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

Réka Juhász (UBC, NBER & CEPR) Nathan Lane (Oxford) Emily Oehlsen (Oxford) Verónica C. Pérez (Boston University)

December 2023

Abstract

Since the 18th century, policymakers have debated the merits of industrial policy (IP). Yet, economists lack basic facts about its use. This study sheds light on industrial policy by measuring and studying global policy practice for the first time. We first create an automated classification algorithm for categorizing industrial policy practice from text. We then apply it to a global database of commercial policy descriptions and quantify policy use at the country, industry, and year levels (2009-2020). These data allow us to study fundamental policy patterns across the world. 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. We show that industrial policy is a prominent feature of the global economy and a far cry from industrial policies of 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, Penny Goldberg, Douglas Gollin, Amy Handlan, Gordon Hanson, Tarek Hassan, Giammario Impullitti, Myrto Kaloupt-sidi, Guy Lalanne, Colin McAuliffe, Phillip McCalman, Suresh Naidu, Ashwin Narayan, Gaurav Nayyar, Nina Pavcnik, Tristan Reed, Dani Rodrik, Francesco Squintani, Chris Woodruff, and seminar audiences at American University, Berkeley, Brown, Chapman University, Columbia, Council of Economic Advisors (U.S.), Duke University, Econometric Society Summer Meetings, ETH Zurich, Harvard, the IGC-Stanford Firms, Development and Trade conference, Michigan, MIT Sloan, National University of Singapore, NBER Summer Institute, New Thinking in Industrial, Innovation and Technology Policy Conference at Columbia, Pakistan Institute of Development Economics, STEG-CEPR-IGC conference, UC-Irvine, University of Melbourne, Warwick, and Yale 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, Mikhael Gaster, Max Marczinek, Esha Vaze, Sarah Wappel, Megan Yamoah, and a team of Columbia and Oxford annotators provided outstanding research support. We gratefully acknowledge funding from ISERP, STEG/CEPR, and the Alfred P. Sloan Foundation.

1 Introduction

Across the social sciences, fundamental questions revolve around the state's role in shaping our economic world. The use of state action to transform the composition of economic activity, or industrial policy, has been controversial since the rise of modern capitalism (List, 1856; Taussig, 1914; Chang, 2002).[1] Despite the longevity of these debates, empirical work on industrial policy is notably incomplete (Harrison and Rodríguez-Clare, 2010). This quantitative deficit spans not only questions around efficacy but, importantly, core questions around practice. We lack measures and knowledge surrounding practice. This gap is particularly stark against the march of large new industrial policy programs in China, the United States, Europe, and beyond.

In this paper, we shed light on global industrial policy practice. First, we propose a new text-based approach to measuring industrial policy at scale. Our method uses machine learning to automatically classify industrial policies from policy text. We develop and apply a classification algorithm to a global policy database to quantify new industrial policy activity at the country, industry, and year levels.[2] Second, we use this data to provide new stylized facts about global policy action and highlight its importance. Our descriptive analysis of industrial policy provides the basic contours of policy practice from 2009–2020. Our paper can thus be seen as a first attempt at systematically measuring and capturing global industrial policymaking. We organize this study in two parts.

The first part of the paper presents our text-based approach to measuring industrial policy activity. Measuring industrial policy is difficult and data on practice does not exist. The industrial policy toolbox includes many types of instruments, or levers, many of which are not systematically recorded (see Goldberg and Pavcnik (2016)). Where data are collected, researchers face a fundamental issue of observability. Consider tariffs, a classic industrial policy lever. Although historically used to protect infant industries, tariffs also have a multitude of uses, from raising government revenue (Johnson, 1951; Balassa, 1989; Cagé and Gadenne, 2018) to managing terms

1. Thinkers have long emphasized the importance of industrial policy in economic transformation (List, 1856; Wade, 1990; Amsden, 1992; Evans, 1995). Others have argued that industrial policies are pernicious (Krueger, 1990; Pack, 2000), exemplified by Gary Becker's quip "The best industrial policy is none at all" (Business Week, 26 August 1985).
2. These flow-based measures of industrial policy are similar to the coverage ratios used in studies of non-tariff measures. Such count-based measures or coverage ratios are inputs into quantitative trade and gravity models used to calculate indices of protection and *ad valorem* equivalencies.

of trade (Broda and Weinstein, 2004).[3] Thus, a data set of tariffs is not a data set of industrial policies. The same holds for non-tariff measures (NTMs), which are far more common today. However, unlike tariffs, data on NTMs are notoriously patchy and endemically under-measured.

These issues of measurement mean we lack a picture of global industrial policy practice. This presents a bottleneck for policymakers and social scientists studying questions surrounding real-world activity.[4] These deficits present a bottleneck to understanding the importance and shape of interventions. Likewise, imperfect proxies complicate more normative studies of efficacy. The practice of equating a policy measure (*e.g.*, *tariffs*, *ex-ante*, with industrial policy, has made cross-country work difficult to interpret (Rodriguez and Rodrik, 2001; Harrison and Rodríguez-Clare, 2010; Rodrik, 2012; Lane, 2020). Studying industrial policy activity requires more complete data, and going beyond the data that do exist.

To measure industrial policy activity we turn to policy text. We record instances of policy activity by considering whether policy descriptions convey industrial policy goals. Such objectives often imbue the way political actors describe policy. Consider the following description of a Chinese subsidy from our data:

> "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. [...]"

The goal, italicized, is stated in the opening sentence. The policymaker seeks to shape the composition of the economy by boosting a particular sector (vehicle batteries). This example is textbook industrial policy. Such language is not anomalous and we show that policy descriptions often express these objectives. Our method uses this information to record industrial policies, and, importantly, to distinguish policy measures used for industrial policy or those used for other goals (*e.g.*, public health interventions or welfare policy).

We use machine learning to create global measures of industrial policy practice. We develop a classification algorithm and apply it to a rich database of commercial policy descriptions. Our algorithm classifies industrial policy activity using supervised machine learning and employs a transparent logistic-based classifier. Our textual data comes from the most comprehensive, international corpus of commercial policy, *The Global Trade Alert* (GTA) database (Evenett, 2019). GTA tracks and collects

3. Low fiscal capacity states and developing countries have historically relied on tariffs as a source of revenue (Balassa, 1989). The poorest and least fiscally developed countries may be the same countries most likely to use tariffs for revenue (Cagé and Gadenne, 2018).
4. Our work, such as (Juhasz, 2018), illustrates that scholars can use shocks that *mimic* industrial policy to test mechanisms and evaluate aspects of industrial policy.

new government policy announcements (see Section 3), with English-language descriptions of policies. By applying our classifier to this data, we are able to quantify new industrial policy activity at the country, sector, and year levels.

Our classification process can be summarized in four steps. First, we develop a simple, broad *ex-ante* definition of industrial policy. 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 our classification model and assess its performance. Third, we use this classifier to automatically categorize, or predict, the remaining policies in our database (~28,000 observations). In effect, this supervised classification algorithm replicates manual policy categorization to classify policies at scale.

Our classification process delivers three important results. First, annotators successfully use our formal definition to categorize policies and agree on what industrial policy is. Second, we test our classifiers on labeled data not seen ("held out") by the model during training and show it performs well across performance metrics (accuracy, recall, precision, etc.). Third, our classification process passes multiple validation exercises. We consider the 'face validity' (Grimmer, Roberts and Stewart, 2022, p.31) for our baseline logistic classifier and show that the words ("tokens") most predictive of industrial policy are reasonable.[5] Predictive tokens include 'technology', 'development', and 'green'. Likewise, we consider, 'hypothesis validity' and show that our findings correlate to anticipated, real-world patterns (ibid), such as the COVID pandemic and recent policy episodes [6]

In the second part of our paper, we use our data to explore fundamental patterns of global industrial policy use. Our analysis delivers four key findings. First, we show that new industrial policy activity is important (25% of commercial policies in our database) and on the rise (doubling through the 2010s). New industrial policy activity is a growing part of the post-Global Financial Crisis world.

The second key result is that contemporary industrial policy practice is technocratic and outward-oriented. The majority of industrial policies identified in our data take the form of subsidies and export-related measures. They are also granular, with many measures targeted at individual firms. These forms of policies, unlike tariffs, demand

5. Face validity requires that, at a minimum, the measures pass inspection by a domain expert.
6. 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. 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)) on the proliferation of industrial policy.

fiscal resources and state capacity. These patterns suggest that current industrial policy practice differs substantially from the blunt import tariffs used in the past.

The third key result considers who uses industrial policies. We find that high-income countries are those implementing many new industrial policies. Our analysis shows a robust positive correlation between income and new industrial policy activity. This pattern is robust to controlling for different reporting biases and to alternative measures of industry policy use. The wide use of industrial policies among higher-income countries is striking, given industrial policy's emphasis within development economics. Likewise, autocracies and late-industrializers have long been associated with industrial policy (Gerring, Gjerløw and Knutsen, 2022). Although we record industrial policy activity among many non-democracies (*e.g.,* China, Russia, and Saudi Arabia), a great deal of activity is seen across liberal democracies. Industrial policy is not only the purview of developmental autocracies but is widely practiced among OECD members, which dovetails with recent accounting work by Criscuolo, Gonne, Kitazawa and Lalanne (2022) for select member states.

The fourth key result is that countries are selective in the sectors they target. Industrial policy activity is directed at a specific subset of sectors. 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 with an established international presence.

2 Defining Industrial Policy

Let's first be clear about concepts. We refer to industrial policy as state action directed at changing the structure of economic activity within an 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).[7] In other words, the intentions and goals of industrial policy matter, especially for recording policy activity and practice.

We apply a formal *ex-ante* definition to record industrial policy activity from data. This formal definition allows us to categorize descriptions of industrial policy

7. This conception of industrial policy is useful in that it encompasses different policy views: *e.g.,* those of neoclassical economists and those of scholars of the developmental state. Such views see industrial policy as state action executed in the service of broader goals. See Appendix A.

activity. Our 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) and align with contemporary efforts to clearly define policy (DiPippo et al., 2022; Criscuolo et al., 2022). We use the following formal definition:

Formal Definition for Recording Industrial Policy Activity

> 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), to shift the long-run composition of economic activity.

Thus, a policy description is identified as industrial policy activity if it satisfies the definition above; specifically, if it has an industrial policy objective and, for this study, is implemented at the national level (at minimum), respectively. Our definition captures the fact that industrial policy is goal-oriented, with a vision of shaping forms of economic activity within an economy. By pursuing these goals, the state takes a stand on its desired structure of the economy and plays an active role in achieving it.[8] Such goals may be aimed at correcting a market failure (neoclassical), achieving structural transformation (developmentalist), geopolitical power (neo-mercantalist), and beyond.

Industrial policy has *specificity* in how it shapes the economy and does so along different dimensions. These dimensions can include more sectoral, or "vertical," industrial policies (*e.g.* infant industry policy aimed at steel or textiles). Alternatively, they include forms of industrial policies that impact specific forms of economic *activity* (*e.g.*, exporting, innovation, or de-carbonization) rather than sectors. In short, our definition is agnostic as to whether a policy is sectoral or not.

For this study, we focus on recording *national-level* economic activity. We consider primarily the national-level economy (*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.[9] Thus, we exclude sub-national policies. Doing so excludes forms of place-based policies implemented by sub-national

8. In emphasizing the goals of policy, our definition encompasses different conceptions of industrial policy.
9. We follow the literature and consider the state as an aggregation beyond merely the government or the bureaucracy.

administrations that are disconnected from national economic goals. This is a simplifying assumption meant to highlight national-level comparisons.[10]

Thus, our *ex-ante* definition flexible accommodates many forms of policy making. Our definition is agnostic about which measures are used to perform industrial policy. Instead, our definition emphasizes the goals, rather than the means to achieving these goals. We view this as a key strength of our definition and more generally, our approach to measurement. This coincides with approaches by Criscuolo et al. (2022), who independently take a similarly broad, *ex-ante* view of policy. Online Appendix A.1 describes common types of policies that are included and excluded by our definition. The implications of our definition are standard and comport with the definitions and current practices of accounting studies of industrial policy (DiPippo et al., 2022; Criscuolo et al., 2022).

3 Data

Measuring industrial policy activity from text requires systematic textual data describing economic policies. We make use of the Global Trade Alert (GTA) database, which contains detailed information on international commercial policies, starting from 2008 to the present day. This section describes 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

Our study uses data from The Global Trade Alert project. The GTA initiative is ambitious in scope and coverage and deploys an international network of policy experts to identify state policy measures and *credible* announcements *that discriminate against foreign commercial interests* (Evenett and Fritz, 2020).[11] Since its inception in 2008, the GTA project has strived 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

10. Another rationale is that sub-national interests and objectives may conflict with national interests or reflect competition among sub-national units. In this case, policies may not reflect the objectives of shaping *national* economic activity.

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

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. The remit of GTA's surveillance effort covers a wide range of tariff and non-tariff measures. By capturing instruments that "discriminate against foreign commercial interests," it takes a broad approach, capturing only those instruments seen as classically "protectionist." In doing so, its approach is agnostic as to the form of commerce impacted by acts (*e.g.* imports) and the instrument used (*e.g.* tariffs) (see discussion in Evenett (2019)). Thus, GTA's wide scope covers a multitude of levers not traditionally captured by organizations like the WTO.

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 makes some targeted activity relatively more attractive. Industrial policy activity will 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 language of text analysis, these textual summaries comprise our "corpus" and the policy variables comprise the "metadata." 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 the 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 sector 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 have captured 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 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 machine learning to classify instances of industrial policy at scale. Intuitively, we use our definition (Section 2) to automatically record, or classify, instances of industrial policy activity. We classify observations in our database using English-language summaries of economic policies from the GTA database described above (Section 3). Our approach to classification is "supervised." We train our classification algorithm with data manually labeled using our formal definition and then apply it to our entire database to predict instances of industrial policy use.

Our classification algorithm can be described in four steps. Before detailing each, we summarize each step and key result.

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. This process was based on a simple codebook using our formal definition.

Result - Our definition of industrial policy are valid and reliable. We find that human annotators 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 workflow.

Step 3 - Training and Cross-Validation - We train a classification model using only policy descriptions as inputs. Specifically, we trained the models on a subset of the labeled observations and then evaluated performance on a held-out sample of labeled data. Hyperparameters for our model are chosen using grid search and cross-validation. For exposition and simplicity, we focus on a logistic binary classifier using vectorized text.

Result i) - Our baseline models perform well 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*.

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

We first hand-coded policies that correspond to formal definitions to construct our training and testing data. To do so, we developed a codebook with instructions for determining whether a policy satisfies our definition. 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 we aimed 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 data set. 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 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 (*i.e.,* the policy was assigned an "industrial policy" or "not industrial policy" label). This is important, as summaries are not designed to capture the goal of the policymaker. Yet, this information is often conveyed in general policy descriptions. 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 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 recorded as industrial policy activity 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, these ties are dropped.[12] In future versions of the paper, we will pursue alternative approaches to better incorporate edge cases into our 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 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.[13] The measures for the four rounds of annotations in our workflow take 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).[14] 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).

12. We have also experimented with another approach. We had an 'expert' (one author of the study, Reka) annotate these descriptions and definitively assign 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.
13. 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.
14. 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).

4.2 STEP 2 - TEXT PROCESSING AND NUMERICAL REPRESENTATION

In step two, we transform each policy description text into a numerical array. For clarity, this section focuses on a 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.[15] Stop words, although critical for proper sentence structure and grammar, do not carry critical information about whether an act is an industrial policy activity. 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.[16] Last, following standard practice, we lemmatize words and replace them with their roots.[17]

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.

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.[18] 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.

15. We follow conventional practice and use baseline stop word removal included in the spaCy Python library, an open-source package for natural language processing.
16. 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.
17. See Online Appendix B for an example of pre-processing a sample policy description.
18. We use an industry-standard library in Python, scikit-learn. Unless otherwise specified, we use scikit-learn's default options.

We use a simple binomial logistic regression, which has been shown to perform well in applications such as ours.[19] 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 being 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 sub-samples of the data.[20] 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 the 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 choose 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

19. According to Gentzkow, Kelly and Taddy (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."

20. 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.

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.[21] 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 in the full dataset satisfy our definition of industrial policy. Column B in the same table shows that 97% of the policies are national-level. Conditioning on nation-level policy and our definition, 25% of the policies are industrial policy activity. 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.

21. 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."

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.[22] 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.[23]

22. We drop observations from 2020 here as it is a partial year—we obtained the data in August 2020.
23. 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.

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 is also unlikely that we are simply picking up a general turn towards protective policy since 2008 (Evenett, 2009). 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).[24] 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 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,

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

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.[25] 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., 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

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

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.[26] 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.

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 an 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 to 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.

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

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 sheds light on an essential class of economic interventions: industrial policy. To do so, we provide a new approach to measuring industrial policy practice and describe the emerging countours of practice, globally, since the Global Financial Crisis.

We provide a text-based approach to recording new policy flows. We develop a classification algorithm for classifying industrial policy at scale, and then apply it to a comprehensive corpus of global commercial policies. We then use this data to explore the important dimensions of industrial policy use globally.

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 show 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

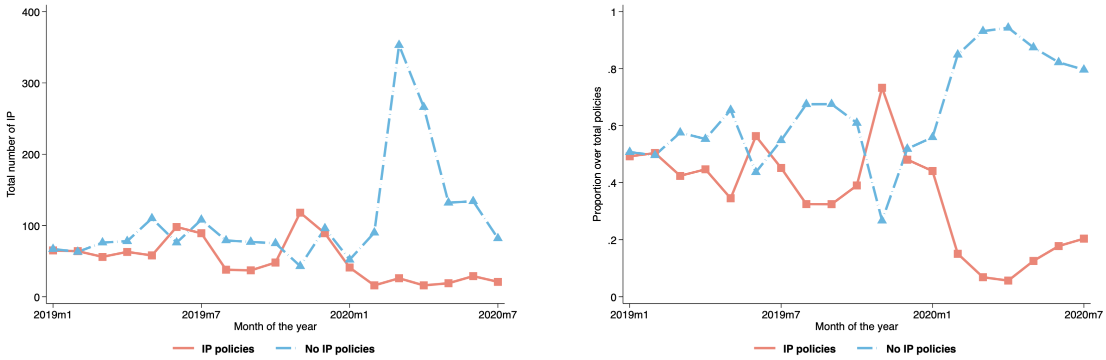

(a) Number of policies (global)

(b) Share of policies (global)

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 if policies through the period (Jan. 2019 - July 2020). Panel (b) shows the share of policies over the same period.

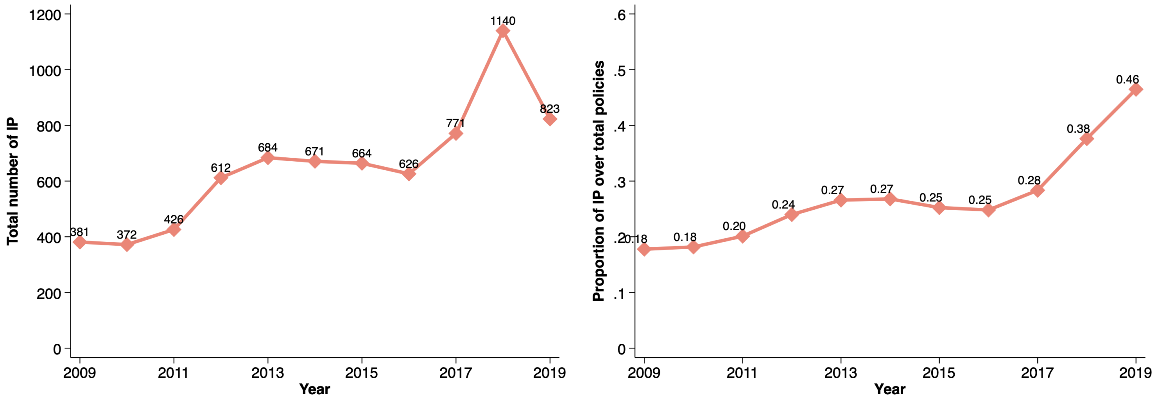

(a) Total industrial policies

(b) Share of industrial policies

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

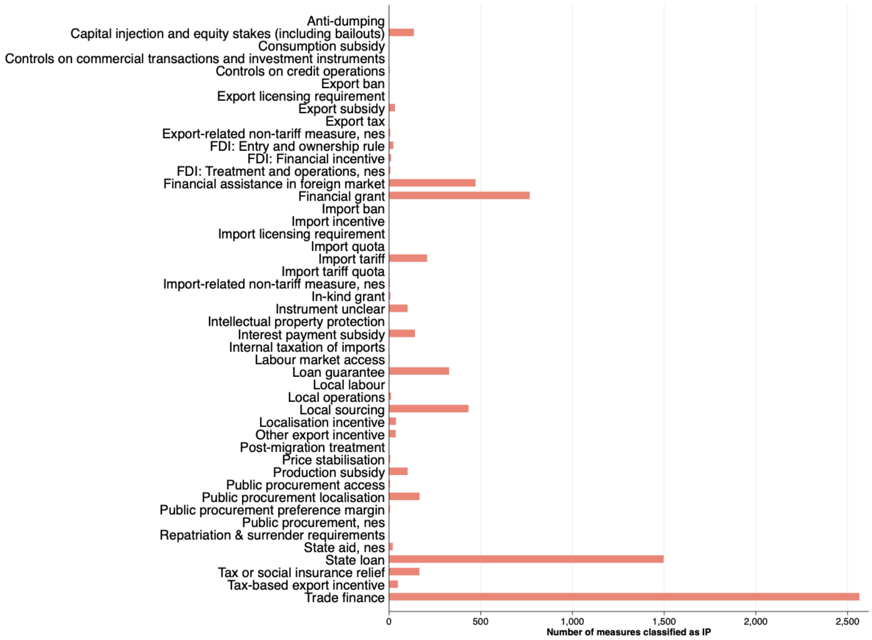

(b) GTA policy measures

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

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

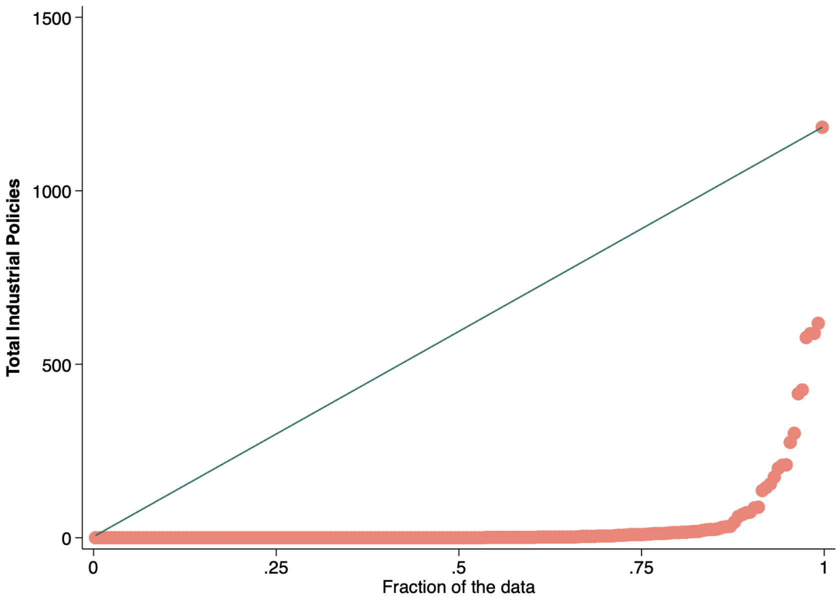

Figure 4: Quantile plot of IP by countries

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

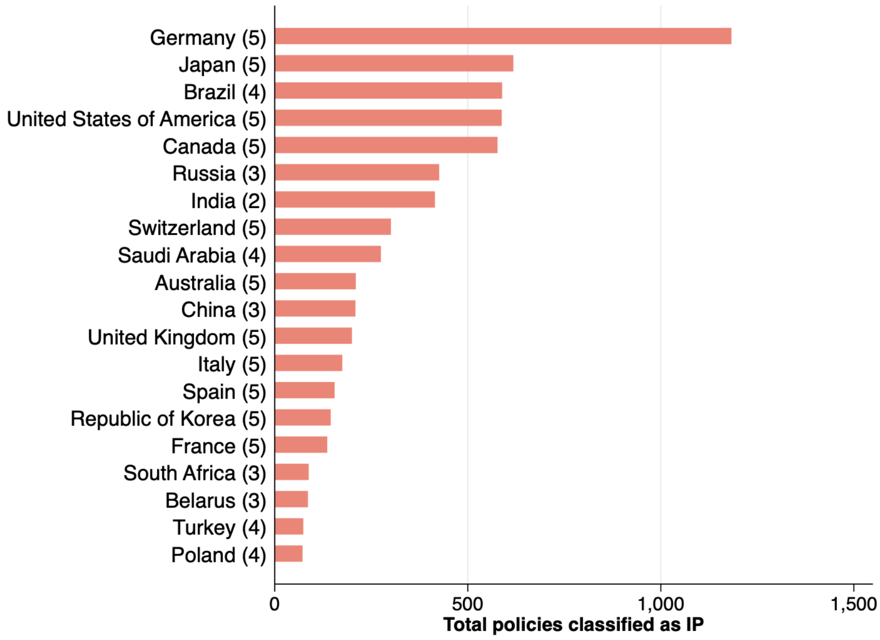

(a) Number of IP policies

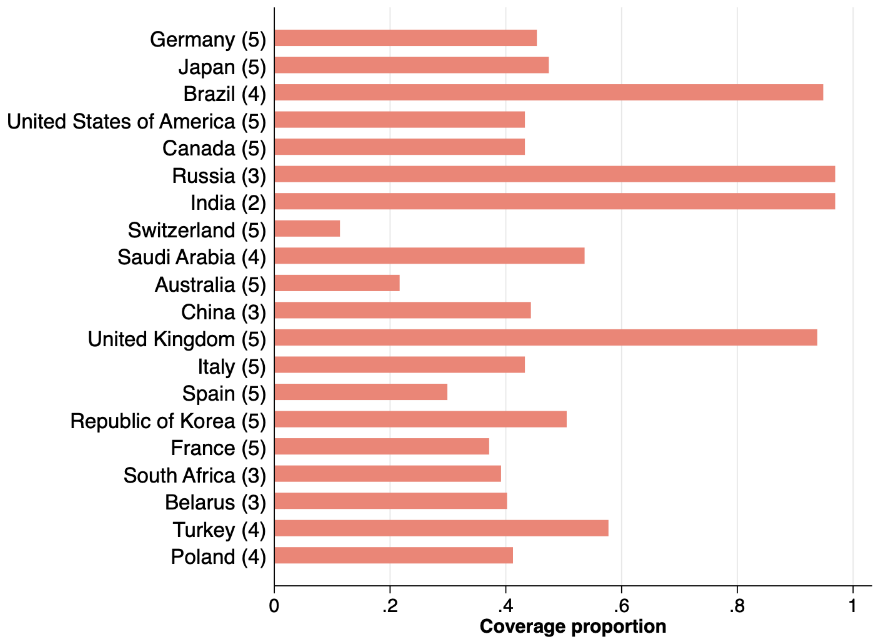

(b) Share of sectors covered by any IP

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.

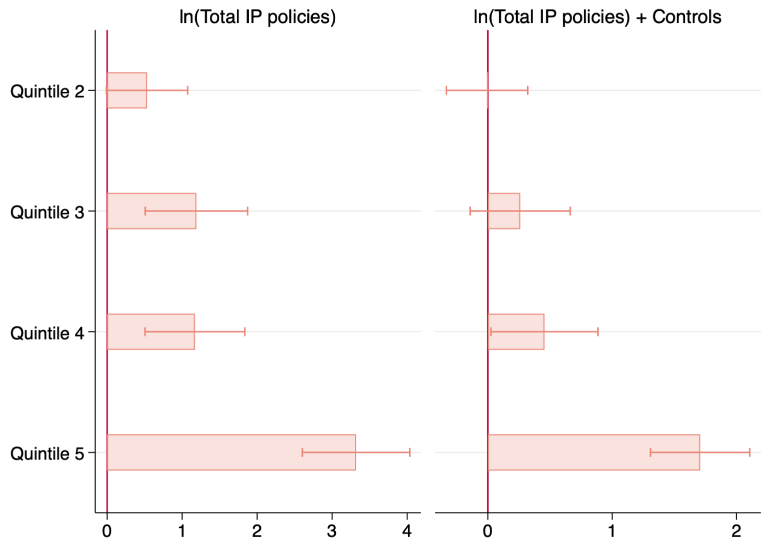

Figure 6: Regression of IP on income quintiles.

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

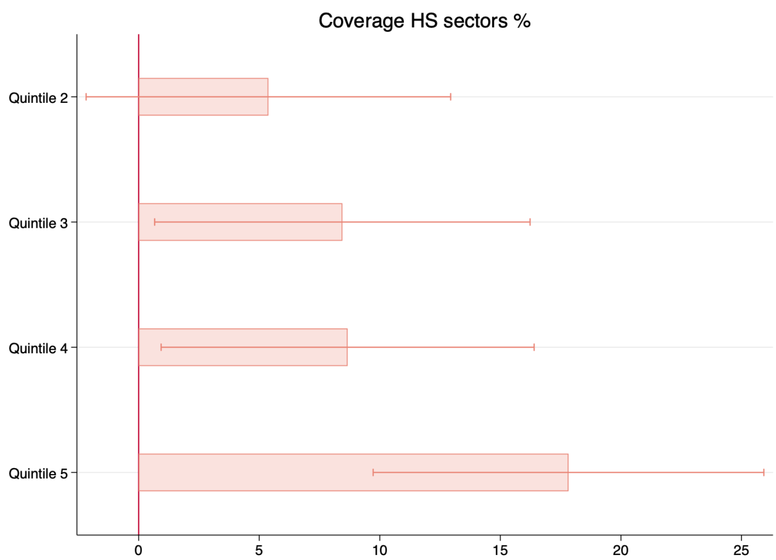

Figure 7: Regression of IP coverage ratios on income quintiles.

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.

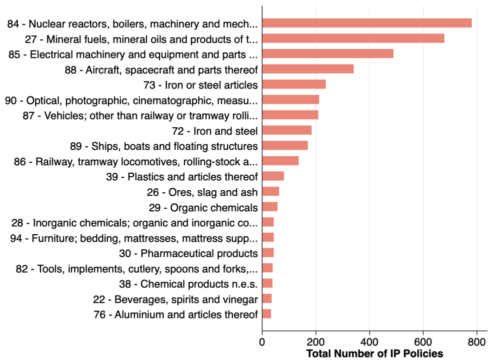

(a) Quintiles 4 and 5

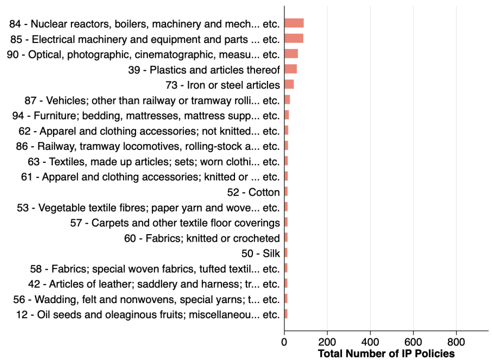

(b) Quintiles 1 and 2

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%		
	Yes		10,156	64.15%		
Firm-Specific policies	No		18,177	35.85%		
	Total		28,333	100%		
	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		
Mast Chaper Code	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(RCA + 1)				
	(1)	(2)	(3)	(4)	(5)
$IP = 1$	0.25***	0.20***	0.59***	0.20***	0.21***
	(0.010)	(0.018)	(0.002)	(0.027)	(0.007)
	Panel (B) Dependent Variable: Dummy RCA = 1 if RCA > 1				
	(1)	(2)	(3)	(4)	(5)
$IP = 1$	0.17***	0.13***	0.31***	0.11***	0.12***
	(0.008)	(0.013)	(0.001)	(0.016)	(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 regfers 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.

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.

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.

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.

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.

_ , "Protectionism, state discrimination, and international business since the onset of the Global Financial Crisis," *Journal of International Business Policy*, 2019, *2* (1), 9–36.

_ 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.

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.

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.

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.

Juhasz, Reka, "Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade," *American Economic Review*, 2018, *108* (11), 3339–3376.

Krippendorff, Klaus, *Content Analysis: An Introduction to Its Methdology*, 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.

List, Frederick, *The National System of Political Economy*, Philadelphia: J. B. Lippincott & Company, 1856.

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.

Warwick, Ken, "Beyond Industrial Policy: Emerging Issues and New Trends," Technical Report, Paris 2013.