.

For the 20 active users of IP, IP use is typically selective and targets a relatively small number of sectors. Recall that we observe hundreds of individual industrial policies for the top IP-using countries over our decade long sample. Yet, most countries target their IP towards a small set of sector codes. Some countries target only a handful of sectors, and most countries have industrial policy in fewer than 40 percent of HS 2 sectors. There are some important exceptions to this selectivity. Industrial policies in Brazil, Russia, and India tend cover a very large share of sectors (over 80 percent). Not far behind is the U.K., where about 70 percent of sectors are covered by industrial policy.

Interestingly, we find that higher and lower income countries tend to target similar sectors. Figures 8a and 8b present the top sectors 20 targeted by IP for the wealthier countries and poorer countries, respectively. The evidence in these plots suggests that, with the exception of greener interventions (richer countries) and textiles (poorer countries), there is overlap in targeted sectors such as heavy and hi-tech industry, across income levels.

Given industrial policy tends to be relatively selective, is there more we can say, systematically, about the pattern of targeting we observe? Indeed, we can. Table 5 shows that IP is systematically correlated with sectors in which countries have a higher revealed comparative advantage (Balassa, 1965). To examine this, we merge our industrial policy measures with trade data. Each observation is an industry-country-year tuple, where industry is defined at the HS 2-digit level or at the HS 6-digit level, depending on the specification. We say an observation is “covered” by industrial policy if we identify any industrial policy there in our dataset.

Table 5 examines the correlation between revealed comparative advantage and industrial policy at different levels of aggregation and with different fixed effects. We use the full variation in RCA (Panel A), but also construct a simpler, easier to interpret binary variable that is one if the value of RCA is greater than one—typically interpreted as a country having revealed comparative advantage (Panel B).

Across Table 5, the message is consistent: we find a positive correlation between comparative advantage and industrial policy. We see this in the simplest specification—without additional controls—at both the HS 2 and the HS 6-digit level. This pattern also holds with increasingly demanding fixed effects. At the 2-digit level, industrial policy is correlated with IP when using only variation within a country-year (column

2). That is, states systematically target policies towards sectors in which their RCA measure is relatively high. At the 6-digit level, we can say even more. Not only do states target sectors with relatively high RCA (column 4), they do so even *within* 2-digit HS industries (column 5, country by year by 2-digit industry effects).

For these specifications, it is particularly valuable that we measure industrial policy as flows. Doing so means that we examine the correlation between new policy with a sector's revealed comparative advantage for that given year. Our analysis then implies that new industrial policy is placed in sectors that have an established international presence.

6 Conclusion

This paper presents a new approach to measuring industrial policy and applies this to a comprehensive corpus of global commercial policies. The key methodological contribution of the paper is to propose a text-based approach to classifying tariff and non-tariff measures into those that are industrial policies and those that are not.

We show in our data that this has practical relevance. Our model, which does not use information on the type of policy measure, classifies only a subset of most policies as industrial policy. We thus demonstrate the very real need to move beyond taking entire classes of policy measures (e.g. tariffs, subsidies) and classifying them, *ex-ante*, as industrial policy.

We also show that getting measurement right has the potential to matter for how we think about industrial policy. Using our measurement approach, our descriptive results suggest that industrial policy, while common, skews very heavily towards rich countries. Importantly, it does so more than commercial policies in general.

The above result is important also for thinking about our findings that shows industrial policies are highly technocratic (using sophisticated policy levers) and often firm-specific. If this is indeed what modern industrial policy looks like, it will be much harder for the poorest countries to muster the fiscal and administrative capacity required to implement similar policies. Moreover, if major economies are indeed moving towards more industrial policy use (as our findings suggest), it will make it even more difficult for poor countries to compete in international markets.

Finally, our data also allow us to make inroads into one of the most controversial aspects of industrial policy: who receives it. We show that on average, the sectors that receive new industrial policies in a given year have systematically higher comparative advantage. Indeed, they are more likely to have established revealed comparative advantage (i.e., an RCA measure above one). This speaks to an age old debate about

the desirability and ability of a policymaker to pick winners. In reality, at least in our data, states support industries that have already established themselves in some way on international markets. While future versions of the paper will dig into this relationship further, we conclude this paper by noting that even these preliminary findings shed light on some of the key debates centered on industrial policy at present.

7 Figures

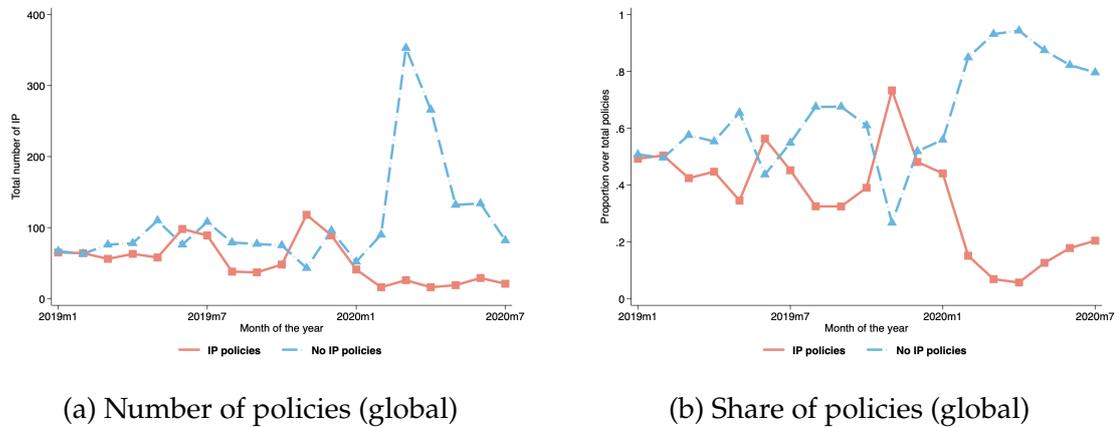


Figure 1: Validation Exercise - The increase in non-industrial policies v. industrial policy during the COVID-19 episode.

Notes: Number and share of industrial policies and non-industrial policies during the COVID-19 crisis. Panel (a) shows the count of policies through the period (Jan. 2019 - July 2020). Panel (b) shows the share of policies over the same period.

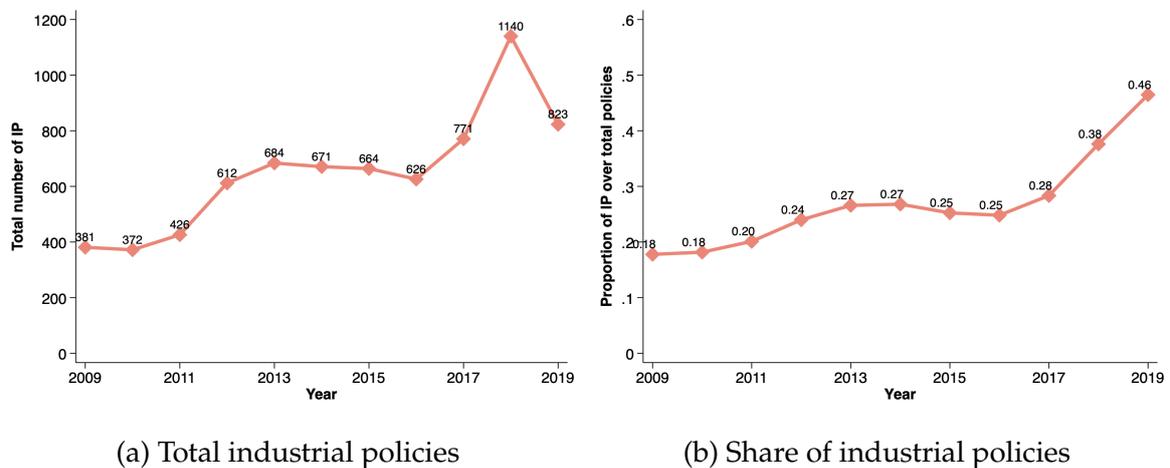
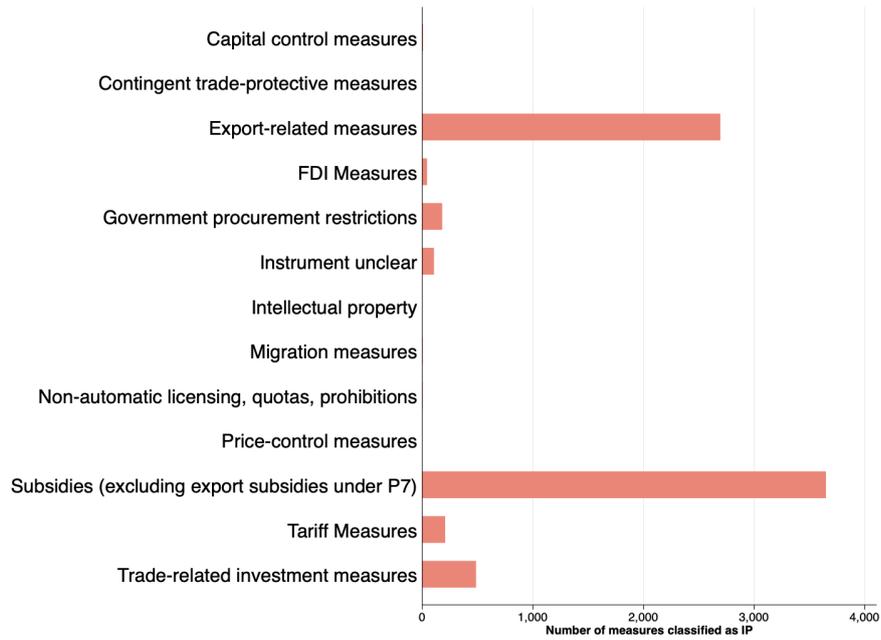
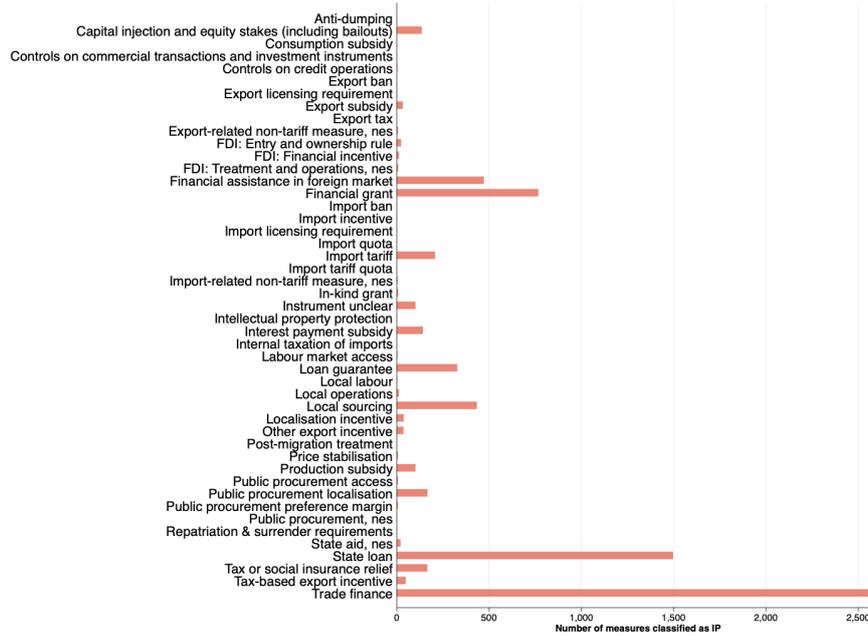


Figure 2: The time trend of industrial policy.

Notes: 2020 dropped as it is a partial year (we obtained the data in August, 2020). We also suspect 2019 may be incomplete because of backfilling in the data. This will be examined as we update the data going forward.



(a) MAST Chapter codes



(b) GTA policy measures

Figure 3: Count of industrial policy by measure type, shown by UN and GTA policy classifications.

Notes: Panel (a) reports counts of industrial policies by UN-MAST code classification. Panel (b) reports counts of industrial policies by the more disaggregated taxonomy used by the GTA project.

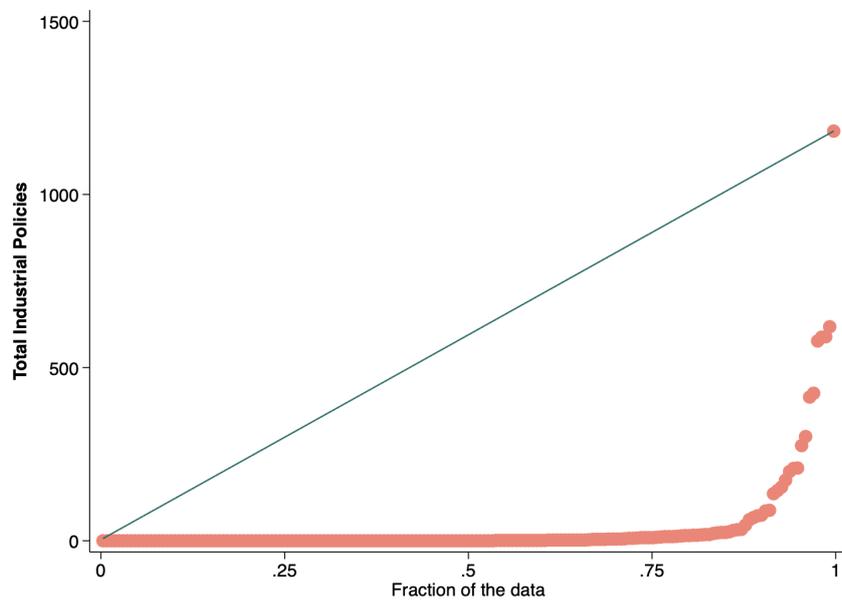
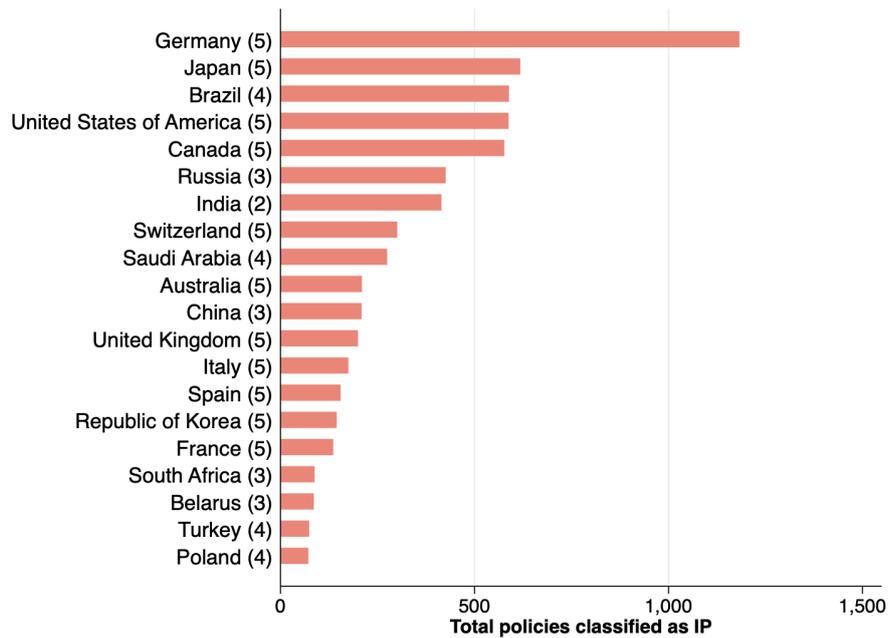
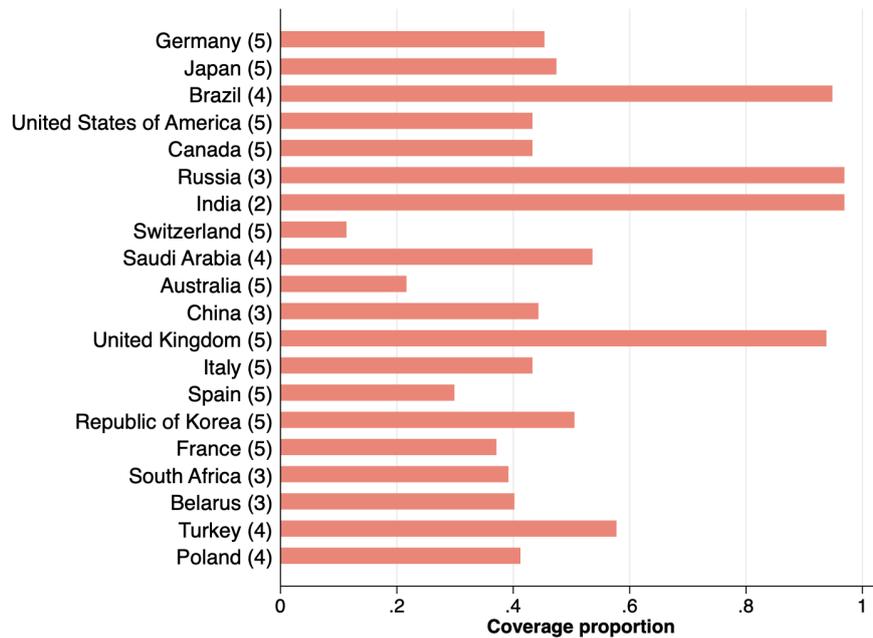


Figure 4: Quantile plot of IP by countries

Notes: A number countries account for many of the industrial policy observations in the data. Industrial policies are unevenly distributed in the data.



(a) Number of IP policies



(b) Share of sectors covered by any IP

Figure 5: Top 20 users of industrial policy and their coverage ratio

Notes: The number in parentheses refers to a country's position in the income distribution. 1 is the poorest and 5 is the richest income quintile based on 2010 GDP per capita. Countries in both panels are ordered according to their total count of industrial policies.

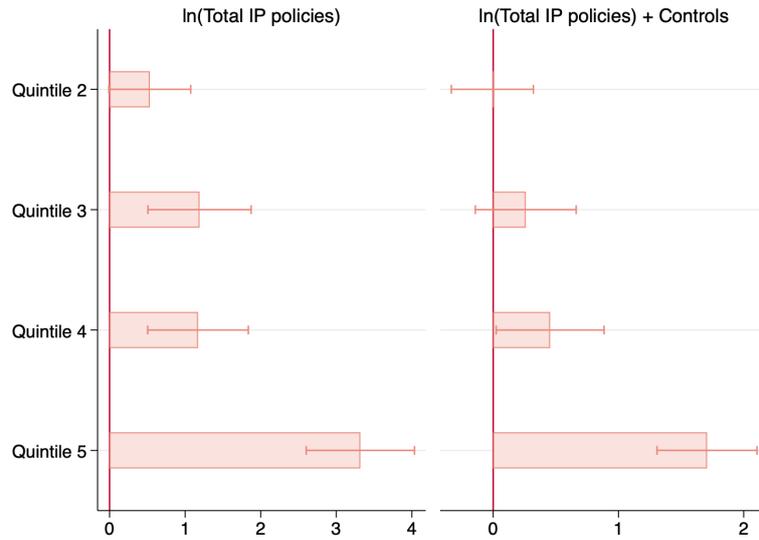


Figure 6: Regression of IP on income quintiles.

Notes: Controls include the (log) total number of policies. 95% confidence intervals. 2010 GDP per capita used to define income quintiles.

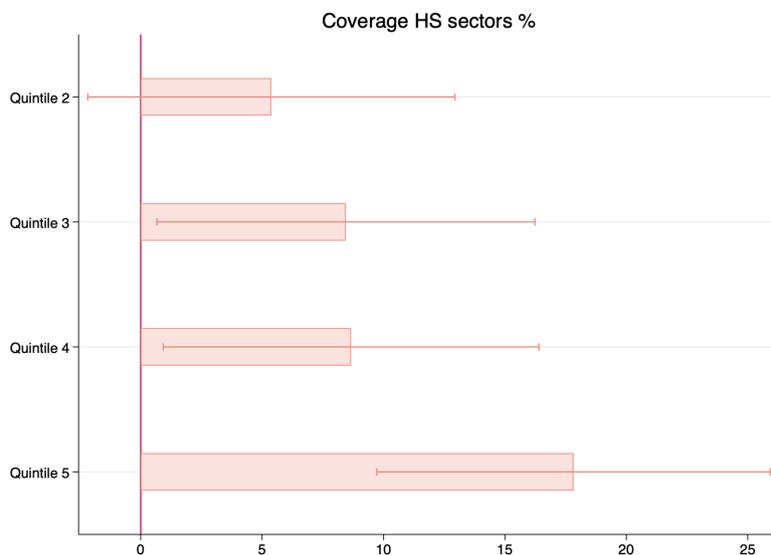
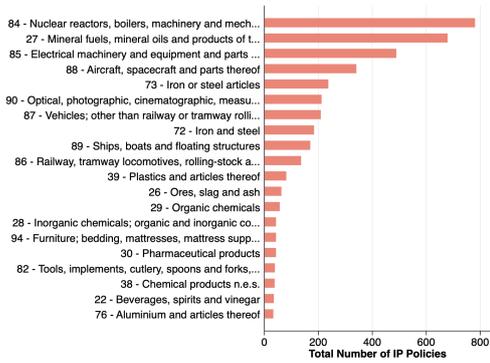
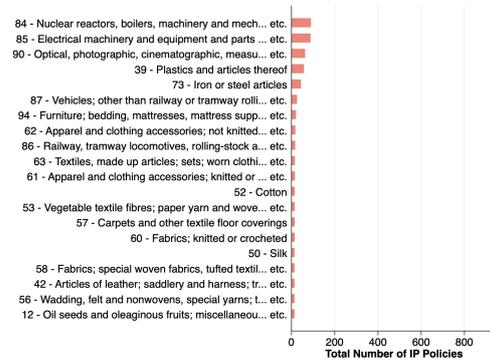


Figure 7: Regression of IP coverage ratios on income quintiles.

Notes: Coverage ratio defined as the number of HS 2 digit sectors that are covered by at least one industrial policy during the sample period. 95% confidence intervals. 2010 GDP per capita used to define income quintiles.



(a) Quintiles 4 and 5



(b) Quintiles 1 and 2

Figure 8: Top 20 sectors based on count of industrial policies

8 Tables

Table 1: Descriptive Statistics from the GTA data

	Mean	Std. Dev	25 pctl	75 pctl	N
Panel (A)					
Length Policy Description	142.18	114.22	69	178	28,333
Number HS - 6 digit affected	1.71	6.20	0	94	28,333
			N	percent	
Panel (B)					
	International Financial Institution	636	2.24%		
	National Financial Institution	4,447	15.70%		
	National	20,504	72.37%		
Implementation Level	Subnational	939	3.31%		
	Supranational	1,802	6.36%		
	Not Specified	5	0.02%		
	Total	28,333	100%		
Firm-Specific policies	Yes	10,156	64.15%		
	No	18,177	35.85%		
	Total	28,333	100%		
Mast Chapter Code	Capital control measures	174	0.61		
	Contingent trade-protective measures	2,761	9.74		
	Export-related measures	4,839	17.08		
	FDI Measures	811	2.86		
	Finance measures	26	0.09		
	Government procurement restrictions	902	3.18		
	Instrument unclear	537	1.9		
	Intellectual property	10	0.04		
	Migration measures	476	1.68		
	Non-automatic licensing, quotas, proh..	1,414	4.99		
	Price-control measures	301	1.06		
	Sanitary and phytosanitary measures	3	0.01		
	Subsidies (excluding export subsidies..	9,792	34.56		
	Tariff Measures	5,460	19.27		
	Technical barriers to trade	4	0.01		
Trade-related investment measures	823	2.9			
	Total	28,333	100%		

Notes: This table presents descriptive statistics for the 28,333 observations in the dataset. Note the missing data issue for the number of 6-digit HS product categories affected. Many (32%) of observations report no sector codes, which we are currently working on systematically predicting these.

Table 2: Performance metrics for baseline model.

	Precision	Recall	F1-Score
No IP Goal	0.96	0.95	0.96
IP Goal	0.87	0.88	0.87

Notes: Performance metrics, Precision, Recall and F1-Score, from baseline binomial logistic regression used predict the industry policy goals and non-industrial policy goals.

Table 3: Ten most predictive terms for baseline industrial policy classifier.

Feature Names	Coefficient Size
Tech	13.91
Green	13.69
Project	12.72
Export	10.89
Million	10.20
Plant	10.12
Lobster	9.84
Loan	9.48
Technology	9.41
Development	9.24

Notes: Coefficients come from baseline binary classifier for text-based logistic regression, and correspond to individual tokens. The text of these features are given in the left column.

Table 4: Results after extending the predictions to all of the GTA.

	Part (A) IP Goal	Part (B) Level of Government	IP: Part (A) + Part (B)
Yes	25.90%	3.31%	74.99%
No	74.10%	96.69%	25.01%

Notes: Precision, Recall and F1-Score resulting from using a binomial logistic regression to predict the industry policy goal.

Table 5: IP positively correlated with Revealed Comparative Advantage.

	2-digit HS		6-digit HS		
Panel (A) Dependent Variable: $\ln(\text{RCA} + 1)$					
	(1)	(2)	(3)	(4)	(5)
$IP = 1$	0.25*** (0.010)	0.20*** (0.018)	0.59*** (0.002)	0.20*** (0.027)	0.21*** (0.007)
Panel (B) Dependent Variable: Dummy $\text{RCA} = 1$ if $\text{RCA} > 1$					
	(1)	(2)	(3)	(4)	(5)
$IP = 1$	0.17*** (0.008)	0.13*** (0.013)	0.31*** (0.001)	0.11*** (0.016)	0.12*** (0.005)
Observations	186,725	186,725	8,783,577	8,783,577	8,783,577
Country-Year FE	No	Yes	No	Yes	No
Country-Year2-digit HS FE	No	No	No	No	Yes

Notes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Significance refers to robust standard errors (columns 1 and 3), standard errors clustered by country - year (columns 2 and 4), standard errors clustered by country - year - HS 2 code (column 5).

References

- Amsden, Alice H., *Asia's Next Giant: South Korea and Late Industrialization*, New York, New York: Oxford University Press, 1992.
- Amsden, Alice Hoffenberg, *The Rise of "The Rest": Challenges to the West From Late-Industrializing Economies*, New York: Oxford University Press, 2001.
- Anderson, JE and E Van Wincoop, "Trade costs," *Journal of Economic Literature*, 2004, 42 (3), 691–751.
- Artstein, Ron, "Inter-annotator Agreement," in Nancy Ide and James Pustejovsky, eds., *Handbook of Linguistic Annotation*, Dordrecht: Springer Netherlands, 2017, pp. 297–313.
- and Massimo Poesio, "Inter-Coder Agreement for Computational Linguistics," *Computational Linguistics*, December 2008, 34 (4), 555–596.
- Baker, Scott R, Nicholas Bloom, and Steven J Davis, "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, November 2016, 131 (4), 1593–1636.
- Balassa, Bela, "Trade Liberalisation and "Revealed" Comparative Advantage1," *The Manchester School*, May 1965, 33 (2), 99–123.
- , "Tariff policy and taxation in developing countries," 1989.
- Baldwin, Robert E., "The Case against Infant Industry Tariff Protection," *Journal of Political Economy*, 1969, 77 (3), 295–305.
- Bartelme, Dominick G, Arnaud Costinot, Dave Donaldson, and Andres Rodriguez-Clare, "The textbook case for industrial policy: Theory meets data," Technical Report 2019.
- Barwick, Panle Jia, Myrto Kalouptsi, and Nahim Bin Zahur, "China's industrial policy: An empirical evaluation," Technical Report 2019.
- Besley, Timothy and Torsten Persson, "State Capacity, Conflict, and Development," *Econometrica*, 2010, 78 (1), 1–34.
- Broda, Christian and David W Weinstein, "Variety Growth and World Welfare," *American Economic Review*, 2004, 94 (2), 139–144.
- Cagé, Julia and Lucie Gadenne, "Tax revenues and the fiscal cost of trade liberalization, 1792–2006," *Explorations in Economic History*, 2018, 70 (November 2017), 1–24.
- Chang, Ha-Joon, *Kicking away the ladder: development strategy in historical perspective*, London: Anthem Press, 2002.
- Cherif, Reda and Fuad Hasanov, "The Return of the Policy That Shall Not Be Named: Principles of Industrial Policy," 2019.
- Corden, W Max, *Trade Policy and Economic Welfare*, Oxford: Oxford University Press, 1997.
- Criscuolo, Chiara, Nicolas Gonne, Kohei Kitazawa, and Guy Lalanne, "An industrial policy framework for OECD countries: Old debates, new perspectives," 2022.

- DiPippo, Gerard, Ilaria Mazzocco, and Scott Kennedy, "Red Ink: Estimating Chinese Industrial Policy Spending in Comparative Perspective," 2022.
- D’Orazio, Vito, Steven T Landis, Glenn Palmer, and Philip Schrodtt, "Separating the Wheat from the Chaff: Applications of Automated Document Classification Using Support Vector Machines," *Political Analysis*, 2014, 22 (2), 224–242.
- Evans, Peter B, *Embedded Autonomy: States and Industrial Transformation*, Princeton, New Jersey: Princeton University Press, 1995.
- Evenett, Simon J, "Global Trade Alert: Motivation and Launch," *World Trade Review*, 2009, 8 (4), 607–609.
- , "Global Trade Alert: Motivation and Launch," *World Trade Review*, 2009, 8 (4), 607–609.
- , "Protectionism, state discrimination, and international business since the onset of the Global Financial Crisis," *Journal of International Business Policy*, 2019, 2 (1), 9–36.
- Evenett, Simon J. and Johannes Fritz, "The Global Trade Alert database handbook," Technical Report 2020.
- Evenett, Simon J and Johannes Fritz, "The Global Trade Alert database handbook," 2020.
- Gaulier, Guillaume and Soledad Zignago, "BACI: International Trade Database at the Product-Level. The 1994-2007 Version," 2010.
- Gennaioli, Nicola and Hans-Joachim Voth, "State Capacity and Military Conflict," *The Review of Economic Studies*, October 2015, 82 (4), 1409–1448.
- Gentzkow, Matthew, Bryan T Kelly, and Matt Taddy, "Text as Data," *National Bureau of Economic Research Working Paper Series*, 2017, No. 23276.
- Gerring, John, Haakon Gjerløw, and Carl Henrik Knutsen, "Regimes and industrialization," *World Development*, 2022, 152, 105791.
- Gerschenkron, Alexander, *Economic Backwardness in Historical Perspective: A Book of Essays*, Cambridge, Massachusetts: Belknap Press of Harvard University Press, 1962.
- Goldberg, P K and N Pavcnik, "Chapter 3 - The Effects of Trade Policy," in Kyle Bagwell and Robert Staiger, eds., *Handbook of Commercial Policy*, Vol. 1A, Amsterdam, The Netherlands: Elsevier B.V. and North-Holland, 2016, chapter 3, pp. 161–206.
- Grimmer, Justin, Margaret E Roberts, and Brandon M Stewart, *Text as data: A new framework for machine learning and the social sciences*, Princeton University Press, 2022.
- Haggard, Stephan, "The developmental state is dead: long live the developmental state!," in James Mahoney and Kathleen Thelen, eds., *Advances in Comparative-Historical Analysis, Strategies for Social Inquiry*, Cambridge: Cambridge University Press, 2015, pp. 39–66.
- Hanlon, W Walker, "The Persistent Effect of Temporary Input Cost Advantages in Shipbuilding, 1850 to 1911," *Journal of the European Economic Association*, December 2020, 18 (6), 3173–3209.

- Harrison, Ann and Andrés Rodríguez-Clare, "Trade, Foreign Investment, and Industrial Policy for Developing Countries," in Dani Rodrik and Mark Rosenzweig, eds., *Handbooks in Economics*, Vol. 5, Elsevier, 2010, chapter 63, pp. 4039–4214.
- Hassan, T A, S Hollander, L Van Lent, and A Tahoun, "Firm-level political risk: Measurement and effects," *Quarterly Journal of Economics*, 2019, 134 (4), 2135–2202.
- Hirschman, Albert O, *The Strategy of Economic Development*, third ed., New Haven, Connecticut: Yale University Press, 1958.
- Johnson, Chalmers, *MITI and the Japanese Miracle: The Growth of Industrial Policy : 1925-1975*, Stanford, California: Stanford University Press, 1982.
- Johnson, Harry G, "Optimum Welfare and Maximum Revenue Tariffs," *The Review of Economic Studies*, August 1951, 19 (1), 28–35.
- Juhász, Reka, "Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade," *American Economic Review*, 2018, 108 (11), 3339–3376.
- Kalouptsi, Myrto, "Detection and Impact of Industrial Subsidies: The Case of Chinese Shipbuilding," *The Review of Economic Studies*, April 2018, 85 (2), 1111–1158.
- Kohli, Atul, *State-Directed Development: Political Power and Industrialization in the Global Periphery*, Cambridge, U.K.: Cambridge University Press, 2004.
- Krippendorff, Klaus, *Content Analysis: An Introduction to Its Methodology*, first ed., Thousand Oaks: Sage Publications Inc, 2004.
- Krueger, Anne O., "Government Failures in Development," *The Journal of Economic Perspectives*, 1990, 4 (3), 9–23.
- Lane, Nathan, "The New Empirics of Industrial Policy," *Journal of Industry, Competition and Trade*, 2020, 1 (2), 1–26.
- , "Manufacturing Revolutions: Industrial Policy and Industrialization in South Korea," 2021.
- Lehmann, Sibylle H and Kevin H O'Rourke, "The Structure of Protection and Growth in the Late Nineteenth Century," *The Review of Economics and Statistics*, August 2011, 93 (2), 606–616.
- List, Frederick, *The National System of Political Economy*, Philadelphia: J. B. Lippincott & Company, 1856.
- Liu, Ernest, "Industrial Policies in Production Networks," *Quarterly Journal of Economics*, 2019, 134 (4), 1883–1948.
- Maloney, William F and Gaurav Nayyar, "Industrial Policy, Information, and Government Capacity," *The World Bank Research Observer*, August 2018, 33 (2), 189–217.
- Mitrunen, Matti, "War Reparations, Structural Change, and Intergenerational Mobility," 2019.

- Nelson, Laura K, Derek Burk, Marcel Knudsen, and Leslie McCall, "The Future of Coding: A Comparison of Hand-Coding and Three Types of Computer-Assisted Text Analysis Methods," *Sociological Methods & Research*, May 2018, 50 (1), 202–237.
- Pack, Howard, "Industrial Policy: Growth Elixir or Poison?," *The World Bank Research Observer*, 2000, 15 (1), 47–67.
- Passonneau, Rebecca J and Bob Carpenter, "The Benefits of a Model of Annotation," *Transactions of the Association for Computational Linguistics*, 2014, 2, 311–326.
- Reidsma, Dennis and Jean Carletta, "Reliability Measurement without Limits," *Computational Linguistics*, September 2008, 34 (3), 319–326.
- Rodriguez, Francisco and Dani Rodrik, "Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence," in "NBER Macroeconomics Annual 2000, Volume 15," Vol. 15 University of Chicago Press January 2001, pp. 261–338.
- Rodrik, Dani, "Why We Learn Nothing from Regressing Economic Growth on Policies," *Seoul Journal of Economics*, 2012, 25 (2), 137–151.
- Stiglitz, Joseph E., Lin, Justin Yifu, Monga, Célestin, "Introduction: The Rejuvenation of Industrial Policy," in Joseph E Stiglitz, Justin Yifu Lin, and Celestin Monga, eds., *The Industrial Policy Revolution I: The Role of Government Beyond Ideology*, New York, New York and Hampshire, UK: Palgrave Macmillan, 2013, pp. 1–18.
- Taussig, Frank William, *The Tariff History of the United States*, New York: GP Putnam's Sons, 1914.
- United States International Trade Commission, "FOREIGN INDUSTRIAL TARGETING AND ITS EFFECTS ON U.S. INDUSTRIES PHASE 1: JAPAN.," Technical Report, Washington, D.C. 1983.
- Wade, Robert H., *Governing the market: Economic theory and the role of government in East Asian industrialization*, second ed., Princeton, New Jersey: Princeton University Press, 1990.
- , "Rethinking industrial policy for low income countries," *African Development Review*, 2009, 21 (2), 352–366.
- Warwick, Ken, "Beyond Industrial Policy: Emerging Issues and New Trends," Technical Report, Paris 2013.

Online Appendix

The Who, What When and How of Industrial Policy: A Text-Based Approach

Réka Juhász Nathan Lane Emily Oehlsen Verónica C. Pérez

A Figures

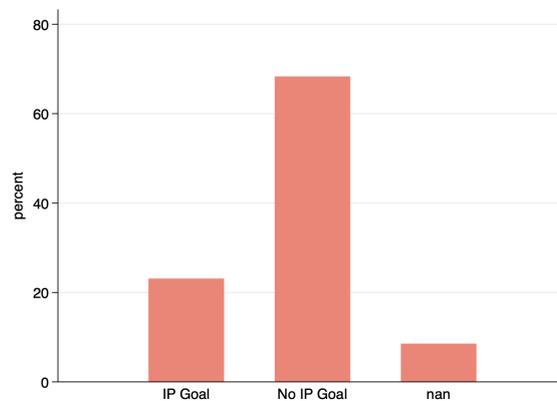


Figure A.1: Labels from hand coders



Figure A.2: Dealing with imbalance