Industrial Policy: Lessons from Shipbuilding

Panle Jia Barwick*, Myrto Kaloupstidi†, Nahim Bin Zahur‡§

September 2024

Abstract

Industrial policy has been used throughout history in some form or other by most countries. Yet, it remains one of the most contentious issues among policy makers and economists alike. In part, this is because the empirical evidence on whether and how it should be implemented remains slim. Scant data on government subsidies, conflicting theoretical arguments, and the need to account for governments' short and long-run objectives, render research particularly challenging. In this article, we outline a theory-based empirical methodology that relies on estimating an industry equilibrium model to measure hidden subsidies, assess their welfare consequences for the domestic and global economy, as well as evaluate the effectiveness of different policy designs. We illustrate this approach using the global shipbuilding industry as a prototypical example of an industry targeted by industrial policy, especially in periods of heavy industrialization. Just in the past century, Europe, followed by Japan, then South Korea, and more recently China, developed national shipbuilding programs to propel their firms to global leaders. Success has been mixed across programs, certainly by welfare metrics, and sometimes even by growth metrics. We use our methodology on China to dissect the impact of such programs, what made them more or less successful, and how we can justify why governments have chosen shipbuilding as a target.

*University of Wisconsin. Email: pbarwick@wisc.edu
†Harvard University. Email: myrto@g.harvard.edu
‡Queens University. Email: nz17@queensu.ca
§We are thankful to Elhanan Helpman for helping parse through the literature on industrial policy. We are also grateful to Leran Qi for excellent research assistance.

Introduction

Industrial policy refers to a government's agenda to shape industry structure by promoting certain industries or sectors. Although casual observation suggests that industrial policy can boost sectoral growth, researchers and policymakers have not yet mastered predicting or evaluating the efficacy of different types of government interventions, nor how to measure the overall welfare effects. What are the short and long run welfare implications of such policies for the local and global economy? How should industrial policy actually be done? Systematic and convincing empirical evidence that addresses these questions is disproportionally slim.

In this article, we focus on one particular example of industrial policy, which we believe serves as a revealing case study: government support for shipbuilding in general, and China's industrial policy to support shipbuilding in particular. As we describe below, shipbuilding has been historically a classic target of industrial policy, pursued by several countries that devised national programs for heavy industrialization, such as Japan in the 1950s, and South Korea in the 1980s. Interestingly, shipbuilding has now entered industrial policy agendas in both the European Union (Folkman, 2024) and the United States (Forooher, 2024), with calls for re-shoring shipbuilding production.

We begin this essay with an overview of global production patterns in shipbuilding, and how these patterns have shifted in the last century or so. The rich and tumultuous history of shipbuilding presents puzzles and leaves us with open questions: Why have governments subsidized shipbuilding throughout history? This is not obvious at first glance: the global market for sales of newly built ships is about $120 billion annually, which by global standards is not large. Was industrial policy successful? Although several national programs, such as Japan's or South Korea's heavy industry drives, seem to be behind tremendous sectoral growth, others, like the US's, have seemingly failed. Finally, how large was government support? Even this apparently simple question is surprisingly difficult to answer due to an "obscure jungle of subsidies" [Stråth, 1987].

We offer a brief overview of the two primary research approaches used to study industrial policy. First, a largely descriptive empirical analysis that explores the impact of specific policies on available outcomes (primarily output growth and employment) exploiting natural or quasi-experimental variation. Second, work based on strategic trade theory. However, we focus primarily on a third approach: structural modeling of the industry, and in particular an application of this approach to China's subsidies in support of shipbuilding. China is of particular interest more generally, given its strong advocacy for industrial policy, as well as its numerous trade conflicts across various industries and with many countries. In recent years especially, the government is explicitly targeting sectors with the goal of turning its firms into world leaders (e.g. "Made in China 2025" plan).

We outline the combination of structural modeling and data work needed to tackle the challenge of measuring explicit and implicit subsidies; our answer to how the global industry would have evolved if China had not subsidized shipbuilding and in particular how China's subsidies for shipbuilding affected shipbuyers, shipping costs, and world trade; the seemingly low economic payoffs that China seemed to receive from its subsidies in the beginning and how it altered the policy design over time; and the broader economic and geopolitical reasons China may have chosen to subsidize shipbuilding. In the conclusion, we offer some provisional lessons for researchers and policymakers considering industrial policies, whether in shipbuilding or in other industries.

A Brief History of Global Shipbuilding

The history of shipbuilding is as tumultuous as the seas themselves. A symbol of industrialization, maritime trade, and military strength, shipbuilding has always held a special allure for governments and has indeed been one of the most subsidized sectors throughout history. Figure 1 gives us a glimpse of its modern history and illustrates how it is defined by the succession of different countries as the world's dominant shipbuilding nation. The UK holds the lion's share of the industry for the better part of the 19th and 20th centuries, fending off competition from Western European economies (mainly, Germany and Scandinavia) at times. After World War II, it is swiftly overtaken by Japan, which prevails as a world leader via a forceful industrialization wave. Japan persists until the 1980's, when South Korea follows suit and dominates the global market. Most recently, in the mid-2000's, China comes in to claim its stake. This succession gets even more dramatic by the notorious "shipbuilding cycles"; a succession of booms and busts, more volatile than any stock market [Kalouptsidi, 2014]. Building large ships takes years, but demand for ships is governed by volatile macroeconomic fluctuations. When demand spikes, but shipbuilding capacity is fixed and sluggish, shipping rates skyrocket, and shipbuilders pile up orders for new ships. Ship prices soar. But in the bust, shipbuilding capacity idles, and prices hit rock bottom.

Why have certain countries dominated global shipbuilding? Was it the outcome of central planning? And if so, why do governments target this sector? As we discuss in this section, the economics of shipbuilding certainly plays a role in fostering the high concentration of production in certain locations. Ship production requires a base on land, easy access to water (sea or rivers), as well as materials such as steel and engines, and (skilled) labor. It is also surely not a coincidence that the countries that dominate shipbuilding often have a strong maritime tradition, and are important players in global shipping and trade. Nor that they are often undergoing a phase of heavy industrialization. But a quick inspection of the history of shipbuilding also reveals the decisive role that industrial policy, of different forms, and with different motivations, has played in shaping the global market for ships.[1]

The UK's leading position in shipbuilding, which lasted for decades up until the 1950s, initially stemmed from having access to cheap iron and steel in the late 1800s [Hanlon, 2020].[2] From the late 19th century onwards, though, the UK's dominance in shipbuilding was in large part due to its strong maritime trade: UK trade flows dominated global trade and the British Empire required ships to execute this trade volume with its colonies and other trading partners. The UK also benefitted from strong integration between its shipowners and shipbuilders - British owners almost never bought foreign ships and instead maintained close relationships with domestic shipbuilders [Pollard, 1957, Stopford,

[1] The major types of ships currently produced include containerships, (oil) tankers, bulk carriers, as well as more niche products like cruise ships, liquefied natural gas carriers, and "Ro-Ro's," which are ships that allow vehicles to be rolled on and off the ship.

[2] The US dominated wood shipbuilding in the first half of 1800s until wood ships were replaced by metal ships in the 1850s.

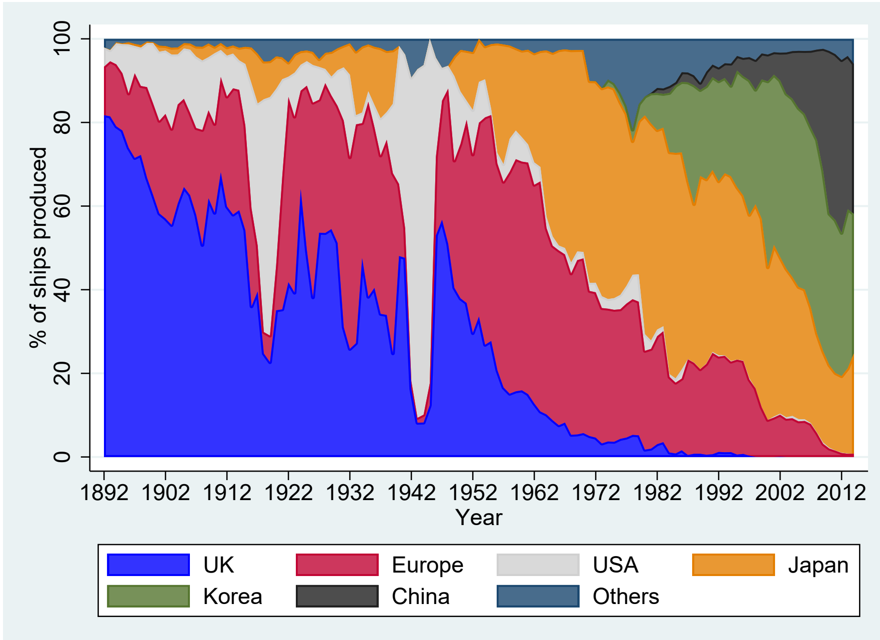

Figure 1: This figure plots the share of commercial ships produced by each country, from 1892 - 2014. Data for 1892 - 1997 was obtained from historical issues of the World Fleet Statistics published by Lloyd's Register, while the data from 1998 onwards is based on Clarksons data. We group together all European shipbuilding countries except for the United Kingdom under "Europe".

2009]. Access to a large and loyal home market allowed British shipbuilders to enjoy economies of scale, both internal and external: it facilitated the formation of a highly skilled pool of labor, enabled shipbuilders to specialize in producing specific ship types, and reduced their exposure to shipbuilding boom-and-bust cycles, allowing them to utilize their capacity more effectively [Pollard, 1957, Hanlon, 2020].

While government policy may have played a role in explaining the UK's leading position in global trade and shipping (which in turn boosted its shipbuilding sector), it is unclear if the UK directly subsidized the shipbuilding sector itself during this period. But this would turn out to be a historical anomaly. Every major shipbuilding player that subsequently emerged did so at least in part with the aid of industrial policy.

The first major challenge to UK's hegemony came from other Western European countries with strong maritime traditions, such as Scandinavian countries. Initially, many of these competitors were unable to compete globally with the significantly more productive UK yards, and instead were propped up through the use of subsidies and various protective policies [Pollard, 1957]. Over the first half of the 20th century, however, other European shipbuilders built up a substantial market share in ship production, while British shipbuilding declined. Part of the reason was that British shipbuilders failed to transition away from the "craft" style of production popular in the early 20th century to the more "industrial" style of production that became more common in the mid 20th century as ship sizes increased, ships became more standardized and ship production became more capital-intensive; that said the protectionist policies used by Britain's competitors (including generous subsidies to shipbuilders) also contributed to Britain's decline [Lorenz, 1991].

During World War I and II, the United States achieved striking production levels of ships: motivated by its military needs, massive shipbuilding capacity was built and later dismantled in a very short timespan [Thornton and Thompson, 2001, Thompson, 2001]. Beyond these two incidents, the US was never globally competitive as a ship producer. This is in spite of the fact that the US had long been a proponent of industrial policy in shipbuilding. For instance, the Jones Act prohibits foreign vessels from transporting goods between two US ports; while the "construction differential subsidies" that were provided to US shipbuilders until the 1980s amounted to between 30 and 50 percent of the cost [Stopford, 2009]. Once the high wartime demand for ships had evaporated, US shipyards, with their much higher costs, were unable to compete commercially with European and Japanese yards [Stopford, 2009]. That said, an important lesson can already be learnt from this example: a massive shipbuilding program can be set up at an astonishing speed, but, there is no guarantee that such a program can be sustained in the long run.

After World War II, in its efforts to rebuild its industrial base, Japan developed national programs for its shipbuilding industry, alongside several other heavy industries, such as steel and coal (which were viewed as complementary to one another), chemical fertilizers, petrochemicals, and automobiles [Okuno-Fujiwara, 1991, Flath, 2022]. An island nation, with a very strong maritime tradition, Japan swiftly became the world's dominant ship producer. Through a series of interventions, including subsidized financing, export credits and protectionist measures, a number of large Japanese conglomerates active in shipbuilding (as well as several other sectors) became global leaders. This includes household-name firms such as Mitsubishi, Kawasaki, and Sasebo. The program was so massive, that during the 1950's, 30% of the total loans made by the Japan Development Bank were for marine transportation (Stopford, 2009). Japanese shipbuilders were also well-positioned to take advantage of structural changes underway in global shipping (such as increased demand for transportation of oil) that increasingly favored the use of larger ships: Japan's shipyards were larger and more integrated with steelmaking facilities than their European competitors, and the Japanese industry was at the forefront of innovation in ship construction that greatly reduced construction time [Stråth, 1987]. By 1970, Japan's global market share in shipbuilding had increased to 48 percent (from only 4.7 percent in 1949), while Europe's market share had fallen from 75 percent to 48 percent; this was in spite of the heavy subsidies devoted to the European shipbuilding industry as a response to increased Japanese competition [Stråth, 1987].

By the 1980's, however, Japan was losing ground to South Korea. South Korean shipbuilding, similarly to Japanese shipbuilding, grew as part of the government's large-scale push for heavy industrialization in the late 1970's (Choi and Levchenko, 2021, Lane, 2022). Recognizing shipbuilding as a strategic industry, the government provided support primarily in the form of favorable financing (such as low-interest loans and government debt guarantees), as well as direct investment in shipbuilding facilities. Also like Japan, major South Korean conglomerates such as Hyundai, Samsung, and Daweoo grew quickly and competed fiercely with their Japanese counterparts. Indeed, Korea's shipbuilding program was from the outset focused on producing very large ships in a small number of large shipyards owned by these conglomerates [Stopford, 2009]. Within 20 years, by 1995, South

Korea's market share grew to 28% (from less than 1% in the early 1970s), reducing Japan's market share from 50% in 1975 to 41%, and Europe's share from 32% to 23%.

However, unlike its predecessors in global shipbuilding, South Korea did not have a maritime tradition, nor a large national fleet. Its maritime trade was much smaller than that of Europe, or Japan. Even as South Korea emerged as one of the two leading shipbuilding countries, the share of the global shipping fleet registered in South Korea never exceeded 2 percent, whereas Japan's share had reached 10 percent by 1984 (Stopford, 2009).[3] In other words, this was the first time that shipbuilding had been targeted as a primarily exporting sector. In 1995, 78% of Korea's ships were exported, compared to 42% for Japan (Lloyd's, 1999). This shift can be partially attributed to the growing adoption of "flags of convenience": from the late 1960's, shipowners began choosing flags of countries that provided tax and licensing benefits, such as Panama or Liberia, instead of remaining loyal to their national fleet registry. This trend rendered shipowning a more "global" industry, breaking the link between shipowning and shipbuilding. Today, demand for ships remains globally fragmented and comes from many different countries and hundreds of different shipowning firms or fleet operators.[4]

Why did South Korea choose to subsidize shipbuilding? The development of heavy industries (including not just shipbuilding, but also steel/metals, machinery, electronics and petrochemicals) was seen by Korea's government as an essential prerequisite for long-term economic growth [Lane, 2022]. National security was another key motive: changing U.S. foreign policy in the early 1970s, and in particular the withdrawal of one-third of all US troops from South Korea in 1971, led the Korean government to prioritize sectors perceived as being important for defense [Bruno and Tenold, 2011]. Finally, shipbuilding may have been targeted because of the shipbuilding production process requiring skilled labor and sophisticated capital and machinery, and even just following in the steps that Japan had taken.

By 2000, European shipbuilders were focusing on niche high-tech products, such as cruise ships. Their overall market share is 14%. Japan and South Korea compete head to head, with an overall market share of 38% each, but focused on different segments. Japanese yards dominate production of bulk carriers (with a share of 70%), while South Korean yards lead the production of higher-end, specialized oil tankers (61%) and containerships (50%).

In the 2000's, China enters the shipbuilding scene. At the turn of the century, China's nascent shipbuilding industry accounted for less than 10 percent of world production. In 2002, former Premier Zhu inspected the China State Shipbuilding Corporation (CSSC), one of the two largest shipbuilding conglomerates in China, and pointed out that "China hopes to become the world's largest shipbuilding country (in terms of output) [...] by 2015." Within a few years, China overtook Japan and South Korea to become the world's leading ship producer in terms of output. By 2009, China's market share had reached 53 percent, from less than 10 percent in 2000; the combined market share of Japan and South Korea decreased from 75 percent in 2000 to 42 percent in 2009.

[3] This estimate based on whether ships are registered in certain countries is likely a gross underestimate of the ships owned by companies from a given country, because of flags of convenience; for instance, according to Stopford [2009] in 2005, 90% of Japanese-owned ships sailed under foreign flags.

[4] Industry structure for bulk carriers (tankers and dry bulk carriers) is highly fragmented with hundreds of firms operating globally (see Kalouptsidi, 2014). In container shipping, although operators are fairly concentrated, they often lease their vessels from a large number of small shipowners.

Industrial Policy and Shipbuilding: Some Key Questions It is evident from the above narrative that industrial policy has played a pivotal role in shaping the evolution of the modern shipbuilding industry. Since the start of the 20th century, each new successor to the throne of the world's biggest shipbuilding region - continental Europe, Japan, South Korea and China - has done so on the back of a deliberate program of supporting shipbuilding. Yet this dive into history also leaves us some puzzles and open questions:

1. *Why do governments subsidize shipbuilding*? Our narrative suggests a wide variety of reasons why governments might target shipbuilding: the connection between trade, shipping, and shipbuilding; the development of heavy manufacturing as a strategy for promoting economic growth; employment; national security and military considerations; and the desire for national prestige (or as Stråth [1987] puts it, "pride and machismo"). Yet in none of the historical cases is it self-evident exactly why industrial policy was actually practiced, and what the underlying objectives of the government were.

2. *Was industrial policy successful?* It is even more challenging to evaluate if industrial policy worked. There are certainly several examples of "apparent success" (Japan, South Korea, China), where a country with a negligible initial share of the global industry embarks on a program of industrial policy and rapidly becomes a global leader. But the history of shipbuilding is also filled with examples of clearly unsuccessful industrial policy, such as the US's long-standing policy of protecting its shipbuilding sector through cabotage laws, European governments' prolonged and costly attempts to subsidize their shipbuilders in the face of Japanese and Korean competition [Stråth, 1987], or South Korea's earlier unsuccessful attempts to promote shipbuilding in the 1960s [Amsden, 1989]; these are just among the countries that were historically prominent in shipbuilding, and we have not even touched on, for example, Brazil's failed attempt to launch its own shipbuilding sector in the late 1970s [Bruno and Tenold, 2011]. And even the apparent success stories become more dubious once we take into account the massive size of policy support required to successfully take over the global shipbuilding industry, leading to the question (rarely answered in the literature) of whether the benefits from subsidizing shipbuilding are worth its large cost.

3. *How large is government support for shipbuilding?* To evaluate if industrial policy worked, one would need to know how much governments spend on industrial policy. Yet even this seemingly straightforward question does not have an easy answer, due to the plethora of different policies that governments have used to subsidize shipbuilding (what Stråth [1987] refers to as the "obscure jungle of subsidies"). This difficulty arises in part because governments have little incentive to be transparent about the full extent of policy support, especially if there is a risk of triggering retaliatory protectionist policy by rivals. Many subsidies favored by governments in practice - such as preferential access to land, or favorable financing - are implicit in nature, with their true cost difficult to uncover.

In our research, we address each of these questions through a detailed study of industrial policy in the context of the Chinese shipbuilding industry. We begin with an overview of what form China's industrial policy took.

Chinese Industrial Policy in Shipbuilding As early as 2003, China's National Marine Economic Development Plan proposed constructing three shipbuilding bases centered at the Bohai Sea area (Liaoning, Shandong, and Hebei), the East Sea area (Shanghai, Jiangsu, and Zhejiang), and the South Sea area (Guangdong). However, China's 11th National 5-year Economic Plan 2006-2010 was the first to appoint shipbuilding as a "strategic industry" in need of "special oversight and support"; the central government "unveiled an official shipbuilding blueprint to guide the medium and long-term development of the industry." As part of the national plan, the central government set specific output and capacity goals: annual production was to reach 15 million deadweight tons (DWT) by 2010 and 22 million deadweight tons by 2015. Remarkably, both goals were met several years in advance. In Figure 2, Panel A shows the rise in China's global market share of shipbuilding, by plotting China's total shipbuilding output as a share of global output. During this period, a boom-and-bust cycle took place in global shipbuilding. In the early 2000's China's international imports (mostly commodities) and exports (mostly manufacturing) boomed, commodity prices soared, and as a result shipping rates spiked to a historical high. Shipowners placed heaps of new ship orders and shipyard backlogs grew exponentially; by the end of 2008, the global ship backlog was more than five times larger than in 2001. But the shipbuilding boom was stopped short by the Great Recession of 2008. The crisis led to an idling of the existing fleet, at the same time that another 70 percent of that fleet was still scheduled for delivery by 2012. Shipbuilding prices plummeted and threatened the survival of many shipyards.

China's national and local governments provided numerous types of subsidies for shipbuilding, which we classify in three groups. First, below-market-rate land prices along the coastal regions, in combination with simplified licensing procedures, acted as "entry subsidies" that incentivized the creation of new shipyards. As shown in Panel (b) of Figure 2 depicting the number of new operating shipyards every year, between 2006 and 2008, annual construction of new shipyards in China exceeded 30 new shipyards per year; in comparison, during the same time period, Japan and South Korea averaged only about one new shipyard per year each. In the booming mid-2000s, many of the orders for new ships were placed in these Chinese "greenfields," which were taking orders as they were getting built themselves.

Second, regional governments set up dedicated banks to provide shipyards with favorable financing terms. These "investment subsidies" in the form of low-interest long-term loans (a common industrial policy tool, as illustrated also by the programs of Japan and South Korea) and preferential tax policies allowed for steep capital accumulation in the 2000s. This rise in investment in shipbuilding is illustrated in Panel (c) of Figure 2, depicting total capital invested by Chinese yards every year.

Finally, China's government also employed "production subsidies" of various forms, such as subsidized material inputs, export credits, and buyer financing. The government-buttressed domestic steel industry provided cheap steel, which is an important input for shipbuilding. Export credits and buyer financing by government-directed banks made Chinese shipyards more attractive to global buyers: when Chinese shipyards were young and unfamiliar, receiving a favorable loan to purchase from them

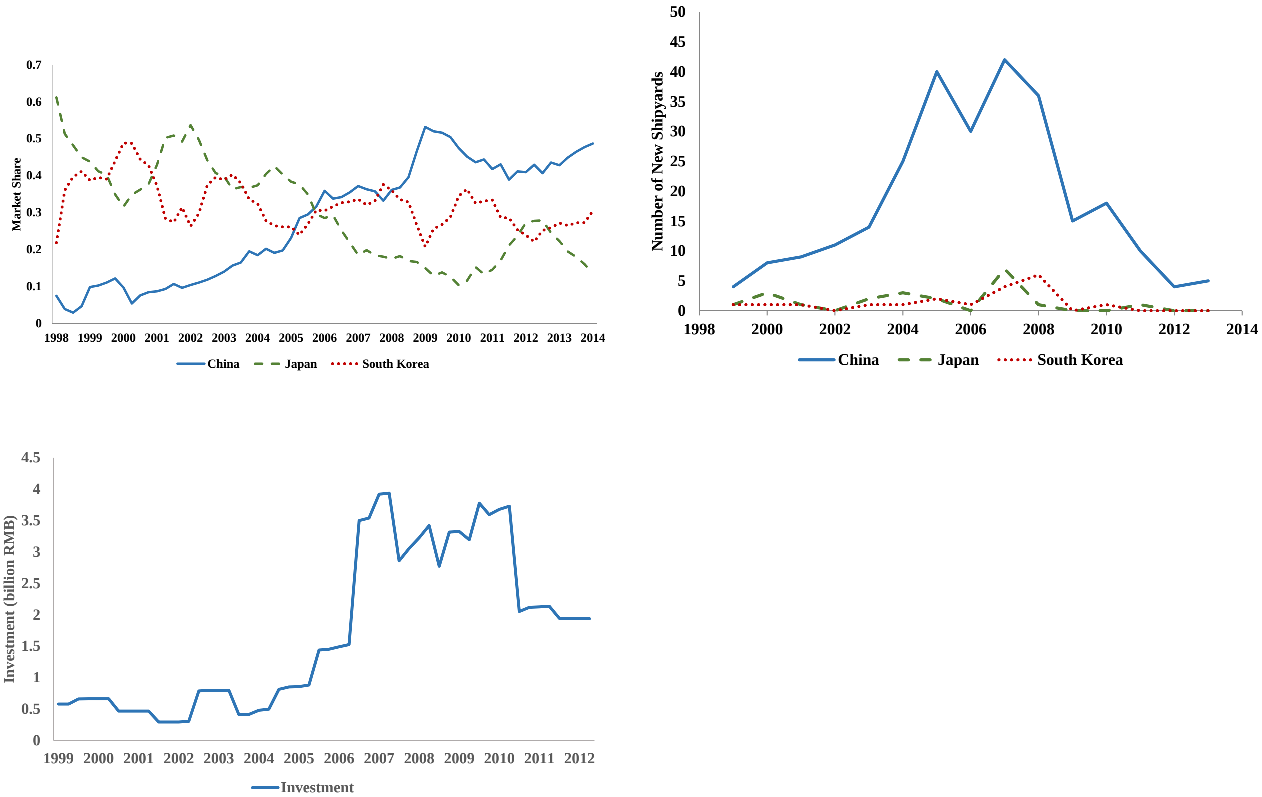

Figure 2: Figures (a) - (c) illustrate the rapid expansion of China's shipbuilding industry. (a) Market shares by country are computed from quarterly ship orders. (b) Number of new shipyards annually and by country. (c) Industry aggregate quarterly investment by Chinese shipyards in billions of 2000 RMB. Source: Barwick et al. [2024a], using data from Clarksons Research and China's National Bureau of Statistics.

proved an effective incentive.

The combination of these policies was followed by a sharp expansion in China's shipbuilding production, market share, and capital accumulation. China's market share grew from 14 percent in 2003 to 53 percent by 2009, while Japan shrunk from 32 percent to 10 percent and South Korea from 42 percent to 32 percent. This impressive output growth was partially achieved via a massive entry wave of new firms: there were 173 new Chinese firms, a 230 percent increase in five years. Indeed, one intriguing characteristic of China's industrial policy in shipbuilding—which also applies to its industrial policy in sectors such as solar panels, auto manufacturing and steel—is that the industrial policy led to a large number of small firms. This pattern is in sharp contrast to the policies adopted by Japan and South Korea, which relied on promoting a handful of large conglomerates that became global industry leaders. Most of China's shipbuilding growth at this time was concentrated in the least high-tech ships (50 percent global share in bulk carriers) vs. oil tankers or containerships (28 percent and 23 percent respectively), where it also concentrated on smaller sizes.

Then came the Great Recession of 2007-09, which drove the global shipping industry to a historic bust. The large number of new Chinese shipyards exacerbated low capacity utilization and contributed to plummeting ship prices around the world. The effectiveness of China's industrial policy was questioned. In response to the crisis and in an effort to promote industry consolidation, the government unveiled the "2009 Plan on Adjusting and Revitalizing the Shipbuilding Industry" that resulted in an immediate moratorium on entry, and subsequently shifted support towards only selected firms in an issued "White List". Concentration in the shipbuilding industry started increasing.

Since 2013, China continues to be the world's leading shipbuilder, accounting for 54 percent of all tonnage delivered in 2022, compared to 28 percent for South Korea and 18 percent for Japan. Moreover, China has begun to slowly move up the product ladder: between 2018 and 2022, Chinese shipbuilders delivered 45 percent of all new containerships (versus 36 percent for Korea), up from 23 percent between 2006 and 2010; albeit, Chinese shipyards still tend to build the smaller containerships, although they have begun to build bigger ships in the simpler ship types, such as bulk carriers.

How to Estimate Effects of Sectoral Industrial Policy

The case of shipbuilding illustrates how industrial programs can at least in some cases lead to rapid and substantial sectoral growth. But how can researchers evaluate the welfare impact of such growth, both domestically and internationally, as well as the efficacy of different types of government interventions? There are considerable challenges in answering these questions.

Any methodology for estimating effects of industrial policy must grapple with a basis data challenge: that is, government subsidies to industries are notoriously difficult to detect and measure. Indeed, "systematic data are non-existent; reliable sources of information are scarce and mostly incomplete [...] because governments do not systematically provide the information" (WTO, 2006). Thus, researchers and policymakers end up relying mostly on data reported by governments. This may include budgetary subsidies that are generally exempt from regulation, such as R&D, environmental and agricultural subsidies. There are sometimes crude reported measures of output subsidies that however tend to be untrustworthy. Even worse, certain industrial support initiatives may be almost unmeasurable: consider a government-built airport in a small city, which builds a longer-than-needed runway (at considerable cost) that is used by an aircraft manufacturer for large plane trials. Quantifying the size and effect of this indirect subsidy would be extremely challenging.[5]

Beyond the data limitations, we need a methodology to assess the impact of industrial policy. The research in this area falls into three broad categories.The first approach relies on reduced-form analyses that regress firm or sectoral outcomes on available measures of industrial policy [Aghion et al., 2015, Harris et al., 2015]. Recently, there has been a proliferation of studies that leverage natural or quasi-experimental variation in historical contexts to derive causal estimates of industrial policies on outcomes such as industry output, revenue, employment, exports, and sometimes productivity.[6] Juhász [2018] uses a natural experiment, the Napoleonic blockade of Britain from 1806-1813, to estimate the causal effect of temporary trade protection on long-term economic development. Though the blockade was not an active industrial policy, it created exogenous and differential variation in trade costs, effectively protecting northern French cotton spinners more than southern ones against British imports. Another notable example is Lane [2022]'s study of South Korea's Heavy and Chemical Industry (HCI) Drive in 1973-1979. The study highlights the role of industrial policy in shaping South Korea's long-term dynamic comparative advantage. Kim et al. [2021] studied the same historical event with more detailed firm-level data for longer periods and reported more nuanced findings: in particular, that misallocation worsened significantly during the HCI drive.[7]

Criscuolo et al. [2019] utilize exogenous variation in the eligibility criteria of place-based investment subsidies in the UK and find that these subsidies increased manufacturing employment for small firms but not large firms, with positive effects on investment but not productivity.[8] Similarly, Rotemberg [2019] leverages changes in the eligibility criteria for India's small-firm subsidy programs and finds almost complete output crowd-out among domestically consumed products but much less crowd-out for those that were exported. Aggregate productivity rose by $1-2$ percent due to the higher productivity of newly eligible firms. This finding echoes one of the obvious but often forgotten messages in Barwick et al. [2024a]: firms are inherently heterogeneous; policies that promote the participation of more efficient firms tend to be more effective.

These studies greatly enriched our understanding of the long-run effects of industrial policy, providing examples that span both developed and developing countries. As noted by Juhász et al. [2023], this literature marks a substantial improvement over the earlier, largely correlational studies. On the other hand, bounded by their research designs, none of these papers can conduct a cost-and-benefit analysis, leaving open the question of whether these industrial policies improved social welfare and

[5] To address this challenge, researchers have in recent years compiled information on industrial policies used by different countries across multiple sectors using Global Trade Alert (Juhász et al. [2022] and Evenett et al. [2024]); these fill an important data gap, since it is often hard to even know which sectors are targeted or which policies are in place; but these databases do not in general have complete information on the size of government subsidies.

[6] See Noland and Pack [2003] and Pack and Saggi [2006] for reviews of earlier work; Lane [2020] and Juhász et al. [2023] for reviews of more recent studies. There is also a growing literature that examines the effect of industrial policy on quality upgrading [Bai et al., 2024]. See Verhoogen [2023] for a review.

[7] Earlier literature (Krueger and Tuncer [1982],Birdsall et al. [1993]) and plenty of anecdotal evidence indicate that industrial policies often lead to misallocation and excess capacity, a topic we return to below.

[8] See Neumark and Simpson [2015] for a review of place-based policies.

passed the *Bastable test*, i.e. whether the discounted future gains in consumer surplus and producer surplus exceed the costs of protection [Harrison and Rodríguez-Clare, 2010]. Additionally, these studies are limited in their ability to make counterfactual predictions, such as how industries would have evolved in the absence of industrial policy, and are sometimes uninformative in terms of mechanisms by which industrial policy affects the economy. Finally, these papers are silent on the comparison across policy instruments and how to design effective and welfare-enhancing policies.

A second approach builds on the significant body of work on "strategic trade," whereby government interventions alter the strategic interactions of firms that compete globally. The literature focused on when such policies benefit the domestic economy. Early prominent work in this area included Helpman and Krugman 1989 and Brander 1995 and was based on models that were simple enough to remain analytically tractable; however, it turned out that different specifications of the model (such as firm conduct) could entirely alter the conclusions, thus not providing clear answers as to whether and how industrial policy should be employed. For instance, the theoretical predictions about the welfare implications of industrial policies depend on whether firms compete in prices or quantities. If firms are choosing quantities produced (Cournot competition), subsidies can improve national welfare. However, subsidies would reduce the welfare of national players if firms competed in prices (i.e. Bertrand competition) because they would intensify competition. This ambiguity in terms of when it is optimal to use subsidies or taxes to help the national champion in a globally oligopolistic industry obstructed the evolution of this literature.

That said, in a follow-up literature [Baldwin and Krugman, 1988b,a, Head, 1994, Irwin, 2000, Luzio and Greenstein, 1995], empirical studies applied these models to study the effect of industrial policies in specific industries, often using calibration or simulation to assess welfare effects (see Harrison and Rodríguez-Clare, 2010 for an excellent review).[9] This literature acts as a precursor to the third approach we discuss below and in contrast to studies of the first approach, the messages here are more tempered. While industrial policies led to industry growth, several of the policies evaluated [Baldwin and Krugman, 1988a, Luzio and Greenstein, 1995, Irwin, 2000] did not pass the Bastable test. Given the multiple objectives of industrial policy—such as distributional considerations, national security, and international competitiveness—the Bastable test may be too narrowly focused. Nevertheless, welfare calculations serve as a valuable benchmark for assessing the efficiency loss (or gain) from industrial policy and allow us to compare different policy tools. Finally, although these studies were informative, they often relied on aggregate data, stylized models, and calibration techniques. Importantly, few examined investment and capacity decisions, which are arguably the most critical margins that industrial policy seeks to influence. Without an increase in productive capital, infant industries cannot flourish, dynamic comparative advantages cannot emerge, and aggregate economic growth remains elusive.

A third approach applies structural methodology in the field of industrial organization [Ackerberg

[9] Baldwin and Krugman [1988a] and Baldwin and Krugman [1988b] analyze the US-Japan rivalry in the 16K RAM sector during 1978-1983, and Airbus' rise of world market share in the 1970s, as a result of trade protection and subsidies, respectively. Head [1994] studies the effect of tariff protection and learning-by-doing on the emergence of the steel rail industry in the US between the civil war and World War I, Irwin [2000] analyzes the effect of tariff protection on the growth of U.S. tinplate industry in the 1890s, and Luzio and Greenstein [1995] study the effects of Brazil's protection of the microcomputer industry in the 1980s.

et al., 2007, Ho et al., 2021]. This approach essentially seeks to put the two previous approaches together, harnessing the benefits of each. Unlike purely descriptive analysis, model-based empirical work looks at the data through the lens of an empirical industry equilibrium model, which allows for welfare calculations. And unlike the strategic trade models, it relies instead on devising the simplest model that is realistic enough to capture the main features of the environment under study. Most important, the combination of data and theory allows the researcher to test or validate the modeling assumptions imposed [Nevo and Whinston, 2010]. Finally, in the case of of industrial policy, structural methodology allows the researcher to infer a measure of subsidies that are otherwise unobserved (more on that below), as well as evaluate welfare and other (policy) counterfactuals of interest.

How does this approach work? The researcher builds a "custom" model for the industry under study: say, shipbuilding. This modeling requires a deep understanding of the environment: for example, the market structure of producers and buyers, how firms compete, how prices are formed, the production cost function and other important firm decisions, is the industry selling a homogeneous or differentiated good—and how government subsidies affect firms. Possible answers to these questions come from extensive discussions with industry participants, industry press or reviews, as well as the prior academic literature. The model must account for the key features of the industry under study, and allow for the key mechanisms the researcher is interested in analyzing. The model must then "meet" the data. The researcher collects data, usually in the form of firm actions, such as quantities produced, prices, investments, and product characteristics. Then, depending on the idiosyncrasies of the environment under study, one estimates key relationships like the demand curve, and the firm's cost curve. The goal is to "assign numbers" to key parameters of interest, such as the demand elasticity and the marginal cost of production or investment. Finally, this quantitative model is used to compute counterfactuals of interest. For example, this may include predicting the evolution of an industry with subsidies (as observed) or without, or it may involve changing the policy mix in different ways and compare the new equilibrium in the model to the observed outcomes.

A small but growing literature applies the third approach to evaluate industrial policies. Miravete and Moral [2024] examine Spanish's decade-long effort to promote its domestic automobile industry, Barwick et al. [2024b] evaluate the optimal design of China's EV subsidies, and Barwick et al. [2024c] document the positive global spillovers of the EV subsidies implemented in China, Europe, and the U.S., as well as the negative implications of local content requirements in the presence of steep learning-by-doing in the upstream EV battery sector. Chen et al. [2021] analyzed China's R&D subsidies (the InnoCom Program).

The structural approach is not without its concerns. Does the model capture the key characteristics of the industry? Does it allow for potential confounding factors, and is there rich enough data variation to reject them? Put differently, the findings are conditional on modeling assumptions; and the data must allow one to test these assumptions. One of the biggest hurdles involves the government's objectives: for instance, how should we incorporate geopolitical considerations, which are prevalent in many industrial policy agendas? Finally, this approach focuses on partial equilibrium or sector-specific analysis, which allows the researcher to exploit rich data and institutional details to answer important "what-ifs". However, general equilibrium analysis may be important in exploring spillovers of industrial policy (Liu 2019, Choi and Levchenko 2021). Finding ways forward combining benefits from different approaches seems to us a fruitful research objective.

Illustration: China's 21st Century Shipbuilding Program

We illustrate the structural approach to studying the impact of industrial policy on industry evolution and global welfare in the context of China's 21st Century shipbuilding program. This is undertaken in Kalouptsidi, 2018 and Barwick et al., 2024a which to our knowledge is the first attempt at evaluating quantitatively industrial policy in shipbuilding globally, and among the first papers employing the structural IO methodology to understand the welfare implications and effective design of industrial policy more generally. We build a model of a global market for ships. On the demand side, a large number of shipowners across the world are deciding whether to buy new vessels. Their willingness-to-pay for new ships depends on present and expected future market conditions, notably world trade and the current fleet level. On the supply side, our model considers shipyards located in China, Japan, and South Korea (that account for 90% of world production). Each shipyard decides how many ships to produce, by comparing the ship market price, dictated by the willingness-to-pay of shipowners on the demand side, and its production costs. The shipyards are price-takers (an assumption we relax below), and will keep producing as long as the price exceeds the marginal cost of the additional vessel. Thus, we can use the assumption of profit-maximizing behavior, along with observed ship prices and observed firm-level production, to uncover the shipyard's cost function. We assume the cost function is convex—that is, marginal costs of production rise with quantity—a feature that may capture capacity constraints.

To bring this model to the data, we employ a rich dataset consisting of firm-level quarterly ship production between 1998 and 2014, firm-level investment, entry and exit, and new ship market prices by ship type (containerships, tankers, and dry bulk carriers, accounting for 90 percent of all global sales). In other words, for each shipyard in the world we know how many ships are ordered and delivered every quarter, how much new capital is obtained, whether the shipyard exits or is a new entrant, as well as the prevailing global ship price. Such data on the shipbuilding industry are available from several data providers, at a relatively small cost to the researcher, such as Clarksons Research (which we use), Lloyds, and S&P Global.

As a starting point, we need to measure the size of China's subsidies for shipbuilding, both direct and indirect. We can then quantify the impact of subsidies on the economy. But how can one gain insight on governmental support offered to firms, when the policies applied are poorly measured? We know that China's subsidies for shipbuilding started in 2006. By comparing observed firm-level production before and after 2006 for Chinese and non-Chinese shipyards, we can estimate the level of subsidies. We use data from before 2006 to estimate a shipbuilding cost function. This cost function should then continue to predict shipbuilding outside of China after 2006; however, it can only predict China's shipbuilding after 2006 by including a measure of what subsidies must have been. With the model in hand, we can compare China's observed shipbuilding production with subsidies, to the outcome that would have arisen had China not subsidized shipbuilding.

In other words, our methodology relies on combining available data on firm choices and an economic model to detect the presence of subsidies. It aims at uncovering a "gap" between the observed production and the optimal production the economic model implies, i.e. the production that equates price to marginal production cost. We estimate the marginal cost from variation in prices and examine its behavior around 2006.[10] We are particularly interested in whether this cost function exhibits a "break" in China in 2006, i.e. an abrupt, and otherwise inexplicable change that makes Chinese shipyards produce as if their costs are all of a sudden lower. We illustrate our model using a toy example below.

Example. Suppose the demand curve for ships is given by $P_t = A_t - bQ_t$, where P_t is the price of a new ship in time period t (t may be a quarter or a year), Q_t is the total number of ships produced in t, A_t captures the "size of the market" in period t and may be governed by macroeconomic fluctuations, while b is a parameter determining the price elasticity for ships. Suppose that shipyards are price-takers in the market for ships and choose how many ships to produce in each period, subject to the convex cost function, $c_{jt}(q_{jt}) = c_1 q_{jt} + c_2 q_{jt}^2$, for shipyard j in period t, where the convexity may capture capacity constraints. Shipyard j then solves

$$\max_{q_{jt}} P_t q_{jt} - \left(c_1 q_{jt} + c_2 q_{jt}^2\right)$$

This implies that the optimal quantity produced is $q_{jt} = (P_t - c_1)/2c_2$, and the global ship price is $P_t = A_t - bNq_{jt} = (2c_2 A - bc_1 N)/(2c_2 - bN)$, where N is the number of shipyards.

When production subsidies are introduced for some shipyards in 2006, the cost function becomes, $c_{jt}(q_{jt}) = (c_1 - s)\, q_{jt} + c_2 q_{jt}^2$, where s is the per unit subsidy. Shipyards are now facing lower per unit costs, and their optimal production is higher than before. It is now equal to $q_{jt} = (P_t - c_1 + s)/2c_2$; equilibrium prices fall due to the higher total production.

By comparing observed firm level production before and after 2006 (as well as for Chinese and non-Chinese shipyards, which would have been captured by two different types of firms in this toy model), we can estimate the level of s. Here is the intuition: use data before 2006 to estimate the cost function, i.e. the parameters c_1, c_2 from the optimal production level; then, use the post-2006 level and the known cost parameters to back out s.[11] In other words, our model delivers the optimal firm-level production in the absence of subsidies. A deviation from the optimal production post-2006 informs us on what the subsidy s must have been. The idea is to essentially ask whether Chinese firms are "over"-producing, compared to our theoretical prediction. With the model in hand, we can then evaluate the welfare effects of industrial policy, by comparing the observed outcome with the outcome that would have arisen had China not subsidized shipbuilding.

[10] The cost function relates output to operating expenditures. As in many industries, however, costs of production are not readily observed. Standard methodology in IO ([Bresnahan, 1982, Berry et al., 1995] estimates costs from rich variation in demand. In our context, price is equal to marginal cost of production. If we impose a functional form on the marginal cost (say quadratic in quantity), we can use the observed price, and the level of output to estimate the cost function coefficients.

[11] Note that in the expression $q_{jt} = (P_t - c_1)/2c_2$, we observe q_{jt} for all j, t, as well as the price P_t. Hence we can essentially run the regression, $q_{jt} = \beta_0 + \beta_1 P_t + \epsilon_{jt}$, where $\beta_0 = 1/2c_2$ and $\beta_1 = -c_1/2c_2$. From these we can uncover β_0, β_1.

How Big Were China's Production Subsidies? Our estimates suggest that $23 billion in production subsidies was provided between 2006 and 2013. This finding is driven by the cost function obtained from this analysis, which exhibits a significant drop for Chinese producers equal to about 13-20 percent of the cost per ship. Simply put, Chinese shipbuilding firms were "over"-producing after 2006 compared to our prediction of output without subsidies.

Might some confounding factor, like a technological advance or a particular type of ship being built, explain this pattern? A general change in technology of shipbuilding does not seem to be the cause, because the drop in costs is only present for Chinese shipyards– there are no "breaks" in the estimated cost functions of Japanese or South Korean shipyards. In addition, the results are robust to many different specifications for the cost function, as well as different ways of accounting for other temporal changes. For example, the results hold when only shipyards that existed prior to 2001 are considered—which in turn suggests that cost declines are not driven by new shipyards, which may have a different technology or may be learning by doing. Moreover, the results hold if we focus on the smallest size category of bulk vessels (called Handysize), where China was already an important producer before 2006, China's production process is not characterized by significant technological advances, and product differentiation is very limited.

Finally, one may wonder about certain modeling assumptions made, such as the price-taking assumption for shipyards. This industry is globally fragmented, with the HHI ranging from 230 to 720 between 2006 and 2013, and China alone having more than 250 yards during the peak of the shipbuilding boom. Nonetheless, we carry out a version of our analysis where we relax the price-taking assumption, since industrial policy is often motivated by strategic trade considerations, and these considerations only come into play if firms can exert market power. Here we face a choice between assuming Bertrand competition (firms compete in prices) and Cournot competition (firms compete in quantities); economic theory suggests that whether or not strategic trade is effective depends on whether firms compete in prices or quantities. In the shipbuilding industry, capacity constraints are first-order; a typical shipyard can only work on a handful of ships at a same time; we therefore think it makes more sense to assume ships compete by choosing quantity. However, when we estimate a model of Cournot competition and back out markups, we find these are tiny: the average markup for bulk carriers is only 6.39% of the price of a new bulk carrier, and markups are even smaller for tankers (4.26%) and container ships (2.12%). Similarly, the assumption that ships are homogeneous (conditional on category and size) is motivated by the empirical pattern that prices of new ships are almost perfectly predictable from ship type, ship size, and time fixed effects.

Taking into account dynamics This simple framework misses several implications of industrial policy, which can affect the long-run behavior of firms. In the context of China's industrial policy in shipbuilding, favorable financing from state-owned banks allowed firms to invest and accumulate capital, as shown in Panel (c) of Figure 2. Along with the provision of land, they also led to the creation of many new firms, as shown in Panel (b) of Figure 2. Entry and investment have long-run implications for industry structure, especially under the severe volatility that characterizes shipbuilding. This is because entry, exit and investment are sluggish to adjust. For instance, during busts, firms do not necessarily exit or divest: these decisions are largely irreversible and firms may delay them in the hope that demand will recover. This in turn leads to excess capacity that stagnates during downturns. Finally, even the decision of how many ships to produce is subject to dynamic considerations, as ship production takes time: building a ship takes two to five years and thus shipyards accumulate backlogs, which can affect their future production costs, either negatively (capacity constraints), or positively (expertise acquisition or larger input orders).

The model used in our empirical analysis of industrial policy is flexible enough to capture the dynamic features of the market. Both the demand for shipping and the supply of ships are at the mercy of large macroeconomic swings, and firms operate in the shadows of severe uncertainty regarding both demand for international trade, as well as input cost shocks (for example, steel prices). Demand for new ships is driven by demand for international sea transport, which is uncertain and volatile. On the supply side, we employ a dynamic model of industry evolution, where firms can enter, exit, invest to increase their capital, and compete by producing ships. Shipyards decide whether to enter by comparing their lifetime expected profitability to entry costs, i.e. costs incurred to set up a new firm including the cost of land acquisition, shipyard construction costs (including the costs of any initial capital investments), and the implicit cost of obtaining regulatory permits. They exit if their expected profitability from remaining in the industry falls below a given threshold, capturing the shipyard's "scrap" value (that is, the costs or proceeds from liquidating the business, as well as any opportunity costs or option values of the firm). Optimal production decisions amounts to comparing current margins to expected costs, given input price fluctuations and backlog accumulation. The industry is globally fragmented enough, that we assume firms do not engage in strategic dynamic interactions; that said, they form expectations about the evolution of the industry, and in equilibrium these are correct on average [Hopenhayn, 1992].

Entry and investment subsidies are identified following the same strategy as for production subsidies. As before, we use the observed entry and investment behavior of firms to back out entry and investment costs. This is done by finding the cost parameters that bring observed behavior as close as possible to the optimal behavior implied by the dynamic model. Once entry and investment costs are estimated, subsidies are estimated by comparing Chinese to non-Chinese firms, before and after 2006. The three types of subsidies can be separately identified because they affect different decisions that firms make on which we have rich firm-level data.

Our estimates suggest that $91 billion in subsidies along all three margins—production, entry and investment— was provided over between 2006 and 2013, averaging over $11 billion per year, which totaled nearly 50 percent of Chinese shipbuilding industry revenue over that period. This is considerably larger than the $23 billion that (according to our estimates) was provided in production subsidies alone. Thus, thinking about decisions such as entry, exit and capacity investment, which all seek to take an uncertain economic future into account, will matter in evaluating the impact and effectiveness of industrial policy. Entry costs for shipbuilding exhibit a 51-64 percent decline in 2006 (depending on the province). Indeed, the entry subsidies were 69 percent of total subsidies, while production subsidies were 25 percent, and investment subsidies accounted for only 6 percent. This is driven by the empirical patterns that shipbuilding firms "over-entered"– recall the astonishing entry rates during the boom years of 2006-2008, as shown in Panel (b) of Figure 2– and "over-invested" – recall the striking increase in investment *during the bust*, post the Great Recession, as shown in Panel (c) of Figure 2.

Output and Welfare Perhaps the biggest advantage of building a quantitative model for an industry is that it can be used to estimate hypothetical scenarios. For example, what would have happened to China's shipbuilding industry absent Chinese subsidies? Presumably, China's market share would still have increased, especially during the boom years; but by how much?

Our structural model suggests that China's industrial policy in support of shipbuilding boosted China's domestic investment in shipbuilding by 140 percent, and more than doubled the entry rate: 143 shipbuilding firms entered with subsidies vs. 64 without subsidies from 2006 to 2013. It also depressed exit. Overall, industrial policy raised China's world market share in shipbuilding by more than 40 percent.

Calculating whether this increase in sectoral output should be counted as an increase in welfare is a more delicate question. Here are several slices at an answer. First, 70 percent of China's output expansion occurred via business taken from rival countries. From a global perspective, Chinese subsidies reduced South Korea's world market share from 48 percent to 39 percent and Japan's market share from 23 percent to 20 percent during 2006-2013, with profits earned by shipyards in these two countries falling by RMB 144 billion (in US dollars, roughly $21 bn). There is evidence (backed by our cost estimates) that Chinese shipyards are less efficient than their Japanese and South Korean counterparts; thus, the transfer of shipbuilding to China that occurred constitutes a misallocation of global resources.

Second, China's industrial policy for shipbuilding led to considerable declines in ship prices. Lower ship prices benefited world ship-owners somewhat, though only a modest amount accrues to Chinese ship-owners, as they accounted for a small fraction of the world fleet.

Finally and most important, although the subsidies were highly effective at achieving output growth and market share expansion, we find that they were largely unsuccessful in terms of welfare measures. The program generated modest gains in domestic producers' profit and worldwide consumer surplus. In the long run, the gross return rate of the adopted policy mix, as measured by the increase in lifetime profits of domestic firms divided by total subsidies, is only 18 percent, meaning that for every dollar the government spends, it gets back 18 cents in profitability. In other words, the net return when incorporating the cost to the government was a negative 82 percent, with entry subsidies explaining a lion's share of the negative return.

Alternative Design of Industrial Policy Our structural model also allows us to consider how different industrial policy designs may have yielded better results. Policy design is bound to be crucial, especially when some producers are more efficient than others, and when demand in the industry is highly volatile.

We first begin by comparing the efficacy of production, investment and entry subsidies (Panel A of Table 1). Although none of the policies (in isolation nor all together) yield positive returns in

Table 1: Welfare Effects of Industrial Policy

	Δ Net Profit/Subsidy	Δ Revenue/Subsidy
Panel A: Comparison of different policy instruments		
Investment Subsidy	74%	153%
Production Subsidy	50%	153%
Entry Subsidy	32%	66%
All Subsidies	18%	72%
Panel B: Industrial policy and the business cycle		
Pro-cyclical subsidies	38%	189%
Counter-cyclical subsidies	70%	168%
Panel C: Targeted industrial policy		
Subsidize all firms	37%	85%
Subsidize "White List" firms	71%	105%

Note: Panel A compares the actual policy mix where firms received all three kinds of subsidies ("All Subsidies"), with counterfactual policy mixes where the government provides only one type of subsidy. In the counterfactuals reported in Panel B and Panel C, firms receive a combination of production and investment subsidies. Panel B compares the effect of providing subsidies during the boom (2006 - 08) vs. the bust (2009 onwards). Panel C compares a policy where all firms are eligible to receive subsidies, versus a "White List" policy where only selected firms (chosen on the basis of how profitable they are) are eligible for subsidies. "ΔNet Profit/Subsidy" is our measure of the gross return on the policy, and equals the change in the discounted sum of net profits (relative to the scenario with no subsidies), divided by the discounted sum of subsidies. Likewise, "ΔRevenue/Subsidy" refers to the change in the discounted sum of revenue (relative to the scenario with no subsidies), divided by the discounted sum of subsidies. Source: Barwick et al. [2024a].

terms of lifetime profits, interestingly, production and investment subsidies can be justified if the goal is industry revenue maximization instead (the ratio of increased industry revenue to subsidies is 153 percent, meaning that $1 in subsidies generates about a $1.5 increase in lifetime revenue). This finding might explain the popularity of these subsidies in China, because the promotions of local individuals are often linked to quantity and revenue targets. Indeed, local governments played an important role in shipbuilding growth. Several shipbuilding programs were local, motivated in part by the important implications of increased shipbuilding output on local employment and industrial growth. This led to regional duplication, with several provinces having their own local shipbuilding industry.

Entry subsidies are wasteful—even by the revenue metric—and lead to increased industry fragmentation and idleness. Entry subsidies attract small, inefficient, low productivity firms. In contrast, production and investment subsidies increase the backlog and capital stock, which lead to economies of scale and drive down both current and future production costs. As such, they favor large and efficient firms. Indeed, the take-up rate for production and investment subsidies is much higher among firms that are more efficient: 82 percent of production subsidies and 68 percent of investment subsidies is allocated to firms that are more efficient than the median firm, whereas only 49 percent of entry subsidies goes to more efficient firms.

Another important consideration is the volatile nature of the shipbuilding industry, which is subject to boom and bust cycles. Our model suggests that a counter-cyclical policy would outperform the pro-cyclical policy that was adopted by a large margin: strikingly, subsidizing firms in production and investment during the boom leads to a gross rate of return of only 38 percent, whereas subsidizing firms during the downturn leads to a much higher gross return of 70 percent (Panel B of Table 1). In boom periods, the industry is operating close to full capacity, and so further expansion is costly and entails the utilization of high-cost resources. During a bust, on the other hand, the industry operates well below capacity and subsidies mobilize underutilized facilities, resulting in smaller distortions. In addition, subsidies during a boom attract inefficient firms, which pushes down the rate of return. Despite the benefits of a counter-cyclical policy, the actual policy mix was overwhelmingly pro-cyclical: 90 percent of total subsidies was handed out between 2006 and 2008 vs. 10 percent between 2009 and 2013.

Finally, we examine the consolidation policy adopted in the aftermath of the financial crisis, whereby the Chinese government implemented a moratorium on entry into shipbuilding and issued a "White List" of firms prioritized for government support (Panel C of Table 1). Indeed, this strategy was adopted in several industries to curb excess capacity and create national champions that can compete globally, following the examples of Japan and South Korea. In our calculations, if an "optimal White List" is formed—that is, the most productive firms are chosen for subsidies—the net rate of return climbs to 71 percent vs. 37 percent when all firms are subsidized. Why? Subsidizing all firms encourages suboptimal entry; in contrast the White List subsidizes existing firms and does not distort entry. Moreover, firms on the White List have lower production costs; shifting support to more efficient firms reduces misallocation. However, China's actual White List was suboptimal, as it favored state-owned enterprises. This illustrates a further difficulty with designing industrial policy: regulatory capture.

Our results highlight why industrial policies have worked better for some countries. In East Asian countries where industrial policy was considered successful, the policy support was often conditioned on export performance. In contrast, in Latin America where industrial policies often aimed at import-substitution, there were no mechanisms to weed out non-performing beneficiaries (Rodrik, 2009). In China's modern-day industrial policy in the shipbuilding industry, the policy's return was low in earlier years when output expansion was primarily fueled by the entry of inefficient firms, but increased over time as the government relied on 'performance-based' criteria via its White List). Such targeted policy design can be substantially more successful than open-ended policies that benefit all firms.

China's Industrial Policy in Shipbuilding: Why? If industrial policy related to shipbuilding had a low payoff, why did China do it? What are the objectives that a government is trying to attain via industrial policy? To what extent are policies based on "economics", as opposed to other considerations?

Many of the standard arguments for industrial policy do not seem to apply especially well to shipbuilding. The shipbuilding industry is fragmented globally, market power is limited, and markups are slim; thus, there are no "rents on the table" that, when shifted from foreign to domestic firms, outweigh the cost of subsidies. We find no evidence of industry-wide learning-by-doing (Marshallian externalities) in ship types that China expanded the most (such as bulk ships whose production technology was mature), another common rationale for industrial policy. Similarly, we find limited evidence for significant spillovers generated by shipbuilding to other domestic sectors (like steel production or the labor market), that would justify the subsidies we uncover in this context. More than 80 percent of ships produced in China are exported, which limits the fraction of subsidy benefits that is captured domestically. Dynamic considerations, whereby Chinese output growth in shipbuilding eventually forces competitors to exit, does not seem first-order either: by 2023, no substantial foreign exit has been observed.

Of course, one possibility is that the structure of China's incentives for local political leaders rewards readily observed results, and so the observed growth in shipbuilding output and global market share are sufficient to offer a political justification for the subsidies.

However, we also find support for two different rationales. First, as China became the world's biggest exporter and a close second largest importer during our sample period, transport cost reductions from increased shipbuilding and reduced shipping costs can lead to substantial increases in its trade volume. China's imports consist mainly of raw materials and are carried by bulk carriers and tankers; while its exports are mostly manufactured goods and are transported in containerships. To evaluate this argument, we carry out a back-of-the-envelope calculation of the subsidies' impact on China's trade flows. Subsidies reduced bulk carrier freight rates by 6 percent and containership freight rates by 2 percent between 2006 and 2013. Using trade elasticities from the literature (Brancaccio et al., 2020, Jeon, 2022), the industrial policy raised China's annual trade volume by 5 percent ($144 billion) between 2006 and 2013. This increase in trade was certainly large relative to the size of the subsidies (again, which averaged $11.3 billion annually between 2006 and 2013). Of course, "more trade" does not translate directly into economic wellbeing, but the relative magnitudes are suggestive.

Second, China's military ship production might have benefited from industrial policy with regard to shipbuilding. Military ship production is concentrated at state-owned yards: in particular, at 13 subsidiaries of China State Shipbuilding Corporation (CSSC) and China Shipbuilding Industry Corporation (CSIC), the two largest conglomerate shipyards that are also state-owned.[12] These subsidiaries are typically dual-use, producing both commercial and military ships in the same complex. Figure 3 plots the annual deliveries of naval and commercial ships from 2006 to 2013. Both types of deliveries experienced a several-fold increase during this period, although military production appears to have accelerated after the financial crisis and continued to increase throughout the sample period, providing suggestive evidence that China's supportive policy might have benefited its military production as well.

[12] Our primary sources are the yearly report known as IHS Jane's Fighting Ships, produced by the intelligence company IHS Jane's Saunders [2015] , as well as "Chinese Naval Shipbuilding: an Ambitious and Uncertain Course", a 2017 book about the Chinese naval shipbuilding industry Erickson [2017]. We are grateful to Elliott Mokski for discovering and collecting these datasets.

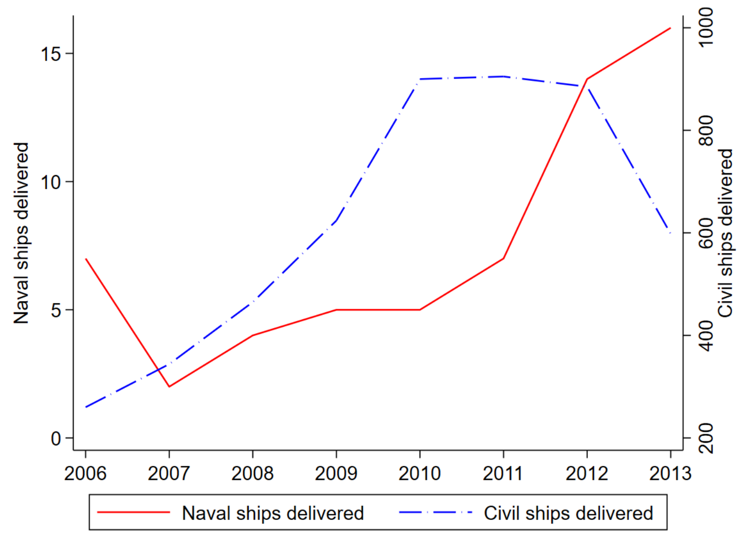

Figure 3: Military Ship Production. This figure plots the number of commercial ships and naval ships delivered by Chinese shipyards over time. Source: Barwick et al. [2024a], using data from Clarksons and IHS Jane's.

Ongoing Challenges of Research on Industrial Policy

In many ways, the example of China's industrial policy with regard to shipbuilding echoes patterns observed in other countries and industries, and it thus illustrates the academic and policy-making challenges of industrial policy.

First, it showcases the issue of scant and mismeasured data on government interventions; thus, we had to derive estimates of subsidies from a structural model and firm-level data. More generally, government subsidies to industries are notoriously difficult to detect and measure. Indeed, partly because international trade agreements prohibit direct and in-kind subsidies, "systematic data are non-existent" [WTO, 2006] and thus the presence and magnitude of industrial subsidies is often unknown. This lack of both information and compliance has obstructed the role of global policy-makers such as the WTO, and prompted them to reevaluate their guidelines.

Second, it showcases the problem of designing industrial policy; indeed, given the pressures that politicians face to support certain industries, the relevant question they face with regard to industrial policy may not be whether, but how to do it [Rodrik, 2009]. China's industrial policy for shipbuilding started with large subsidies, such as cheap land and subsidized credit, without many restrictions on the firms that could access them. This policy approach proved costly and inefficient, as it led to a massive entry wave of new firms that were not high performers. This poor design was hidden during the boom years, but painfully revealed when in the Great Recession idleness and excess capacity plagued the industry. The government then subsidized only firms on a White List; this pattern seems to continue today in the "Made in China 2025" program.

Third, economists need improved methodology for assessing the welfare impact of industrial policy, domestically and globally. Measuring welfare effects necessitates the development of a modeling framework, but there are also non-trivial choices to be made as to what is included (for example, partial vs. general equilibrium). Most important, in our view, is that evaluations of industrial policy need to reach beyond purely economic objectives: for instance, how do researchers incorporate the geopolitical considerations that are so common in industrial policy agendas today? Understanding noneconomic objectives requires economists to think outside their standard toolbox, and thus poses both a great challenge and an opportunity for research.

References

Daniel Ackerberg, C Lanier Benkard, Steven Berry, and Ariel Pakes. Econometric tools for analyzing market outcomes. *Handbook of econometrics*, 6:4171–4276, 2007.

Philippe Aghion, Jing Cai, Mathias Dewatripont, Du Luosha, Ann Harrison, and Patrick Legros. Industrial policy and competition. *American Economic Journal: Macroeconomics*, 7(4):1–32, 2015.

Alice H Amsden. Asia's next giant: South korea and late industrialization, 1989.

Jie Bai, Panle Jia Barwick, Shengmao Cao, and Shanjun Li. Quid pro quo, knowledge spillover, and industrial quality upgrading: Evidence from the chinese auto industry. 2024.

R. E. Baldwin and P. R. Krugman. Market access and international competition: A simulation study of 16k random access memories. In R. C. Feenstra, editor, *Empirical Methods for International Trade*. Cambridge, Mass.: MIT Press, 1988a.

Richard Baldwin and Paul Krugman. Industrial policy and international competition in wide-bodied jet aircraft. In *Trade policy issues and empirical analysis*, pages 45–78. University of Chicago Press, 1988b.

Panle Jia Barwick, Myrto Kalouptsidi, and Nahim Bin Zahur. Industrial policy implementation: Empirical evidence from china's shipbuilding industry. *Review of Economic Studies (accepted)*, 2024a.

Panle Jia Barwick, Hyuk-Soo Kwon, and Shanjun Li. Attribute-based subsidies and market power: an application to electric vehicles. 2024b.

Panle Jia Barwick, Hyuk-soo Kwon, Shanjun Li, and Nahim Bin Zahur. Drive down the cost: Learning by doing and government policy in the global electric vehicle battery industry. 2024c.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

Nancy M. Birdsall, Jose Edgardo L. Campos, Chang-Shik Kim, W. Max Corden, Lawrence MacDonald, Howard Pack, John Page, Richard Sabor, and Joseph E. Stiglitz. *The East Asian Miracle: Economic Growth and Public Policy.* A World Bank policy research report. World Bank Group, Washington, D.C., 1993.

Giulia Brancaccio, Myrto Kalouptsidi, and Theodore Papageorgiou. Geography, transportation, and endogenous trade costs. *Econometrica*, 88(2):657–691, 2020.

James E. Brander. Strategic trade policy. In G. Grossman and K. Rogoff, editors, *Handbook of International Economics, Chapter 27, Vol. III*. 1995.

Timothy F Bresnahan. The oligopoly solution concept is identified. *Economics letters*, 10(1-2):87–92, 1982.

Lars Bruno and Stig Tenold. The basis for south korea's ascent in the shipbuilding industry, 1970–1990. *The Mariner's Mirror*, 97(3):201–217, 2011.

Zhao Chen, Zhikuo Liu, Juan Carlos Suárez Serrato, and Daniel Yi Xu. Notching r&d investment with corporate income tax cuts in china. *American Economic Review*, 111(7):2065–2100, 2021.

Jaedo Choi and Andrei A Levchenko. The long-term effects of industrial policy. Technical report, National Bureau of Economic Research, 2021.

Chiara Criscuolo, Ralf Martin, Henry G Overman, and John Van Reenen. Some causal effects of an industrial policy. *American Economic Review*, 109(1):48–85, 2019.

Andrew S Erickson. *Chinese Naval Shipbuilding: An Ambitious and Uncertain Course*. Naval Institute Press, 2017.

Simon Evenett, Adam Jakubik, Fernando Martín, and Michele Ruta. The return of industrial policy in data. *The World Economy*, 47(7):2762–2788, 2024.

David Flath. *The Japanese Economy*. Oxford University Press, 2022.

Varg Folkman. Shipbuilders ask for eu help on chinese subsidies. *Politico*, 2024.

Rana Forooher. Shipbuilding: the new battleground in the us-china trade war. *Financial Times*, 2024.

W Walker Hanlon. The persistent effect of temporary input cost advantages in shipbuilding, 1850 to 1911. *Journal of the European Economic Association*, 18(6):3173–3209, 2020.

Richard Harris, Ian Keay, and Frank Lewis. Protecting infant industries: Canadian manufacturing and the national policy, 1870–1913. *Explorations in Economic History*, 56:15–31, 2015.

Ann Harrison and Andrés Rodríguez-Clare. Trade, foreign investment, and industrial policy for developing countries. *Handbook of development economics*, 5:4039–4214, 2010.

Keith Head. Infant industry protection in the steel rail industry. *Journal of International Economics*, 37:141–165, 1994.

Elhanan Helpman and Paul Krugman. *Trade policy and market structure*. MIT press, 1989.

Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri, editors. *Handbook of Industrial Organization, Volume 5*, volume 5. Elsevier, Princeton, NJ, United States; Chicago, IL, United States, 2021.

Hugo A Hopenhayn. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1127–1150, 1992.

Douglas A Irwin. Did late-nineteenth-century us tariffs promote infant industries? evidence from the tinplate industry. *The Journal of Economic History*, 60(2):335–360, 2000.

Jihye Jeon. Learning and investment under demand uncertainty in container shipping. *The RAND Journal of Economics*, 53(1):226–259, 2022.

Réka Juhász. Temporary protection and technology adoption: Evidence from the napoleonic blockade. *American Economic Review*, 108(11):3339–3376, 2018.

Réka Juhász, Nathan Lane, Emily Oehlsen, and Verónica C Pérez. The who, what, when, and how of industrial policy: A text-based approach. *What, When, and How of Industrial Policy: A Text-Based Approach (August 15, 2022)*, 2022.

Réka Juhász, Nathan Lane, and Dani Rodrik. The new economics of industrial policy. *Annual Review of Economics*, 16, 2023.

Myrto Kalouptsidi. Time to build and fluctuations in bulk shipping. *American Economic Review*, 104 (2):564–608, 2014.

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

Minho Kim, Munseob Lee, and Yongseok Shin. The plant-level view of an industrial policy: The korean heavy industry drive of 1973. Technical report, National Bureau of Economic Research, 2021.

Anne O Krueger and Baran Tuncer. An empirical test of the infant industry argument. *The American Economic Review*, 72(5):1142–1152, 1982.

Nathan Lane. The new empirics of industrial policy. *Journal of Industry, Competition and Trade*, 20: 209–234, 2020.

Nathan Lane. Manufacturing revolutions: Industrial policy and networks in south korea. *Quarterly Journal of Economics (accepted)*, 2022.

Ernest Liu. Industrial policies in production networks. *The Quarterly Journal of Economics*, 134(4): 1883–1948, 2019.

Lloyd's. Lloyd's register world fleet statistics, 1999. Various Issues, from 1892 - 1999. Data retrieved from `https://hec.lrfoundation.org.uk/archive-library/world-fleet-statistics`.

Edward H Lorenz. An evolutionary explanation for competitive decline: the british shipbuilding industry, 1890–1970. *The Journal of Economic History*, 51(4):911–935, 1991.

Eduardo Luzio and Shane Greenstein. Measuring the performance of a protected infant industry: The case of brazilian microcomputers. *The Review of Economics and Statistics*, 77(4):622–633, 1995.

Eugenio J. Miravete and María J. Moral. Trade protection, industrial policy, and the shaping of local preferences. Work-in-progress, 2024.

David Neumark and Helen Simpson. Place-based policies. In Gilles Duranton, J. Vernon Henderson, and William C. Strange, editors, *Handbook of Regional and Urban Economics*, volume 5, pages 1197–1287. Elsevier, 2015.

Aviv Nevo and Michael D Whinston. Taking the dogma out of econometrics: Structural modeling and credible inference. *Journal of Economic Perspectives*, 24(2):69–82, 2010.

Marcus Noland and Howard Pack. *Industrial policy in an era of globalization: Lessons from Asia*, volume 69. Peterson Institute, 2003.

Masahiro Okuno-Fujiwara. Industrial policy in japan: A political economy view, 1991.

Howard Pack and Kamal Saggi. Is there a case for industrial policy? a critical survey. *The World Bank Research Observer*, 21(2):267–297, 2006.

Sidney Pollard. British and world shipbuilding, 1890–1914: a study in comparative costs. *The Journal of Economic History*, 17(3):426–444, 1957.

Dani Rodrik. Industrial policy: don't ask why, ask how. *Middle East development journal*, 1(1):1–29, 2009.

Martin Rotemberg. Equilibrium effects of firm subsidies. *American Economic Review*, 109(10):3475–3513, 2019.

S.C. Saunders. Jane's fighting ships 2014-2015: Yearbook. Technical report, 2015. 117th Edition, IHS Inc., London.

Martin Stopford. *Maritime Economics*. New York: Routledge, 2009.

Bo Stråth. *The politics of de-industrialisation: the contraction of the West European shipbuilding industry*. Croom Helm, 1987.

Peter Thompson. How much did the liberty shipbuilders learn? new evidence for an old case study. *Journal of political Economy*, 109(1):103–137, 2001.

Rebecca Achee Thornton and Peter Thompson. Learning from experience and learning from others: An exploration of learning and spillovers in wartime shipbuilding. *American Economic Review*, 91 (5):1350–1368, 2001.

Eric Verhoogen. Firm-level upgrading in developing countries. *Journal of Economic Literature*, 61(4): 1410–1464, 2023.

WTO. *World Trade Report 2006: Exploring the links between the subsidies, trade and the WTO.* WTO, 2006.