Innovation Spurred: Evidence from South Korea's Big R&D Push*

Luis F. Jaramillo†
Job Market Paper

Chan Kim‡

August 8nd, 2024

Abstract

We study how South Korea's first "mission-oriented" R&D program shaped innovation and economic outcomes after its implementation between 1992 and 2001. Using new textual data from archival sources and a language model to identify targeted and control technological classes, we exploit that some research projects were planned but not implemented due to budget shocks. We use a local projections event study to compare the outcomes of targeted technological classes relative to control classes. Despite the absence of differential trends before the program, targeted classes doubled their future-citation-weighed patenting output and tripled their real exports relative to control classes ten years after receiving program support. These results stand when we study cross-country evidence in both outcomes. Technological classes with less concentrated patenting output before the program drive our results. Our findings suggest that technology policy played a central role in South Korea's transition to a knowledge-intensive economy.

* We are grateful to Sebastián F. Galiani, Nuno Limão, and John Wallis for their guidance, encouragement, and thoughtful suggestions. Jaedo Choi, Ginger Zhe Jin, Pablo Querubín, Steve Redding, Mateo Uribe-Castro, and several participants at the 2024 Mid-Atlantic Trade Workshop and the 2024 Washington Area International Trade Symposium provided useful comments that significantly improved the paper. We thank Nayeon Kwon for her outstanding help in the data collection process. The National Research Foundation of Korea provided timely support in the data-collection process.

† University of Maryland. Corresponding author. lfj@umd.edu
‡ University of Maryland. chankim@umd.edu

1. Introduction

Industrial policy is back: the Inflation Reduction Act, the CHIPS Act, the European Green Deal, and the Made in China 2025 Plan commit trillions of dollars for investments that aim to transform the economy's structure and the direction in which it innovates. Despite theoretical arguments that rationalize interventions to address market and coordination failures, empirical practice lacks a clear notion of what works – and what does not work. As Juhász et al. (2023) point out, the open question is increasingly not *whether* but *how* governments should conduct industrial and innovation policy.

We contribute to this debate by investigating how the G7 Program (G7P, henceforth), South Korea's first "mission-oriented" Research and Development Program, shaped innovation patterns and real outcomes over the last four decades. Though a global innovation powerhouse *today,* South Korea was not always a leader in science and technology. The G7P, active between 1992 and 2001, was the first explicit policy effort to achieve such status. A coordinated effort across different ministries, the program aimed for frontier-level (G7-country-level, hence the name) technological development in selected industries by the 2000s. The G7P was the policy response to the perceived exhaustion of South Korea's catching-up strategy and rising labor costs since the return to democracy in 1987 (Ministry of Science and Technology of the Republic of Korea, 1991).

We find that the G7P shifted the direction in which South Korea innovated. Exploiting that some R&D projects were planned but not implemented, we find that G7P-targeted technological classes doubled their future-citation-weighed patent output relative to control classes ten years *after* they were targeted. Targeted and control classes followed similar trends in this outcome in the five years before they received G7P support. Moreover, G7P-targeting is not informative of underlying economic characteristics that potentially informed sectorial choice - such as value added, output per worker or capital intensity.

Changes in innovation output had important consequences on the real economy. Though less immediate than in the case of patenting, G7P-targeted technological classes tripled their exports relative to control classes ten years *after* they were targeted. We do not find evidence of differential trends in exports the five years *before* G7P-targeting.

Policymakers often frame industrial policy within the realm of strategic competition across countries. Indeed, some industrial policies might be a policymakers' response to their counterpart's program in a rivaling country. We study this dimension and validate our within-country, baseline empirical results by investigating South Korea's performance in targeted technological classes relative to other countries, which effectively serve as placebos. We find that, despite showing similar trends before targeting, South Korean future-citation-weighed patents and exports in targeted technological classes grew substantially faster than in other countries after the G7P. In other words, South Korea outperformed other countries in G7P-supported technological classes.

While an extensive literature in economics studies innovation (Cohen, 2010) and how different policy regimes might spur it (Bloom, 2019), research studying the practice of large publicly funded R&D programs is thin. Indeed, most of the literature studies efforts that change research funding marginally and in which the government has a limited role if any at all. Notable exceptions are Gross and Sampat (2023) and Kantor and Whalley (2023), who, respectively, study how expenditure shocks related to WWII and the Apollo Mission shaped innovation and economic outcomes in the United States over the second half of the twentieth century. These papers deal with programs in times of crisis in the United States. Though highly relevant, these episodes might have limited external validity to policymakers during more mundane times in different contexts.

We study a large, actively managed applied research program that invested $7 billion (2023 dollars) to support government-picked R&D projects. The program provided funds to firms and public research institutes interested in participating in specific, government-commissioned projects. It took place in a developing economy looking to close the technological gap to the frontier in select industries.[4] The country we study had strengths in some of those (electronics, machines, materials), but not in others (energy, biotechnology). Besides the former sectors, in which South Korea had an established comparative advantage, policymakers sought to develop capacities in the latter because self-sufficiency was deemed strategic for any advanced country.

[4] Despite spectacular growth rates in the 1970s and the 1980s, South Korea was still a developing economy by 1992. According to the World Bank (2023), South Korea's per capita GDP (PPP) in 1992 was $8,126.67, which is lower than Mexico's in 2005 ($8,321.85) or Brazil's in 2007 ($8,801.60). By the same year, the Observatory of Economic Complexity (2023) ranked South Korea's export basket just slightly higher (position 19/115) than Mexico's (24/115) or Brazil's (31/115).

We investigate how future-citation-weighed patenting activity in G7P-targeted technological classes evolved relative to other *almost-targeted* classes. To address selection concerns, we exploit that the government unit that managed the G7P selected 24 candidate mega-projects. Of those 24, only 18 were implemented. Even though the remaining mega-projects passed all selection filters and were deemed high-potential by program experts, these did not receive funds due to budgetary and program-fit considerations (KISTEP, 2002). Our treated (control) group comprises technological classes related to 18 (6) mega-projects that were (were not) implemented. We discuss below how these projects map into technological classes.

As the G7P supported many technological classes over time, with new supported *cohorts* in every year from 1992 to 2000, we implement an event-study analysis that allows us to circumvent the concern that single temporal shocks contemporaneous to the G7P drive our results.

We find that the G7P substantially increased the future-citation-weighed patenting output of targeted technological classes relative to control classes. Our estimates show that by the 5th year after receiving program support, targeted classes saw their granted patents increase by 64% relative to control classes on the year before receiving program support. This effect increases to 123% by the 10th year and 232% by the 15th year. We interpret these results as causal since there were no differential trends in patenting *before* targeting. Moreover, and supporting our identification strategy, we fail to identify systematic differences between targeted and control classes in a variety of observables such as output per worker, value added, or capital intensity.

Since these programs typically have important economic motivations, we also study the G7P impact on exports. We focus on this specific dimension of real activity since it is widely acknowledged that these were central to South Korea's outstanding economic performance over the last few decades. We find that the program had *null* effects on exports over the first three years after the G7P targeted a technological class. These, however, picked up quickly. By the 5th year after receiving program support, G7P-targeted classes saw their exports grow by 62% relative to control classes on the year before targeting. This figure increases to 245**%** by the 10th year and 204% by the 15th year. We do not find pre-trends in this outcome variable before the G7P targeted a sector. Our findings underline that R&D programs might take time to yield tangible benefits in the real economy.

Technological classes with less concentrated scientific output drive our results. We show that lower concentration measures in citation shares in the decade leading up to the G7P leads to stronger program effects. Moving from the 25^{th} to the 75^{th} percentile in the Hirschman Herfindahl Index (HHI) citation shares has a detrimental effect equivalent to roughly three-quarters of the baseline effect we identify. This finding suggests that technological classes' pre-existing structures, which likely govern spillovers through knowledge networks, played an important role in determining the G7P effectiveness. Moreover, it might serve as a cautionary tale against policies that focus solely on creating "national champions" while pursuing scale, as they might limit knowledge spillovers.

To move beyond within-country comparisons and study the international dimension of the G7P, we study how targeted technological classes fared in South Korea relative to other countries around the World. Intuitively, we compare our baseline estimates for South Korea with those for other countries, which effectively serve as placebos. Our results here are like those we exposed before. We find substantial effects by the tenth year following targeting, a lack of differential trends in outcomes before G7P support, and a delayed response for exports. The interpretation here is, however, different. Targeted technological classes not only did better, but South Korea did better than other countries ~ accomplishing, in fact, the G7P's primary objective.

A substantial data effort makes our analysis possible. Through a Transparency Law request, we obtained and digitized data on circa 4,800 G7P research projects from the National Research Foundation of Korea. For each of those, we observe the mega-project affiliation, project code, name, description, objectives, the managing public research institute, individual participating firms (if any), start and finish dates, and public and private investment. We would like to observe the technological classes targeted by each project. Unfortunately, we do not. This is problematic since the outcomes we study (future-citation-weighed patents and exports) are available for the International Patent Classification (IPC) codes, which denote technological classes. How do we use the G7P information we gathered to study the important questions that motivate this paper?

We overcome this challenge with a text-based approach that leverages the information we digitized and developments in language models. We use the rich textual data we digitize to classify patents on technological classes based on text. We use the project's name, description, and objectives to feed the World Intellectual Property Organization's (WIPO) International Patent Classification Computer-Assisted Categorization (IPCCAT) tool to retrieve a technology class associated with each project. We input those characteristics into the language model. We then receive a list of related technological classes. We keep only high-quality predictions. Our list of targeted and almost-targeted technology classes follows from this exercise. We later find that our results are robust to using alternative quality thresholds.

IPCCAT is an AI-based tool that enables users, typically resource-constrained patent offices, to consistently classify an invention in one or more IPC categories based on some inputted text. Trained with over 37 million manually classified patent documents, IPCCAT provides a level of confidence for each prediction. The tool correctly classifies over 94% of the documents in the training dataset at the levels of disaggregation of the IPC code we use for this study (WIPO, 2021).

We download the universe of United States Patent and Trademark Office (USPTO) granted patents from 1980 to 2015 from USPTO's PatentView. Our sample starts twelve years before the implementation of the first G7P project and fifteen years after implementing the last one. Though data for later years is available, we choose 2015 to avoid the right-truncation problem that arises from long patent application cycles. We focus on USPTO-granted patents to keep contextual elements, such as relative market attractiveness and strength of property-rights protection, as fixed as possible. Moreover, inventors worldwide typically file important discoveries with the USPTO (Bloom et al., 2021).

We observe each patent's assignee (the legal entity holding ownership interest in the legal rights at the time of application) and research team member location, IPC code, and cited patents. We follow the literature in defining a patent as South Korean when the patent's assignee is located in South Korea. On a similar fashion, our primary outcome of interest for this dataset is the future-citation-weighed count of USPTO-granted patents to South Korea at the 4-character IPC technological class level. The literature prefers this measure, which is a way to adjust innovation output for quality, to raw patent counts which weigh equally highly used and non-used patents. We also explore how using alternative definitions of patent nationality might alter our findings. Our results do not materially change.

We download information on South Korea's exports from the UN-COMTRADE database. Our sample starts in 1980 and ends in 2015. Though this data is not available at the IPC code level, we use Lybbert & Zolas' (2014) correspondence table between the Standard International Trade Classification (SITC) Rev. 2 and the IPC code to conduct our empirical analysis. We also use South Korea's Mining and Manufacturing Survey, which provides plant-level information and is available to us between 1980 and 2003. Given the limited timeframe, we use this dataset to study the validity of our identification strategy.

Our work contributes to several literatures in economics. First, a large body of work deals with the causes and consequences of investment in R&D (Romer, 1990; Grossman & Helpman, 1991; Aghion & Howitt, 1992; Howell, 2017; Acemoglu et al., 2018; Akcigit et al., 2020; Chen et al., 2021; Dechezleprêtre et al., 2023). An important part of it investigates how policy regimes hinder or spur innovation (Bloom et al., 2019). Most of these papers study incentives, typically in the form of marginal subsidies or grants, that firms might get from the government to develop R&D projects conceived in a decentralized manner. In contrast to the existing literature, we study a large R&D program in which a public sector organization made the specific program design and selection in a centralized fashion. The G7P was a representative "mission-oriented" R&D program (Mazzucato, 2013 ; Gruber & Johnson, 2023; Kim, 2020).

Our paper is close to Gross & Sampat (2023) and Kantor and Whalley (2023), who study other "mission-oriented" programs in crisis moments in the United States (WWII and the race to the moon, respectively). However, we study a setting that might be more informative to innovation and technology policymakers in more mundane times. We study a program that did not occur during a crisis, where stakes might be higher and incentives different. Further, we study a program in a developing economy with a developed industrial base and limited innovation activities.

Second, we contribute to a growing body of literature studying industrial policy (Juhász, 2018; Kalouptsidi, 2017; Criscuolo et al., 2019; Giorcelli, 2019; Hanlon, 2020; Mitrunen, 2021; Choi & Levchenko, 2023; Choi & Shim, 2023b; Lane, 2023; Barwick et al., 2023). The work in this area focuses on technology adoption. Our paper speaks to a different phenomenon: technology development, particularly in the setting of a developing country. Though an observer might point out that technological change is exogenous to these countries (Gollin et al., 2002), recent work by Moscona & Sastry (2023) suggests they have incentives to pursue their own R&D due to the high productivity costs of inappropriate technology. Moreover, Choi & Shim (2023b) show that advanced countries might become increasingly reluctant to transfer technology as receiving countries become more affluent and start competing with them. We also contribute by studying how concentration in a country's innovation structure might affect the effectiveness of industrial policy. Altogether, our work speaks to the demand for knowledge of technology development in developing countries, a domain which remains fundamentally under-investigated.

Third, we contribute to an extensive literature studying the role of industrial policy in the economic miracles of East Asia (Johnson, 1982; Wade, 1990; Amsden, 1992; Chang, 1993; Rodrik, 1995; Krueger, 1995; Noland & Pack, 2003; Choi & Levchenko, 2023; Choi & Shim, 2023a, 2023b; Lane, 2023). This work focuses on canonical industrial policy interventions, particularly South Korea's Heavy and Chemical Industry (HCI) Drive in the 1970s. These papers explain how South Korea became a country with an established industrial base but do not address how it became a global innovation powerhouse, a leap middle-income countries often fail to make. Closer to our paper, Choi & Shim (2023b) study why countries might move from adoption subsidies to R&D subsidies. Our work complements theirs in that we document the shift in South Korea's industrial strategy and provide microeconometric evidence on the effectiveness of the country's first "mission-oriented" R&D program. Such policy shift is relevant since South Korean policymakers experimented, without much success, with more traditional and decentralized R&D tax-credits before the G7P (Kim, 2021).

The rest of the paper is as follows: section 2 provides historical context and institutional detail; Section 3 discusses our data collection process; Section 4 outlines our empirical strategy; Section 5 presents and discusses our results; Section 6 concludes.

2. Historical Context and the G7 Program[5]

Though highly successful, South Korea's insertion into the global economy was not linear. The South Korean external sector underwent several boom-and-bust cycles over the last five decades. One materialized in the late 1980s, following the sudden end of the so-called three lows: low oil price, low interest rates, and low (weak, relative to the Japanese yen) dollar. These circumstances enabled a rapid, debt-driven expansion during the second half of the 1980s. As

[5] This section relies heavily on KISTEP (2002) and KISTEP (2003).

these external conditions changed, the external sector took a hit: exports stagnated, and their share of GDP fell from 34.8% in 1987 to 23.8% in 1991 (World Bank, 2023).

Most economists agree that the crisis revealed structural weaknesses in South Korea's industrial development strategy, which favored (debt-driven) input-oriented expansion based on comparatively low labor costs (Kwon, 2021). The end of favorable external conditions and the increase in labor costs that followed the return to democracy exhausted the strategy's sources of competitive edge. Indeed, *real* labor pay rose by 53% between 1987 and 1989, far surpassing the growth rate of labor productivity. These conditions made competition in relatively low-value-added markets with other Asian countries much tougher. Without the edge of low labor costs, the prospect of competing with advanced countries was grim due to relatively poor technology capacities (Ministry of Science and Technology of the Republic of Korea, 1991).

South Korean policymakers identified the need to shift the nature of the markets where South Korea competed abroad towards increased value-added. This necessity, paired with the increased reluctance of developed countries to share technology with South Korean firms (Choi & Shim, 2023b), justified the development of "indigenous" innovation and R&D capacities. Following the relative lack of success of earlier promotion policies (Kwon, 2021), which included a R&D tax credit, policymakers identified the need for a more coordinated and concentrated effort (Ministry of Science & Technology of the Republic of Korea, 1991). The G7 Program, announced by President Roh Tae-wooh in November of 1991, responded to this challenge.

The G7P

Also known as the Highly Advanced National Program (HAN, as the river crossing Seoul), the G7P was South Korea's first National R&D program. It invested over $7 billion (2023 dollars) and mobilized over 100.000 research staff from 1992 to 2000 (Kwon, 2021). The program aimed to take South Korean R&D capacities in select sectors to the level of G7 countries by the 2000s.

The G7P supported research projects looking to address problems in applied technology, not basic science. The program supported two types of "mega-projects": (i) "product technologies" and (ii) "base technologies". Over its course, the G7P supported 18 "mega-projects", nine in each category. Our empirical strategy exploits that budget shocks and program-fit concerns reduced the length of supported “mega-projects” from 24 to 18. Each “mega-project” had smaller individual projects, for which we collected data, that map into IPC technology classes in our regression analysis.

For “product technologies”, the concern was that the private sector would not engage in these projects because they were too large and risky. Indeed, and with few exceptions, the South Korean private sector had been unwilling to engage in R&D and favored instead the continuation of previous input-driven strategies (Kwon, 2021). A government subsidy and a pooling mechanism that enabled the participation of various firms in a single research project would enhance the risk-reward profile of these investments. The distinguishing features of these projects were that they had immediate commercial applications and that South Korea already had product development capacities. Table 1 shows the nine “mega-projects” that fell within this category. Some familiar names like HDTV, a next-generation flat panel display, a high-capacity semiconductor, and an electric vehicle appear here.

For “base technologies”, the main concern was that the private sector would not find these projects profitable because they lacked immediate commercial applications. This phenomenon might lead to a typical under-provision of public goods because private agents do not incorporate the society-wide returns that a given invention brings. Table 1 shows the nine “mega-projects” supported in this category. Technologies with significant environmental and national security externalities, such as a next-generation nuclear reactor, a fuel cell, and next-generation semiconductor materials, are here.

Figure 1 summarizes the “mega-project” selection process, which was implemented in 1992 and 1995. The G7P, thus, implemented projects in two waves. The selection process started with a broad search for candidate projects led by the Research Coordination Department of the Ministry of Science and Technology in conjunction with other ministries. The G7 Expert Planning Team, the government unit created to run the G7P, received this information. The G7P unit then came up with preliminary lists of candidate “mega-projects” that might be worth considering and drafted preliminary plans for each. These were sixty in the first wave and fourteen in the second wave.

The G7P unit sent a questionnaire to hundreds of sectorial experts, mainly in ministries and universities. Among those, the experts were asked to choose the most promising (fifteen in the first wave, ten in the second wave) according to their potential in nine dimensions related to potential externalities, potential to succeed and close the technological gap with frontier countries, market potential, and fit with the program's philosophy.[6] The experts then rated all projects in terms of its potential in each of those nine dimensions.[7]

Using the survey results as input and in consultation with other ministries, the G7P unit came up with a list of projects (fourteen in the first wave, ten in the second wave) needing final approval from South Korea's General Science and Technology Council, the country's highest policymaking body in technology policy.

Though all candidate projects were deemed worthy of support and highly ranked by experts in their questionnaire answers, the Council did not go ahead with implementing some projects in each round. Among those in the first wave were a high-speed maritime ship and an aircraft core technology. Over the second wave, these were a Korean natural language processing system, an automated traffic control system, and an off-shore manufacturing plant.

These projects were not approved due to concerns about the G7P's ability to complete them following some budget changes. There were also some more idiosyncratic concerns about their fit with the G7P philosophy. The Council decided to form commissions to assess the possibility of independently supporting those projects outside the G7P. Support eventually did not materialize.

We exploit that these projects were planned but not implemented to inform our empirical analysis. A concern is that our estimates might reflect the successful selection of profitable technologies. As our previous discussion suggests, that is precisely what experts and policymakers sought. We address that concern by using as control technologies just those that

[6] Those were: technical externalities, comparison of technology level with advanced countries in the early 2000s if supported, comparison of international competitiveness in the early 2000s if supported, size of the domestic market when commercialized, size of the global market when commercialized, contribution to the general welfare, estimated R&D cost, required R&D investment, and fit with "G7P philosophy".
[7] We retrieved the detailed results of the survey for the first wave but were unable to do so for the second wave. Final project proposals seem to follow the expert's choices.

went through the complete selection process, were deemed worthy of support, but were not implemented at the end due to budget shocks.

We map the "mega-projects" into IPC technological classes using a language model. Parallel trends in exports, a variable explicitly targeted in the process, between our targeted and control technologies in the years prior to the implementation of the program supports our identification assumption. Moreover, we do not find systematic differences between the targeted and control technological classes in a variety of observable variables (such as value added, output per worker, or capital intensity) that might have influenced the policymakers' decision-making.

Once the "mega-projects" were selected, the G7P Unit designated a public research institute to run them. These institutes expanded and implemented the research plans developed by the G7P unit. Once these plans were completed and specific research projects defined, the institutes would issue public Request For Proposals in which firms, state-owned and private, would submit budgets and research plans. Research activities would follow after receiving approval from the managing public research institutes.

3. Data

Our empirical analysis relies on newly digitized data from the G7P, a language model to classify research projects in technological classes, patenting and citation data, and exports and manufacturing data. The rest of this section discusses the samples we use, the G7P Files and how we use a language model to determine targeted and control / almost-targeted technological classes, and the sources for the different outcomes we look at.

We look at two types of outcomes. The first one, which we use to study future-citation-weighed patenting (patenting sample, henceforth), is targeted and control technological classes at the 4-digit IPC level.[8] We choose this level because finer levels of disaggregation come at an important cost in terms of precision of our language model and is widely used in the innovation literature.[9] Once we restrict the sample to targeted and almost-targeted technological classes,

[8] An example might be illustrative. The IPC has 5 levels of disaggregation: 1 digit ("domain"), 3 digits ("class"), 4 digits ("sub-class"), 5 digits ("main group"), and 7 digits ("sub-group"). For a hydraulic steering gear, the 1-digit IPC would be "B – Performing operations, transporting", the 3-digit would be "B62 - Land vehicles for travelling otherwise than on rails", the 4-digit would be "B62D – Motor vehicles; trailers", the 5-digit would be "B62D3 – Steering gears", and the 7-digit would be "B62D314 – Hydraulic".

[9] The precision of the IPCCAT model at 96.2% at the 3 digit level, 94% at the 4 digit level, 89.4% at the 5 digit level, and 82% at the 7 digit level.

we keep 520 out of a universe of 646. These 520 technological classes, which we use for our empirical analysis, account for 90.7% of the USPTO-granted patents to South Korea in 1990.[10] This sample starts in 1980 and ends in 2015.

We use a second sample to study export outcomes (export sample, henceforth). Targeted and almost-targeted classes at the 3-digit IPC level comprise this sample. We follow the literature (Liu & Ma, 2023) and choose this level because the correspondence table between the IPC (at which our targeting variable is available) and the SITC (at which export data is available; Lybbert & Zolas, 2014) is noisy at finer levels. After restricting the sample to targeted and almost-targeted technological classes, we have 101 technological classes out of a universe of 131. These accounted for 80.3% of South Korea's exports in 1990. As with our patenting sample, this sample starts in 1980 and ends in 2015.

Finally, we use South Korea's Mining and Manufacturing Survey, which is currently available to us between 1980 and 2003. This source contains yearly plant-level information on sales, inputs, and outputs for South Korean establishments involved in mining or manufacturing employing ten or more employees. Given the limited timeframe to which we have access, we use this source to assess the extent to which our identification strategy addresses selection concerns. Here we focus on observables, such as output per worker our capital intensity, that policymakers might have targeted while selecting the G7P projects.

G7 Program Files and The Language Model

Our primary source for G7P information is the *G7P Yearly Project List*.[11] We obtained a copy for every year the G7P was active (1992 through 2001) through a Transparency Law Request to the National Research Foundation of Korea. We digitized and cleaned these records for information on all 4,787 G7P projects. We observe each project's G7P mega-project affiliation, name, description, objectives, managing research institute, participating firms (if any), start date, end date, and funds provided (public and private). Figure 2 shows an example of a typical page of these records.

[10] In practice, our samples consists of all treated technological classes in all domains and untreated classes in all domains except C (Chemistry; Metallurgy" and "A" (Textiles; Paper). Using IPCCAT on the control "mega-projects" yields all classes but those two.

[11] The publication name in Korean is 선도기술개발사업 과제목록 (G7 프로젝트)

We would also like to observe the specific technological class targeted by each project. We do not have such data. The lack of information presents a challenge since any econometric evaluation requires a notion of the sectors the G7P targeted.

We classify projects into technological classes by using the rich textual information in a language model. The model is IPCCAT, which stands for International Patent Code Computer-Assisted Categorization. The World Intellectual Property Organization (WIPO) developed the first version of IPCCAT in 2002 to assist resource-constrained patent offices in classifying inventions in the IPC technological classes, precisely our task. Improved and refined ever since, IPCCAT today uses data on over 37 million inventions (their abstract and description) and their human-originated classification.

For each G7P project, we do the following: (i) we input the project's name, description, and objectives into IPCCAT, (ii) we choose the level at which we want our classification done (that is, 3-digit, 4-digit, etc.), and (iii) we choose the language in which we are inputting the text.[12] Once we decide on those options, IPCCAT prints the IPC predictions for the text we inputted with a degree of confidence that ranges from 0 to 5. We detail our choices on these items below.

As discussed above, we chose the IPC 4-digit code for our patenting sample because predictions at finer levels might have low confidence due to the increased number of choices the algorithm faces. The number of categories (precision of the algorithm) rises (falls) from 646 (94%) at 4 digits to 7437 (89.4%) in 5 digits and 65158 (82%) in 7 digits (WIPO, 2021). Choosing finer levels makes the classification problem more complex without a perceived gain from increased granularity.

We use the IPC 3-digit code for our export and manufacturing survey sample. This choice follows from the fact that the correspondence between IPC codes and real production variables (exports, in this case) is imperfect at relatively fine levels of detail. This precludes us from using the IPC 4-digit code we use in our patenting sample.

[12] Another decision the researcher needs to make is the IPC version in which IPCCAT will print the predictions. It is not a relevant decision for us since the IPC codes do not change for the levels of disaggregation we use.

Finally, we would like to use high-quality IPC-code predictions only. For both cases, we decide to use predictions made with a confidence of 3 or higher. This level allows us to cover projects accounting for over 97% of total G7P funds. For those projects for which we discard information, we impute their respective IPC codes from all the other projects in the same G7P mega-project each year. We show in the Appendix that alternative quality thresholds do not substantially change our findings.

After we perform this exercise, we have a database of 4,787 research projects with information on G7P mega-project affiliation, name, description, objectives, managing research institute, participating firms (if any), start date, end date, funds provided (public and private), and targeted technology classes at the 3-digit and 4-digit IPC code level. We use this information to determine targeted classes and the time of targeting. We perform a similar exercise for the almost targeted (planned but not implemented) G7P projects, which yields the technological classes we use as controls for our targeted classes. We assume that once a class is targeted, it remains so until the end our study.

Patenting Data

We download the universe of patents granted by USPTO from 1980 to 2015. For each one of the over 7 million patents granted by USPTO, we observe the patent's application and grant years, IPC code(s), the geographic location of the assignee (the legal entity holding ownership interest in the legal rights at the time of application) and each one of the inventors, the patents it cites, and the citations from subsequent patents.

We define a patent as coming from a given country when the assignee is in such country. We show in the Appendix that our results do not change when we use more demanding definitions of patent nationality. Second, we only consider future citations coming from patents (i) classified in different 3-digit IPC codes to the underlying patent and (ii) from countries different to South Korea. We do so to avoid the so-called *home bias* (Kwon et al., 2017) and possible strategic behavior in citation patterns. Third, we divide the patent's future citation equally among all its IPC codes (that is, we use *fractional citations*). We add those citations at the 4-digit IPC code for each year, referencing the application year. We then merge this data with our database on G7P-targeted technological classes we described above.

Export Data

We use UN-COMTRADE export data for South Korea and the rest of the World for the period between 1980 and 2015. We gather this information at the Standard International Trade Classification (SITC, Rev. 2) 4-digit level. We use Lybbert and Zolas (2014) correspondence table between SITC and IPC 3-digit codes. We add exports at the IPC 3-digit code using the probability that each SITC code belongs to an IPC 3-digit code as weights. We end with a panel of exports at the IPC 3-digit code level from 1980 to 2015. We merge this data with the information we retrieved on G7P-targeted technological classes.

Manufacturing Data

We have access to South Korea's Mining and Manufacturing Survey (MMS) between 1980 and 2003. This source gives us access to plant-level information on output and input usage for all mining or manufacturing plants employing ten or more people. The MMS includes information on the Standard Industrial Classification (SIC) in which each plant operates. We use this information and Lybbert and Zolas (2014) correspondence table to determine the IPC 3-digit codes relevant to each plant. We abstract from the effects of entry by limiting our sample to plants existing before the G7P.

4. Empirical Strategy

For both our patenting and export samples, our design features (i) the use of targeted and *almost* targeted technological classes to estimate program effects and (ii) an event-study analysis. As the G7P treated different technological classes over the years it operated, we observe "cohorts" of targeted technological classes every year from 1992 to 2000.

The two features enable us to address concerns about identification that might threaten our regression analysis. The first one relates to selection: perhaps selected technologies were ripe to succeed in any case. Though the overall technology selection process was indeed endogenous, the technologies we use as controls were perceived as equally promising by G7P Experts, as per program records (KISTEP, 2002). Those were not implemented because of budget shocks and concerns about the ability of the program to sustain those projects over the long-term. Our empirical analysis supports our claim as treatment tells us *nothing* about exports, an explicitly targeted outcome, over the pre-G7P period. Moreover, we find that treatment is not informative about output, value added, output per worker, and capital intensity – which are variables policymakers might have considered while selecting the G7P projects.

The second feature, the event-study design, is convenient because it enables the exploration of treatment dynamics and exploits the fact that the G7P targeted different classes over time. Formally, our identification assumption is that targeted classes would have evolved similarly to non-targeted classes, had the G7P not been implemented. This assumption might take different forms. For example, it might relate to our previous discussion on selection. In Section 5, we never reject the null hypothesis that pre-G7P treatment coefficients differ from zero at standard confidence levels. Exports and other variables targeted by the selection process tell us *nothing* about G7P treatment before the implementation of the program.

The conditional independence assumption also relates to contemporary shocks to our causing variables that might bias our estimates. Given the several G7P cohorts, the coefficients we estimate are not derived from single years and are, therefore, less likely to be driven by contemporary shocks. Moreover, we impose a relatively stringent set of controls to account for possibly correlated shocks.

Patenting

We estimate the effect of the G7P on future-citation-weighed innovation output and industry exports. Equation 1 is our baseline specification:

$$\Delta ihs(patents)_{s,g+h} = \alpha + \beta_{g+h}\Delta \mathrm{G7P}_{s,g+h} + \delta_{c,t} + \sum_{j=1987}^{2015} X_s\gamma_j + \varepsilon_{s,g+h} \quad (1)$$

$$\Delta ihs(patents)_{s,g+h} = ihs(patents)_{s,g+h} - ihs(patents)_{s,g-1} \quad (2)$$

$$\Delta \mathrm{G7P}_{s,g+h} = \mathrm{G7P}_{s,g+h} - \mathrm{G7P}_{s,g-1} \quad (3)$$

$\Delta ihs(patents)_{s,g+h}$ is the change in (inverse hyperbolic sine, ihs) future-citation-weighed patents in an IPC 4-digit level technological class *s*, *h* years after G7P-targeting relative to the year *g – 1*, which is the year *before* targeting. Our coefficient of interest on the right-hand side is β_{g+h}, which captures the average G7P effect on treated classes at different points in time. We include $\delta_{c,t}$, a calendar year-IPC 3-digit level technological class *c* fixed effect to account for all shocks at this level. All our specifications include the interaction between technological class' *s* share of patenting output between 1987 and 1991, X_s, and calendar year dummies to account for potentially time-variant unobserved biases towards technologies in which South Korea had an existing research capability. We note that by specifying our model in a difference setting, we account for unobservable attributes at the IPC 4-digit level that do not change over time.

We allow *h* to be between -5 and +15 – that is, we investigate our outcome in the period comprised between the five years *before* a class was targeted and up to fifteen years after targeting. Our identification assumption is that, conditional on the fixed effects we include and other variables on the right-hand side, treated and control classes would have evolved similarly, had the G7P not been implemented. We cluster standard errors at the IPC 4-digit level.

We use a standard local projection approach to estimate Equation 1 above (Jordá, 2005; Dube et al., 2023). This means that (i) we estimate these regressions using OLS separately for each year and (ii) we restrict the sample to comply with the clean-control condition. In practice, we keep only "newly treated" technological classes ($\Delta\mathrm{G7P}_{s,g+h} = 1$) or clean controls ($\mathrm{G7P}_{s,g+h} = 0$). We prefer this approach to alternatives because it prevents us from doing *forbidden comparisons* where some treated observations are controls to other treated observations. These might lead to contaminated coefficient estimates.

Our estimation method choice implies that we do not have to saturate our specification with pre-period coefficients to avoid contamination. We allow for a window of 5 years to assess differential trends across technological classes. Though we include more pre-treatment lags in robustness checks, we choose this timeframe because planning exercises typically consider those time horizons.

Exports

We study exports using our export sample, which is at the IPC 3-digit level. Equation 5 gives our baseline specification:

$$\Delta ihs(exports)_{c,g+h} = \alpha + \beta_{g+h}\Delta\mathrm{G7P}_{c,g+h} + \delta_{d,t} + \textstyle\sum_{j=1987}^{2015} X_c\gamma_j + \varepsilon_{c,g+h} \tag{4}$$

$$\Delta ihs(exports)_{c,g+h} = ihs(exports)_{c,g+h} - ihs(exports)_{c,g-1} \tag{5}$$

$$\Delta\mathrm{G7P}_{c,g+h} = \mathrm{G7P}_{c,g+h} - \mathrm{G7P}_{c,g-1} \tag{6}$$

$\Delta ihs(exports)_{s,g+h}$ is the change in (inverse hyperbolic sine, ihs) of exports in an IPC 3-digit level technological class *c*, *h* years after G7P-targeting relative to the year $g-1$, the year *before* targeting. The coefficient of interest is β_{g+h}, G7P effect on treated classes at different points in time. We include a calendar year-IPC 1-digit technological class d fixed effect to account for shocks at that level. All our specifications include the interaction between technological class' c average share of exports between 1987 and 1991, X_s, and year calendar dummies to account for potentially time-variant unobserved biases towards technologies in which South Korea had an existing export capacity. As in Equation 1, we set up Equation 5 in differences, which allows us to control for unobserved characteristics at the IPC 3-digit class level. Here, we also allow h to be between -5 and +15. We include more pre-treatment lags in the robustness checks in the Appendix.

Our identification assumption is that, conditional on the fixed effects, targeted and control classes would have evolved similarly, had the G7P not been implemented. As with patenting, we use a standard local projection approach, which implies that (i) we estimate these regressions using OLS separately for each year and (ii) we restrict the sample to comply with the clean-control condition. Thus, we keep only "newly treated" technological classes ($\Delta \mathrm{G7P}_{s,g+h} = 1$) or clean controls ($\mathrm{G7P}_{s,g+h} = 0$). We cluster standard errors at the IPC 3-digit level, which is the level at which our causing variable changes in this case.

Cross-Country Evidence

How did G7P-targeted technological classes' patenting output and exports fare in comparison to the rest of the World? We study this question by taking our within-country estimations to cross-country samples in a triple-difference setting. Intuitively, we compare our baseline within-South Korea estimates to those in other countries – which effectively act as placebos.

Patenting

Equation 7 below shows the specification we estimate, as above, using a standard local projection approach:

$$\Delta ihs(patents)_{s,g+h,k} = \alpha + \beta_{g+h}\Delta \mathrm{G7P}_{s,g+h} * \mathrm{I[South\ Korea]} + \delta_{c,t,k} + \sum_{j=1987}^{2015} X_{s,k}\gamma_j + \varepsilon_{s,g+h,k} \qquad (7)$$

Note that it is identical to Equation 1 except for the inclusion of country subscript k and an indicator variable for South Korea. Here $\Delta ihs(patents)_{s,g+h}$ is the change in (ihs) future-citation-weighed patents in an IPC 4-digit level technological class s for country k, h years after G7P-targeting relative to the year *g – 1*, which is the year *before* targeting. We include $\delta_{c,t,k}$, a calendar year-country-IPC 3-digit level technological class c fixed effect to account for all shocks at this level. We also include the interaction between technological class' s share of patenting output in country k between 1987 and 1991, $X_{s,k}$, and calendar year dummies. As in our baseline within-country specifications, we allow for for h to be between -5 and 15.

Exports

Equation 8 shows the specification we estimate, as above, using a standard local projection approach:

$$\Delta ihs(exports)_{c,g+h,k} = \alpha + \beta_{g+h}\Delta \mathrm{G7P}_{c,g+h} * \mathrm{I[South\ Korea]} + \delta_{d,t,k} + \sum_{j=1987}^{2015} X_{c,k}\gamma_j + \varepsilon_{c,g+h,k} \quad (8)$$

As with Patenting, Equation 8 is identical to Equation 4 except for the inclusion of the country k subscript and an indicator variable for South Korea. Here $\Delta ihs(exports)_{s,g+h,k}$ is the change in (ihs of) exports in an IPC 3-digit level technological class c in country k, h years after G7P-targeting relative to the year *g – 1*, the year *before* targeting. We include a country-calendar year-IPC 1-digit technological class d fixed effect to account for shocks at that level. All our specifications include the interaction between technological class' c average share of exports for country k between 1987 and 1991, $X_{s,k}$, and year calendar dummies. As in our baseline within-country specifications, we allow for for h to be between -5 and 15.

5. Results

We find that G7P-targeted classes substantially increased their forward-citation-weighed patenting output and exports relative to non-targeted classes over the long run. The dynamics were, however, different: whereas patenting output increased almost immediately following targeting, exports started increasing just a few years after targeting. We first discuss our within-country results, including an expanded discussion on mechanisms and selection concerns, and then move to our cross-country findings.

Patenting

Figure 3 shows the result of estimating Equation 1 using the empirical strategy outlined above. We find that G7P-targeted classes increased their quality-adjusted (future-citation-weighed) patenting output relative to non-targeted classes. These effects varied over time. Our point estimates suggest that G7P-targeted classes increased their patenting output by 16% the year *after* they first received G7P support relative to control classes in the year *before* treatment. This metric increases to 64% by the 5^{th} year, 123% in the 10^{th} year, and 232% in the 15^{th} year. The evolution of treatment effects over time suggests that the program spurred innovation relatively quickly and had an important long-term effect on targeted classes. These effects are not linear: our point estimates do not vary much between the third and the ninth year after targeting.

Figure 3 also shows that our targeted and control groups did not have systematically different trends in the years *before* the G7P targeted a technological class. We cannot reject the null hypothesis that those coefficient estimates are equal to zero at standard confidence levels. Moreover, all estimated coefficients are very close to zero in all cases.

Exports

Figure 4 shows the result of estimating Equation 4 using the empirical strategy outlined above. We find that G7P-targeted classes increased their exports relative to non-targeted classes over the long run. Unlikely patenting output, which responded almost immediately to targeting, exports took some time to react. Our coefficients are essentially *null* the first three years after treatment and are only statistically different from zero at standard significance levels by the 5^{th} year. These point estimates suggest that targeted classes increased their exports relative to control sectors by 62% in the 5^{th} year, 245% in the 10^{th} year, and 204% in the 15^{th} year. We note that these are *real* changes. Though we measure exports in nominal dollars, our fixed effects absorb price differentials over time.

Figure 4 also shows that treated and control classes did not have differential trends on the years anteceding treatment. We are unable to reject the null hypothesis that pre-targeting coefficients are equal to zero in all cases. As we hint before, we also interpret these results as a sanity check for our research design in the patenting sample. It is widely acknowledged that exports played a central role in South Korea's economic miracle. Conversely, if there was any selection that our design does not capture, we should expect to observe it here as external market potential was a variable that the "mega-project" selection process explicitly considered. As these results confirm, targeted and almost targeted mega-projects passed them. Thus, pre-G7P export performance tells us little about selection.

We implement an important number of robustness checks to assess the extent to which decisions made while collecting data drive our findings. We assess the robustness of our findings to the logarithmic transformation of the dependent variable, alternative definitions of patent nationality, alternative quality thresholds in our language model exercise, and longer pre-treatment lags. We refer the reader to the Appendix while noting that our results are robust to alternative choices.

Further Discussion on Selection

One concern is that our baseline results might reflect successful technology selection. However, we discuss in detail how our identification strategy deals with this matter and how the lack of differential trends in outcomes is informative about this issue. Further discussion is warranted. One way in which selection might occur is comparative advantage. Perhaps our estimates reflect that South Korean policymakers chose sectors prone to succeed because they built on the underlying strengths of the South Korean economy. An example could be the electronics and home appliances sector, where South Korea was a relevant player before the G7P.

We emphasize that our estimates control for such types of pre-existing strengths when we include pre-G7P shares of patenting and exports – and they are not central to our results. Most importantly, however, the G7P supported projects in which South Korea had well-known strengths – and others in which it had no tradition, such as nuclear power and high-speed rail. In some of our control projects, like the high-speed ship project, South Korea had (and continues to have) global relevance. Such choices suggest that comparative advantage was not the sole driver of project selection.

Though we are unable to rule out selection on unobservables, we show that selection in observable economic variables was unlikely. It is what we expect from the nature of the projects that were approved but not implemented. To assess the validity of such expectation, we estimate Equation 9 below for our Manufacturing Sample using the standard local projections approach used throughout the paper:

$$\Delta Y_{f,c,g+h} = \alpha + \beta_{g+h}\Delta \mathrm{G7P}_{g+h} + \delta_{c,t} + Age_f + \varepsilon_{fc,,g+h} \tag{9}$$

Where $\Delta Y_{f,\mathrm{c},g+h}$ is the change in variable Y for plant *f* in technological class c h years after G7P-targeting relative to the year g – *1*, the year *before* the G7P targeted the technological in which the plant operates. We include $\delta_{c,t}$, a 3-digit-level technological-class fixed effect to account for shocks at that level. Age_f is plant *f*'s age. We allow h to exist between -10 and 0. $\Delta \mathrm{G7P}_{g+h}$ is defined as before in the paper. Our coefficient of interest is β_{g+h}, which is informative about differences between the targeted and almost-targeted groups in the variables we look at. Though it wouldn't necessarily invalidate our identification strategy, evidence of systematic differences between targeted and almost-targeted classes might raise concerns about our baseline findings. They might indicate that there could be some selection that was not addressed by our strategy.

We look at (log) output, (log) value-added, (log) output per worker, and relative capital intensity. We choose these variables because policymakers might have targeted them while selecting the projects. Figure 5 shows the results of estimating Equation 9. We find that targeting is not informative about these variables as our coefficient estimates are typically very close to zero and are not statistically significant in any case for the years we investigate. Though we are unable to rule out selection on unobervables, these findings alleviate remaining selection concerns.

Mechanisms

To further understand the economics behind the G7P, we study the nature of the sectors that drive our baseline results. Different theories in economic growth (Romer, 1990) emphasize the role of knowledge spillovers in spurring innovation. Indeed, these spillovers often justify policy interventions to address market failures and align private incentives with societal goals. At the same time, the practice of industrial policy has often (though not always) been associated with creating "national champions" able to exploit scale. For example, Draghi (2024) pushes for such an industrial policy. However, one concern is that with scale and concentration might come more limited spillovers. What can we learn about this tension from the application of the G7P?

We measure the level of concentration of scientific output by computing the Hirschman-Herfindahl Index (HHI) for citation shares at the technological class level in the pre-G7P period. To do so, (i) we retrieve the number of citations that any South Korean assignee received for patents linked to a technological class over the pre-G7P period, (ii) we compute each assignee's share of citations for each class, and (iii) we compute the HHI for each technological class using those shares. We use incorporate this measure in the interaction of this measure with G7P-year dummies.

Figure 6 shows the results of such exercise. For ease of interpretation, we present the results for the HHI normalizing for a change of 6380, which is the difference between the 25th and 75th percentile of the HHI distribution. Our findings suggest that sectors with *lower* levels of concentration in terms of scientific output performed much better than those with higher concentration. An HHI change of 6380 leads to a reduction in the baseline program effect of about three-quarters by the tenth year after receiving program support. These effects are substantial and underline the relevance of spillovers in determining program success. Moreover, these might serve as a cautionary tale for pushes for scale in industrial policy.

Cross-Country Evidence

Patenting

Figure 7 shows the result of estimating Equation 7. Though slightly lower in magnitudes, our findings here are like those in our within-country comparison. We find that South Korea's patenting output in G7P-targeted technological classes increased relative to other countries' following G7P support. The point estimates suggest that the increase was of 34.8% by the 5th year, 82.2% by the 10th year, and 111.7% by the 15th year. Figure 7 also shows that targeted technological classes in South Korea followed similar trends as those in other countries *before* receiving program support.

Exports

Figure 8 plots our findings after estimating Equation 8. Our results here are very similar, both in direction and magnitude, to those in our within-country regressions. We find that exports South Korea's exports in G7P-targeted technological classes increased when compared to other countries' after receiving program support, though it took time for those effects to materialize. We find negligible effects the first three years and only detect a statistically significant increase in exports by the 4th year. Our point estimates imply that exports increased, relative to the year before receiving program support, by 64.8% on the 5th year, by 171.8% on the 10th year, and 215.8% on the 15th year. We also fail to find differential trends in South Korean exports in

G7P-targeted technological classes relative to other countries *before* receiving program support.

Discussion

Our results highlight that the G7P shifted the direction in which the South Korean economy innovated. The quality-weighed patenting output of G7P-targeted technological classes grew substantially faster than control classes after treatment. These effects were quick and remained over time, suggesting that *relative* innovation levels in these classes changed permanently, even after the G7P ended operations in 2001. We find similar results when we estimate a triple difference model in which we compare South Korean patenting output in G7P-targeted technological classes to patenting output in other countries of the World, suggesting important absolute *level* effects. Overall, our analysis shows that the program successfully spurred high-quality innovation in targeted technological classes. Our work rationalizes how South Korea caught up with technological frontier over the 1990s and 2000s (Kwon et al., 2017).

These shifts in the direction of innovation had important impacts on the real economy, even if they took time to materialize. We can only detect a statistically significant effect on exports in the 5th year after the program targeted a technological class. These results contrast with the relatively rapid impact on innovation output. Our findings are similar when we estimate a triple-difference model, suggesting that South Korea, already an export powerhouse in some sectors, managed to improve its standing *relative* to other countries after the implementation of the G7P.

We find that the program was substantially more effective in spurring high-quality innovation in technological classes where scientific output was less concentrated in the advent of the G7P. We interpret this finding as indicative of the relevance of knowledge spillovers as determinants of successful innovation policy. Moreover, it might also act as a cautionary tale against excessive concentration, which might be a by-product of certain policies that, for instance, might look to create "national champions" able to exploit economies of scale. A possible cost from higher concentration might be, for example, more limited spillovers.

Taken together, our results imply that the policy embedded in the G7P was relevant to South Korea's transition into a knowledge-intensive economy. This leap is one that countries often fail to accomplish. Most of the debate about East Asia's economic miracles focused on the role that industrial policy played in enabling heavy industry over the 1970s. Yet, as countries like Brazil or Mexico suggest, the obstacles facing productive development do not stop there. We show that technology policy played a role in increasing the sophistication of the South Korean economy *after* it developed a sizeable (heavy) manufacturing sector. Economic development, as Hirschman (1958) argues, is a complex process that necessitates a strategy, not a plan, and shifting policies to address the ever-changing nature of the hurdles that developing economies face. Our findings are consistent with the story of a developmental state that opportunely shifted its industrial strategy to overcome those ever-changing hurdles.

Conclusion

We study how South Korea's first "mission-oriented" R&D program, the G7 Program, shaped innovation and economic outcomes after its implementation between 1992 and 2001. We establish that the program shifted the direction in which the South Korean economy innovated over the 1990s and 2000s, when South Korea caught up with the frontier of knowledge. Targeted technological classes doubled their forward-citation-weighed patenting output ten years after receiving program support, relative to control classes. The program effects were not limited to patenting: though less immediate than in innovation activities, targeted sectors tripled their real exports ten years after they were targeted, relative to control classes. Our results point out that the G7P had an important role in transforming South Korea's industrial economy into an innovation-driven economy.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., & Kerr, W. (2018). Innovation, reallocation, and Growth. *American Economic Review*, *108*(11), 3450–3491. https://doi.org/10.1257/aer.20130470

Aghion, P., & Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, *60*(2), 323. https://doi.org/10.2307/2951599

Akcigit, U., Hanley, D., & Serrano-Velarde, N. (2020). Back to basics: Basic research spillovers, innovation policy, and Growth. *The Review of Economic Studies*, *88*(1), 1–43. https://doi.org/10.1093/restud/rdaa061

Amsden, A. H. (1992). *Asia's next giant: South Korea and late industrialization*. Oxford University Press.

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

Bloom, N., Hassan, T. A., Kalyani, A., Lerner, J., & Tahoun, A. (2021). *The Diffusion of Disruptive Technologies*. https://doi.org/10.3386/w28999

Bloom, N., Van Reenen, J., & Williams, H. (2019). A toolkit of policies to promote innovation. *Journal of Economic Perspectives*, *33*(3), 163–184. https://doi.org/10.1257/jep.33.3.163

Chang, H.-J. (1993). The Political Economy of Industrial Policy in Korea. *Cambridge Journal of Economics*, *17*(2), 131–157. https://doi.org/10.1093/oxfordjournals.cje.a035227

Chen, Z., Liu, Z., Suárez Serrato, J. C., & Xu, D. Y. (2021). Notching R&D investment with corporate income tax cuts in China. *American Economic Review*, *111*(7), 2065–2100. https://doi.org/10.1257/aer.20191758

Choi, J., & Shim, Y. (2023a). *From Adoption to Innovation: State-Dependent Technology Policy in Developing Countries*. Manuscript.

Choi, J., & Shim, Y. (2023b). *Technology Adoption and Late Industrialization*. Manuscript.

Cohen, W. M. (2010). Fifty Years of empirical studies of innovative activity and performance. *Handbook of The Economics of Innovation, Vol. 1*, 129–213. https://doi.org/10.1016/s0169-7218(10)01004-x

Criscuolo, C., Martin, R., Overman, H. G., & Van Reenen, J. (2019). Some causal effects of an industrial policy. *American Economic Review*, *109*(1), 48–85. https://doi.org/10.1257/aer.20160034

Draghi, M. (2024). Radical Change is What is Needed. Le Groupe d'études Géopolitiques.

Dechezleprêtre, A., Einiö, E., Martin, R., Nguyen, K.-T., & Van Reenen, J. (2023). Do tax incentives increase firm innovation? an rd design for R&D, Patents, and Spillovers. *American Economic Journal: Economic Policy*, *15*(4), 486–521. https://doi.org/10.1257/pol.20200739

Dube, A., Girardi, D., Jorda, O., & Taylor, A. M. (2023). A local projections approach to difference-in-differences event studies. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4433863

Giorcelli, M. (2019). The long-term effects of management and technology transfers. *American Economic Review*, *109*(1), 121–152. https://doi.org/10.1257/aer.20170619

Gollin, D., Parente, S., & Rogerson, R. (2002). The Role of Agriculture in Development. *American Economic Review*, *92*(2), 160–164.

Gross, D., & Sampat, B. (2023). America, jump-started: World War II R&D and the takeoff of the U.S. Innovation System. *American Economic Review*. https://doi.org/10.3386/w27375

Grossman, G., & Helpman, E. (1990). Comparative advantage and long-run growth. *American Economic Review*, *80*(4), 796–815. https://doi.org/10.3386/w2809

Gruber, J., & Johnson, S. (2023). *Jump-starting america: How breakthrough science can revive economic growth and the American Dream*. PublicAffairs.

Hanlon, W. W. (2018). Skilled immigrants and American industrialization: Lessons from Newport News Shipyard. *Business History Review*, *92*(4), 605–632. https://doi.org/10.1017/s0007680519000023

Hirschman, A. O. (1958). *The strategy of Economic Development*. Yale Univ. Press.

Howell, S. T. (2017). Financing innovation: Evidence from R&D grants. *American Economic Review*, *107*(4), 1136–1164. https://doi.org/10.1257/aer.20150808

Johnson, C. (1982). *Miti and the Japanese miracle*. Stanford University Press.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, *95*(1), 161–182. https://doi.org/10.1257/0002828053828518

Juhász, R. (2018). Temporary Protection and Technology Adoption: Evidence from the Napoleonic blockade. *American Economic Review*, *108*(11), 3339–3376. https://doi.org/10.1257/aer.20151730

Juhász, R., Lane, N., & Rodrik, D. (2023). The New Economics of Industrial Policy. *Annual Review of Economics*. https://doi.org/10.3386/w31538

Kalouptsidi, M. (2017). Detection and impact of industrial subsidies: The case of chinese shipbuilding. *The Review of Economic Studies*, *85*(2), 1111–1158. https://doi.org/10.1093/restud/rdx050

Kantor, S., & Whalley, A. (2023). Moonshot: Public R&D and growth. *National Bureau of Economic Research*, *31471*. https://doi.org/10.3386/w31471

Kim, H. C. (2020). (rep.). *National R&D Program Planning and Management in Terms of MOP in Korea*.

KISTEP. (2002). (rep.). *A Study on the Highly Advanced National(HAN) Program Evaluation I*.

KISTEP. (2003). (rep.). *A Study on the Highly Advanced National(HAN) Program Evaluation II*.

Krueger, A. O. (1995). East Asian Experience and Endogenous Growth Theory. In *Growth theories in light of the East Asian experience*. essay, University of Chicago Press.

Kwon, H. (2021). *Changes by competition: The evolution of the south korean developmental state*. Oxford University Press.

Kwon, S., Lee, J., & Lee, S. (2017). International trends in Technological Progress: Evidence from Patent Citations, 1980–2011. *The Economic Journal*, *127*(605). https://doi.org/10.1111/ecoj.12314

Lane, N. (2023). Manufacturing revolutions: Industrial policy and industrialization in South Korea. *Quarterly Journal of Economics*. https://doi.org/10.31235/osf.io/6tqax

Levchenko, A., & Choi, J. (2023). *The Long-Term Effects of Industrial Policy*. https://jaedochoi.github.io/Web/Choi_Levchenko.pdf

Liu, E., & Ma, S. (2023). Innovation networks and R&D allocation. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3994285

Lybbert, T. J., & Zolas, N. J. (2014). Getting patents and economic data to speak to each other: An 'algorithmic links with probabilities' approach for joint analyses of patenting and economic activity. *Research Policy*, *43*(3), 530–542. https://doi.org/10.1016/j.respol.2013.09.001

Mazzucato, M. (2013). *The Entrepreneurial State*. Anthem Press.

Ministry of Science and Technology, Republic of Korea. (1991). (rep.). *Highly Advanced National Project: A Plan Towards 21st Century*.

Mitrunen, M. (2021). *Industrial Policy, Structural Change, and Intergenerational Mobility: Evidence from the Finnish War Reparations*. Manuscript.

Moscona, J., & Sastry, K. (2023). Inappropriate technology: Evidence from Global Agriculture. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3886019

Noland, M., & Pack, H. (2003). *Industrial policy in an era of Globalization: Lessons from Asia*. Institute for International Economics.

Rodrik, D., Grossman, G., & Norman, V. (1995). Getting interventions right: How South Korea and Taiwan grew rich. *Economic Policy*, *10*(20), 53. https://doi.org/10.2307/1344538

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, *98*(5, Part 2). https://doi.org/10.1086/261725

Wade, R. (1990). *Governing the market: Economic theory and the role of government in East Asian industrialization*. Princeton University Press.

World Bank. (2023). *World Development Indicators*. DataBank. https://databank.worldbank.org/source/world-development-indicators

Figures

Figure 1
Selection Process

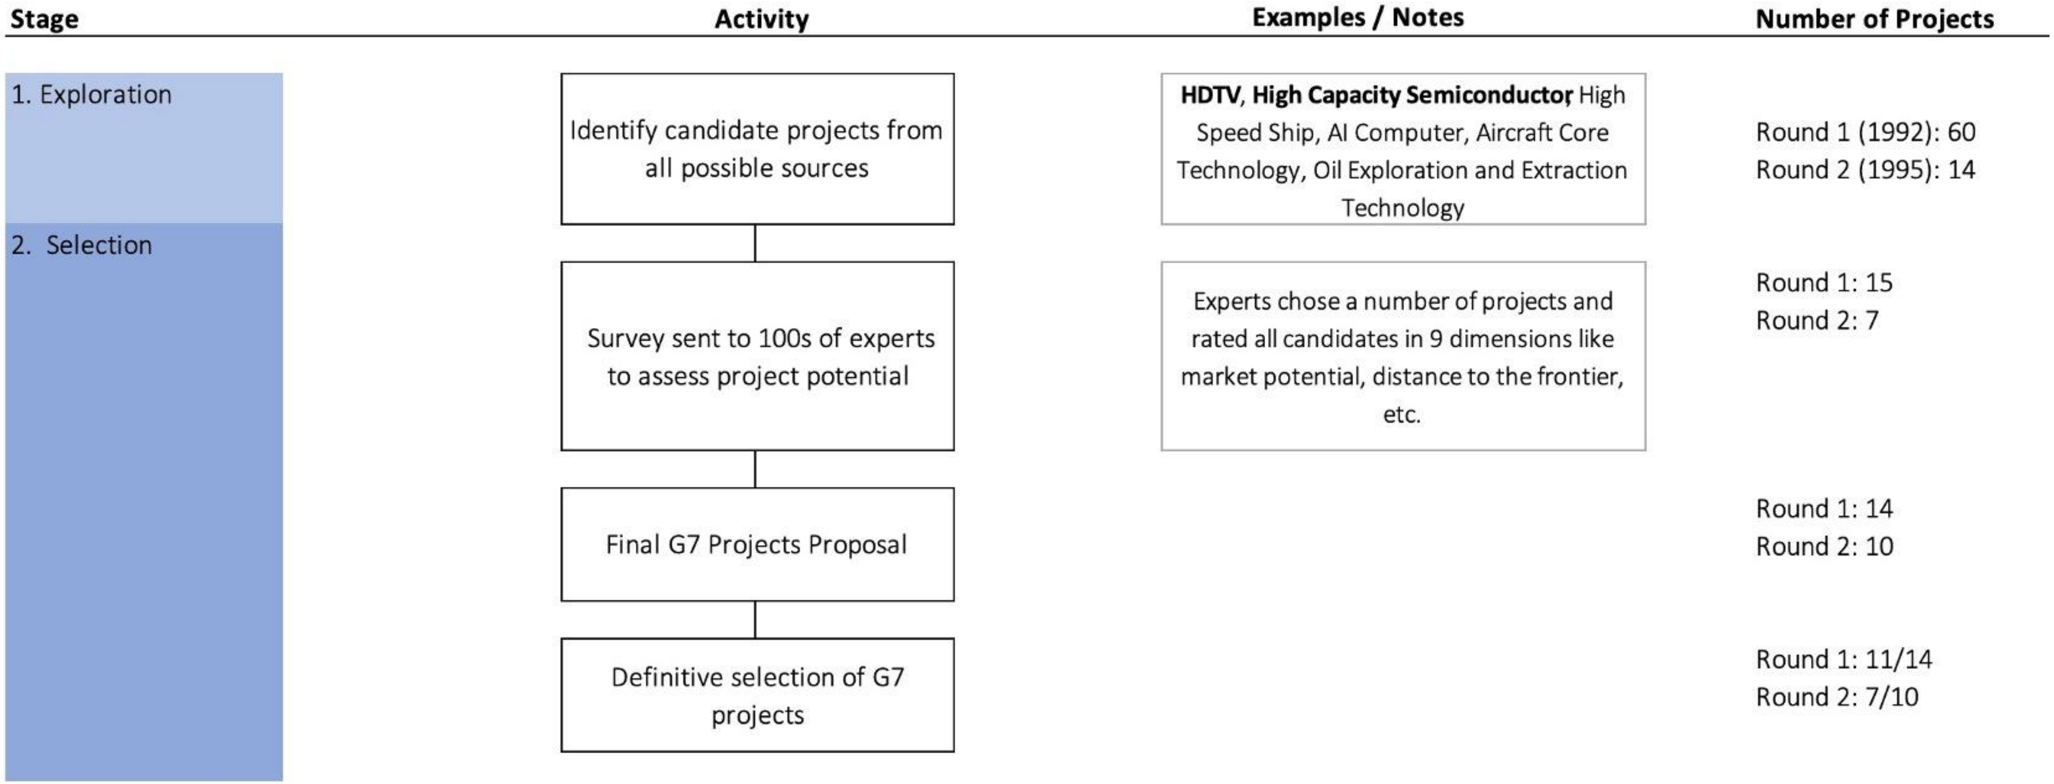

Source: KISTEP (2002, 2003)

Figure 2

Typical Record of the *G7P Yearly Project List* - 선도기술개발사업 과제목록 (G7 프로젝트)

사업구분 | 선도기술개발사업

과제번호	과 제 명	연구기관 (책임자)	참여기업	연구기간	'95 연구개발비(단위:천원)			최종목표	연구내용
					정부	기업	계		
95-G-02-01-A	교환기술분야개발	전자통신연구소 (임주환)	한화정보통신 동아전기 삼성전자(주) 대우통신(주) 우진전자통신(주) LG정보통신(주)	'92 ~ '97 ('95/01/01 ~'95/12/31)	32,868,000	44,237,000	77,105,000		
95-G-02-01-A-01	ATM 교환기 시스템 개발	전자통신연구소 (한치문)	한화정보통신 동아전기 삼성전자(주) 대우통신(주) 우진전자통신(주) LG정보통신(주)	'92 ~ '97 ('95/01/01 ~'95/12/31)	27,520,000	38,669,000	66,189,000	정보화 사회의 구축에 핵심적인 광대역 ATM기술, 광교환기술 등 차세대 교환기술개발	○소형 ATM 교환기 개발 완료 ○중형 ATM 교환기 구조 설계
95-G-02-01-A-01-A	ATM 교환기에서의 과부하제어에 관한 연구	한남대학교 (최진규)					15,000		
95-G-02-01-A-01-AA	운용메시지의 음성화에 관한 연구	과학기술원 (오영환)					15,000		
95-G-02-01-A-01-AB	ATM 교환기의 내진동 설계 및 해석에 관한연구	과학기술원 (엄윤용)					25,000		

–59–

Figure 3

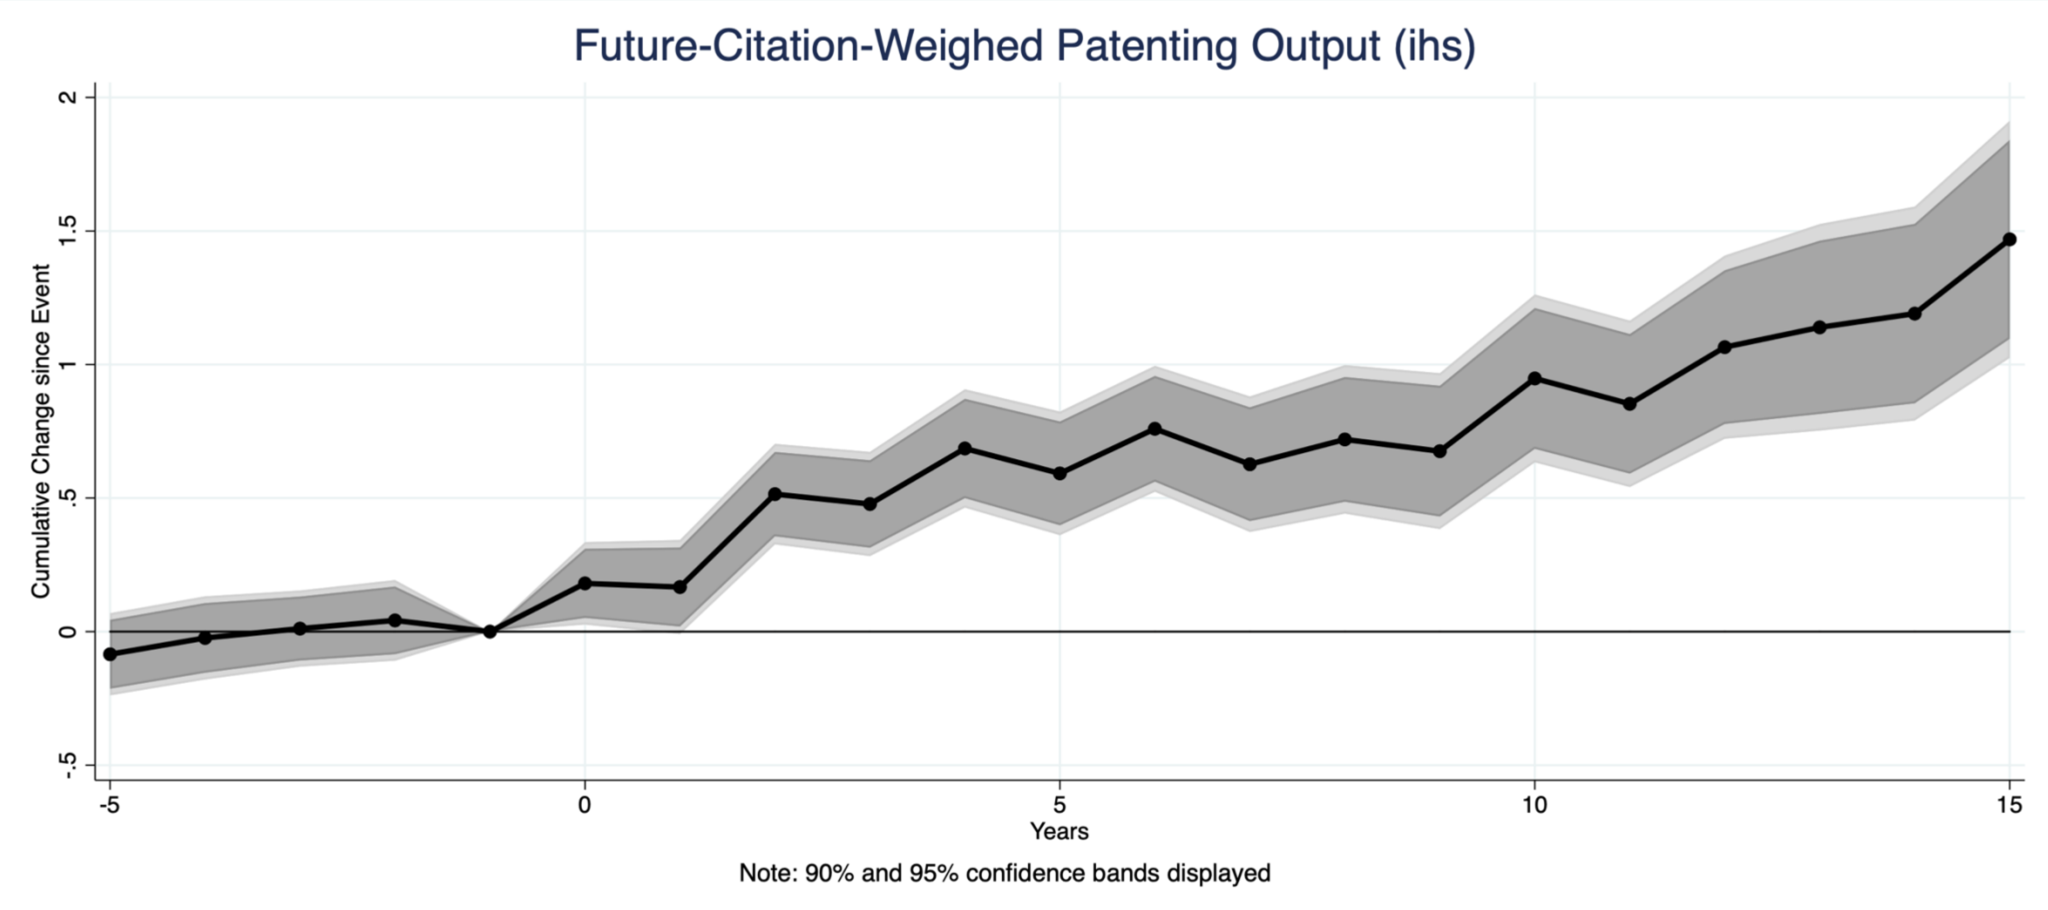

Future-Citation-Weighed Patenting Output (ihs)
Cumulative Change since Event
2
1.5
1
.5
0
-.5
-5
0
5
10
15
Years
Note: 90% and 95% confidence bands displayed

Figure 4

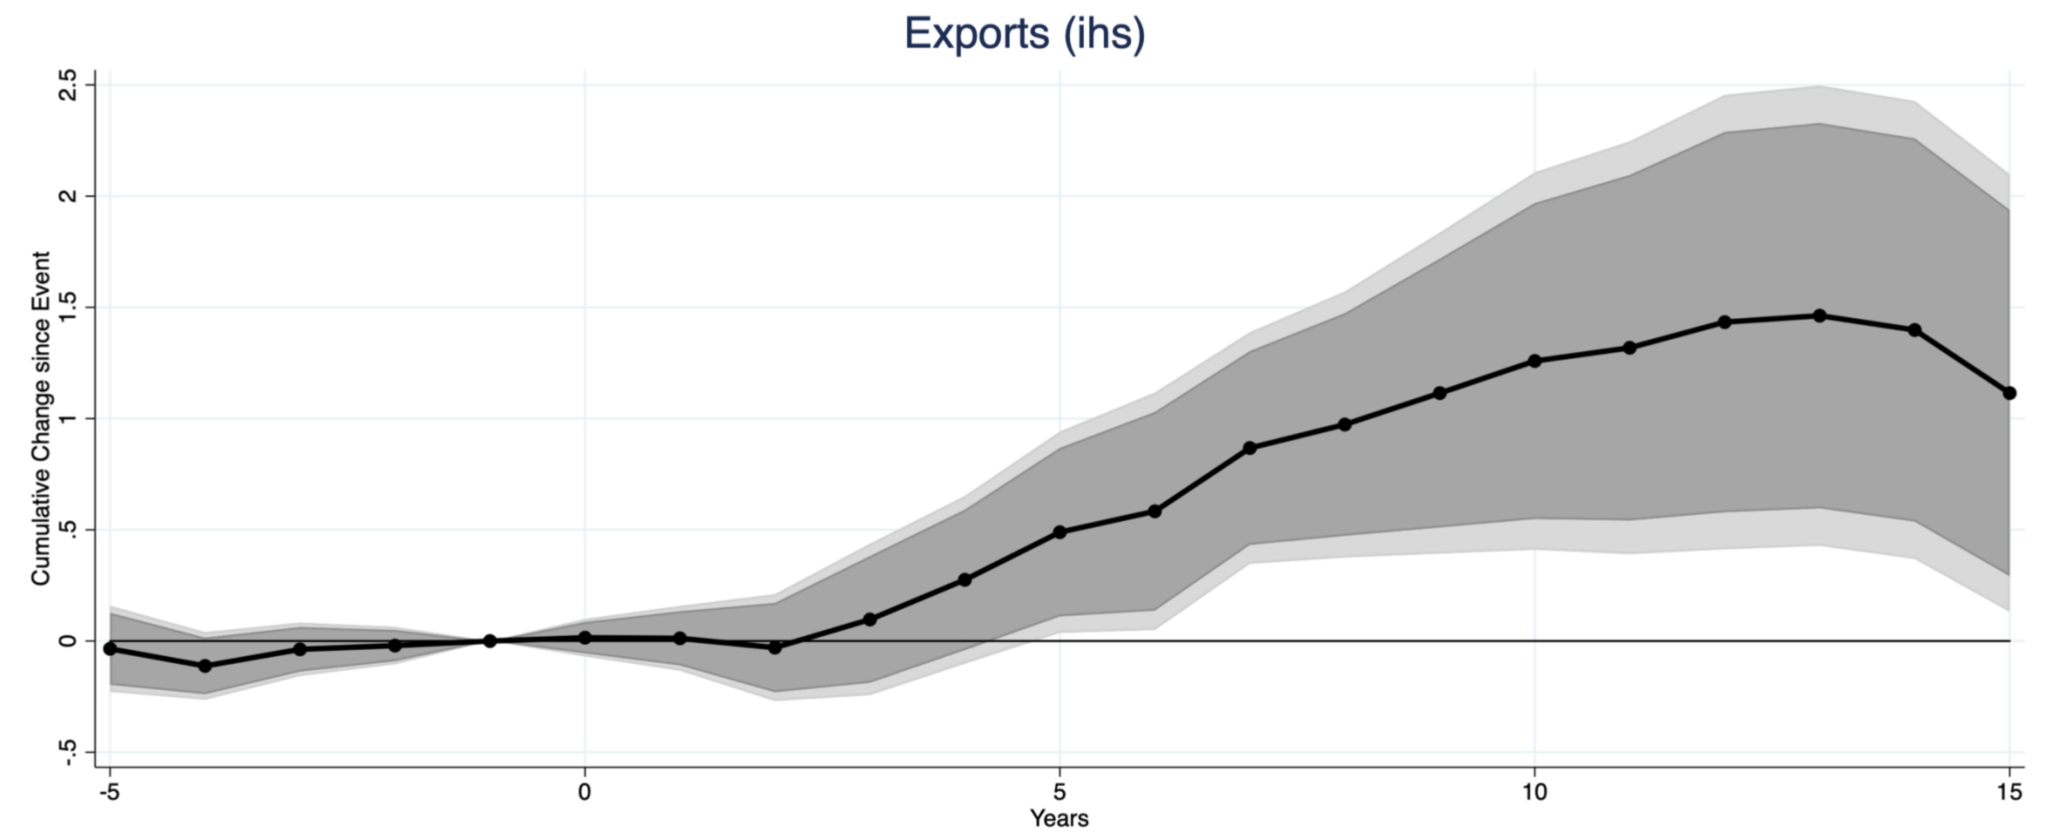

Note: 90% and 95% confidence bands displayed

Figure 5

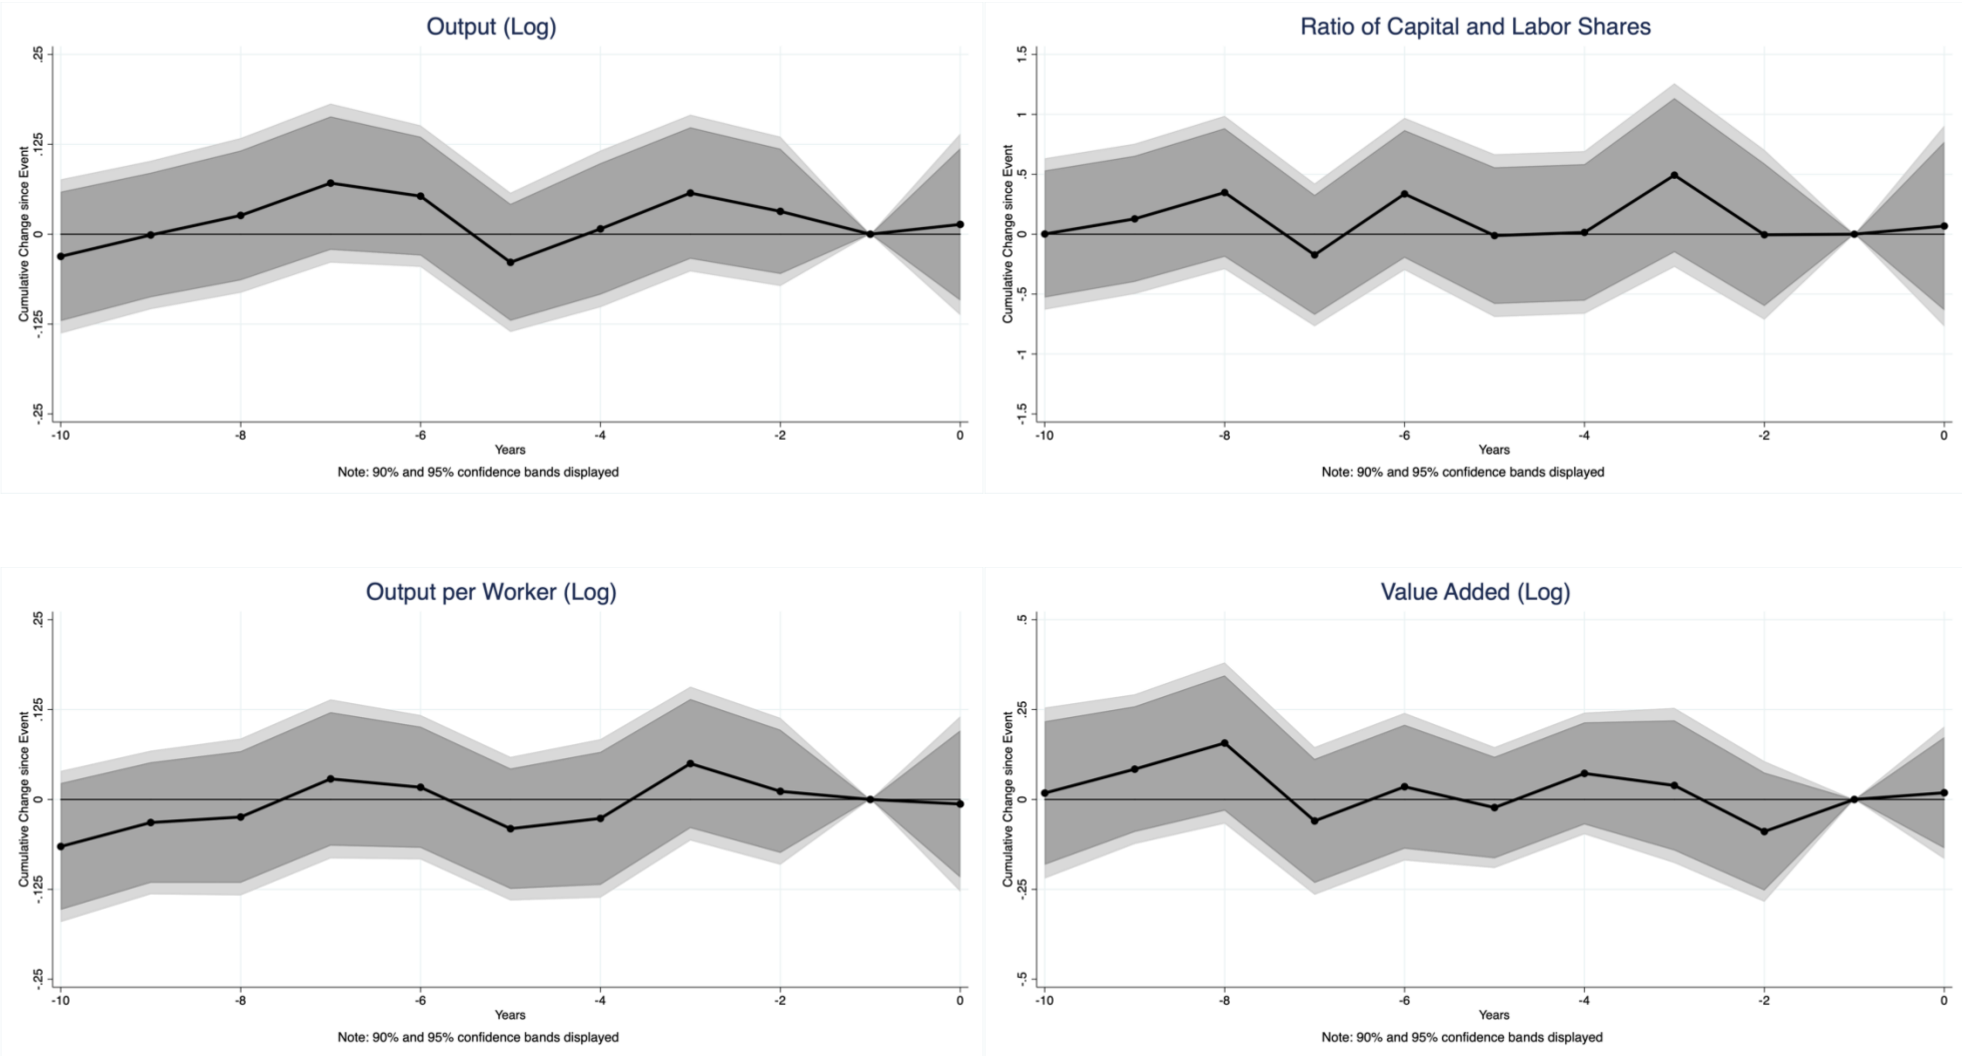
Output (Log)
Ratio of Capital and Labor Shares
Output per Worker (Log)
Value Added (Log)
Cumulative Change since Event
Years
Note: 90% and 95% confidence bands displayed
.25
.125
0
-.125
-.25
1.5
1
.5
-.5
-1
-1.5
-10
-8
-6
-4
-2

Figure 6

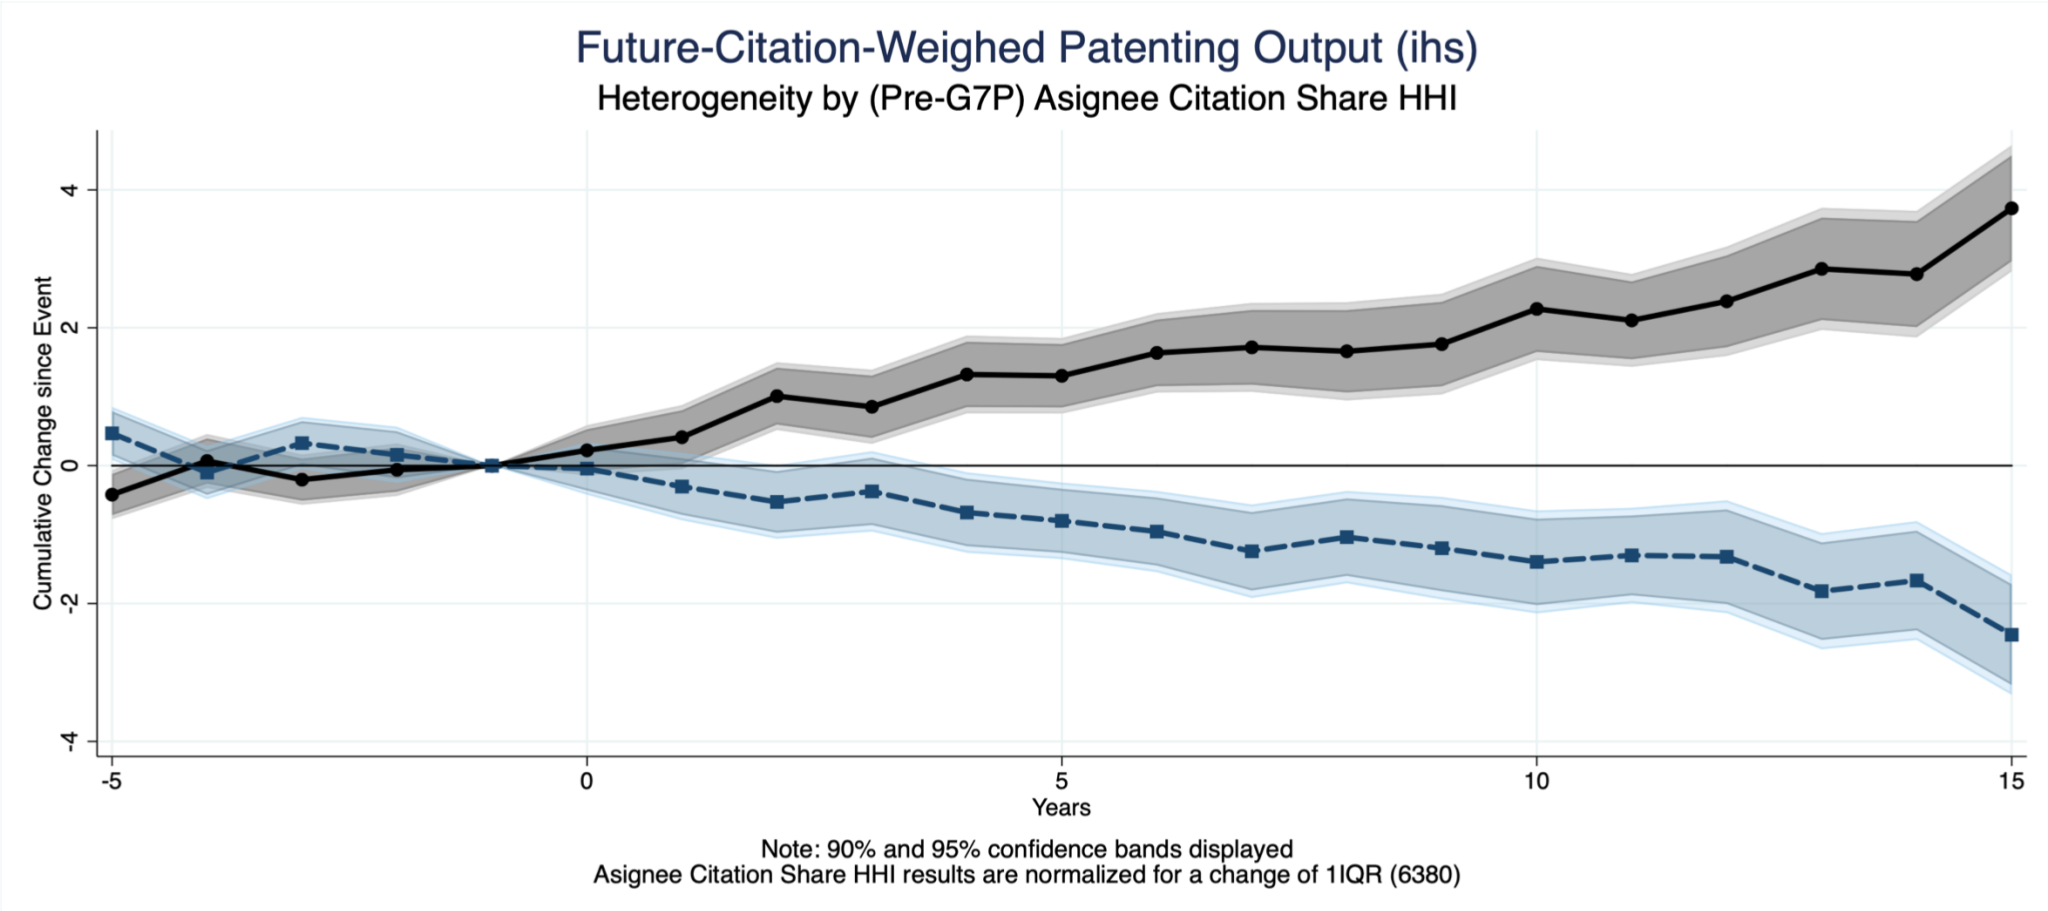

Future-Citation-Weighed Patenting Output (ihs)
Heterogeneity by (Pre-G7P) Asignee Citation Share HHI
Cumulative Change since Event
4
2
0
-2
-4
-5
0
5
10
15
Years
Note: 90% and 95% confidence bands displayed
Asignee Citation Share HHI results are normalized for a change of 1IQR (6380)

Figure 7
Cross-Country Regressions

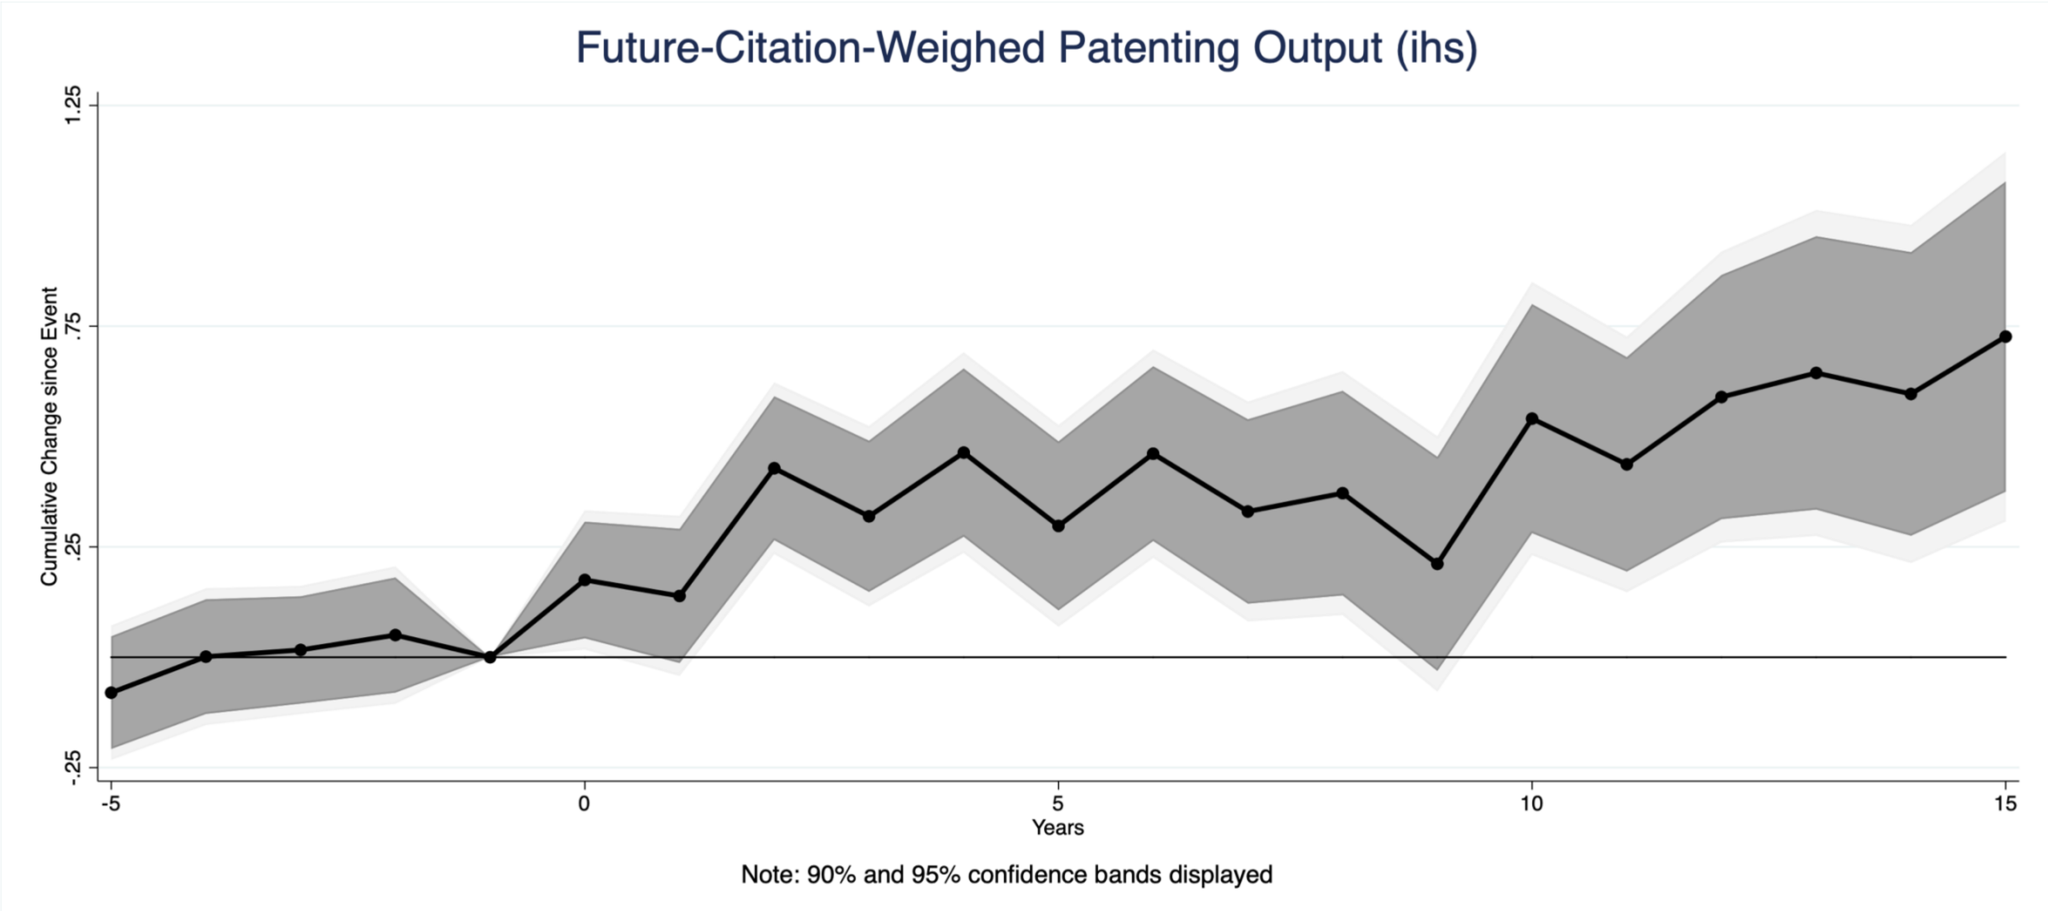

Figure 8
Cross-Country Regressions

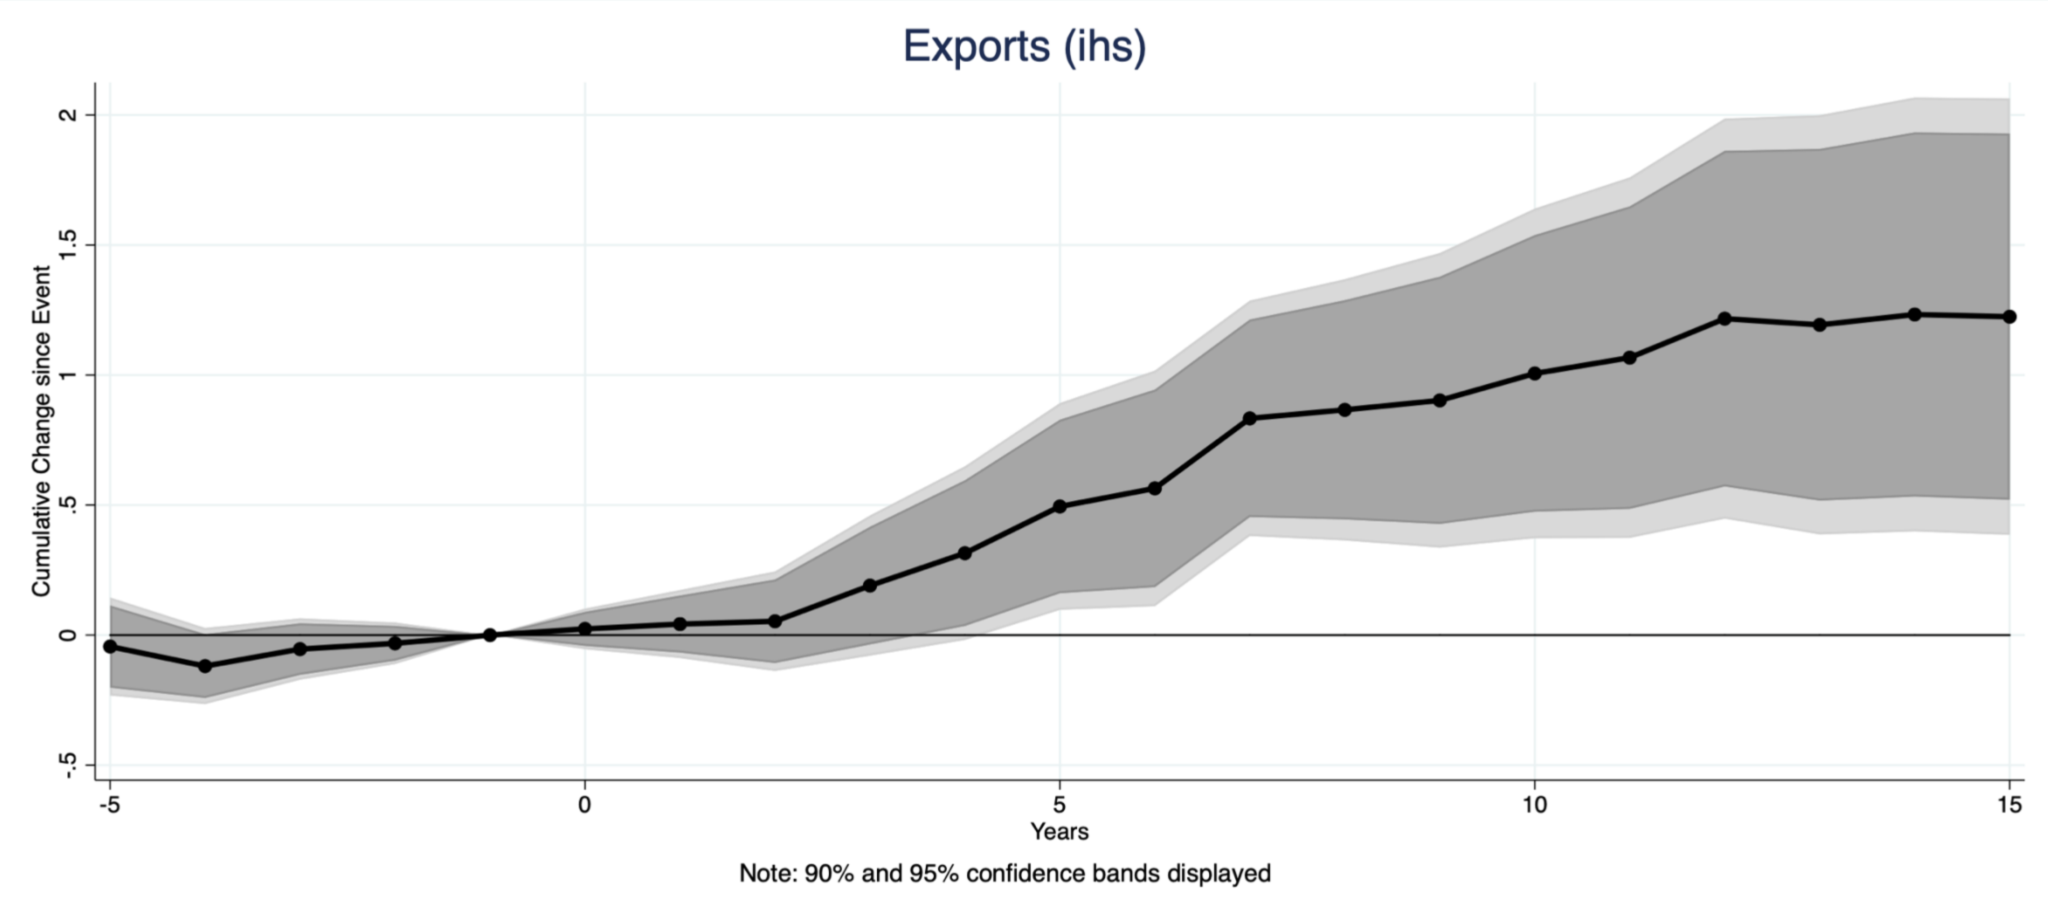

Tables

Table 1
List of G7 Projects

Type	Project Name	Implementation Period
Product	HDTV	1992 - 1994
	High-capacity semiconductor	1995 - 1999
	Next-generation car (electric vehicle)	1992 - 2001
	Next-generation flat panel display	1995 - 2000
	B-ISDN - Broadband Comprehensive Information and Communication Network Devices for 10GB environments	1992 - 2001
	New medicines and agrochemicals	1992 - 1997
	Medical Engineering	1995 - 2001
	Ultra-compact precision machinery	1995 - 2001
	High-speed train	1996 - 2001
Base	Advanced energy and informatic materials	1992 - 2001
	New functional biomaterials	1992 - 2001
	Advanced production system	1992 - 2001
	Next-generation semiconductor	1993 - 1996
	Environmental engineering	1992 - 2001

	New energy (Fuel-cell)	1992 - 2001
	Next-generation nuclear reactor	1992 - 2001
	"Emotional" Engineering	1995 - 2001
	Next-generation superconducting nuclear fusion device	1995 - 2001

<u>**Appendix**</u>

Robustness Checks

We assess the robustness of our findings to alternative definitions of patent nationality, logarithmic transformation of the dependent variable, and alternative quality thresholds in our language model exercise.

Figure A1

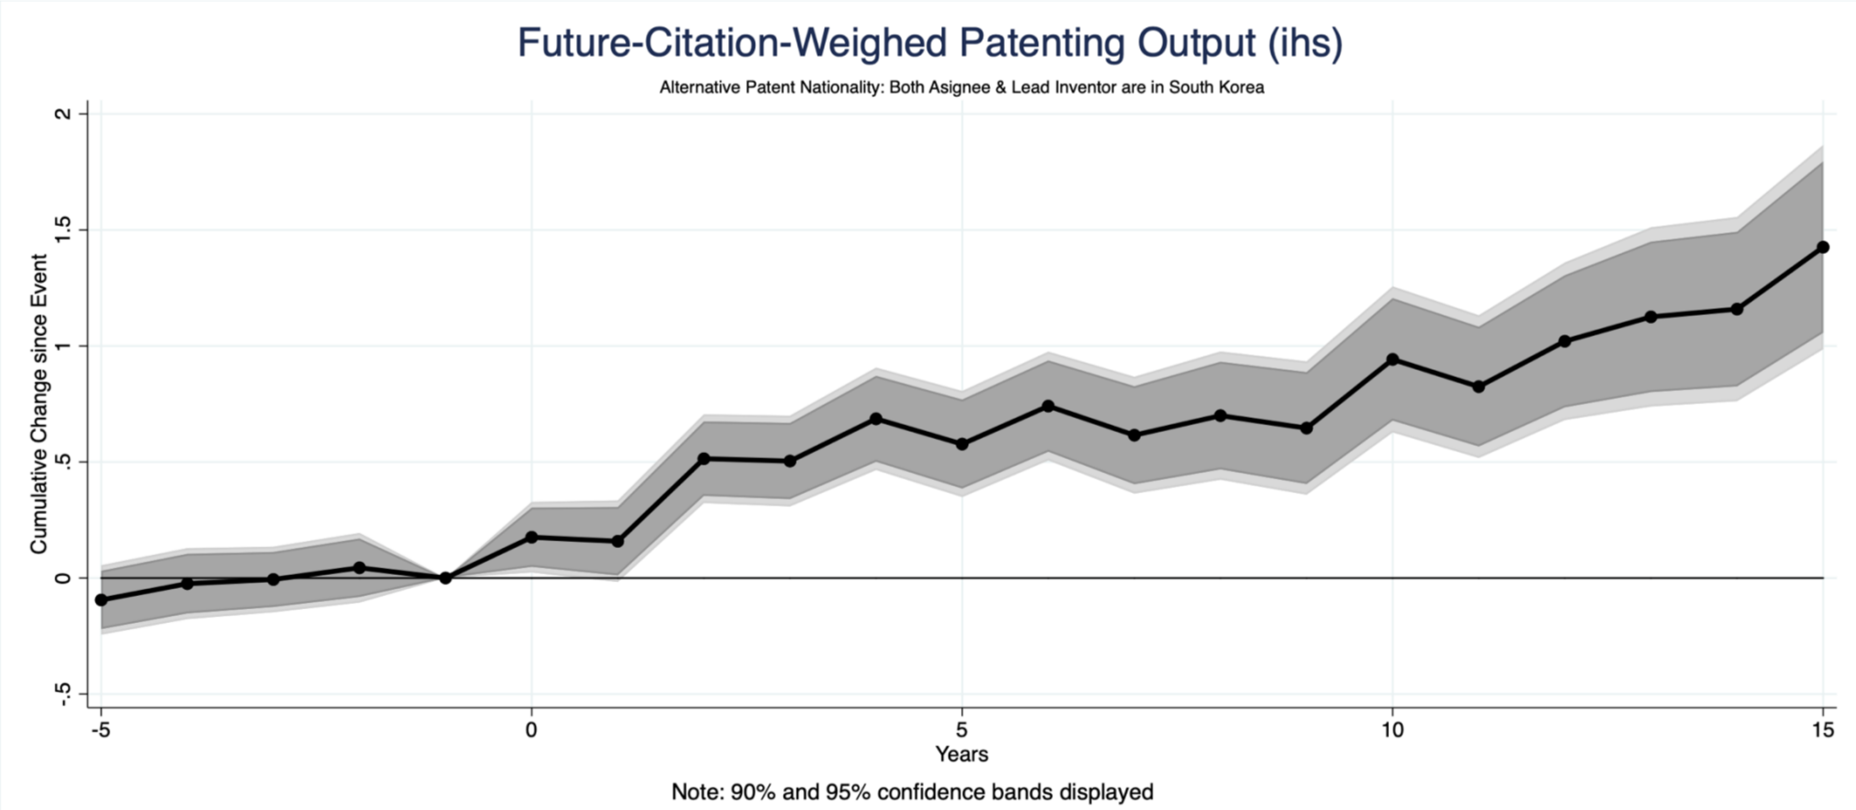

Figure A2

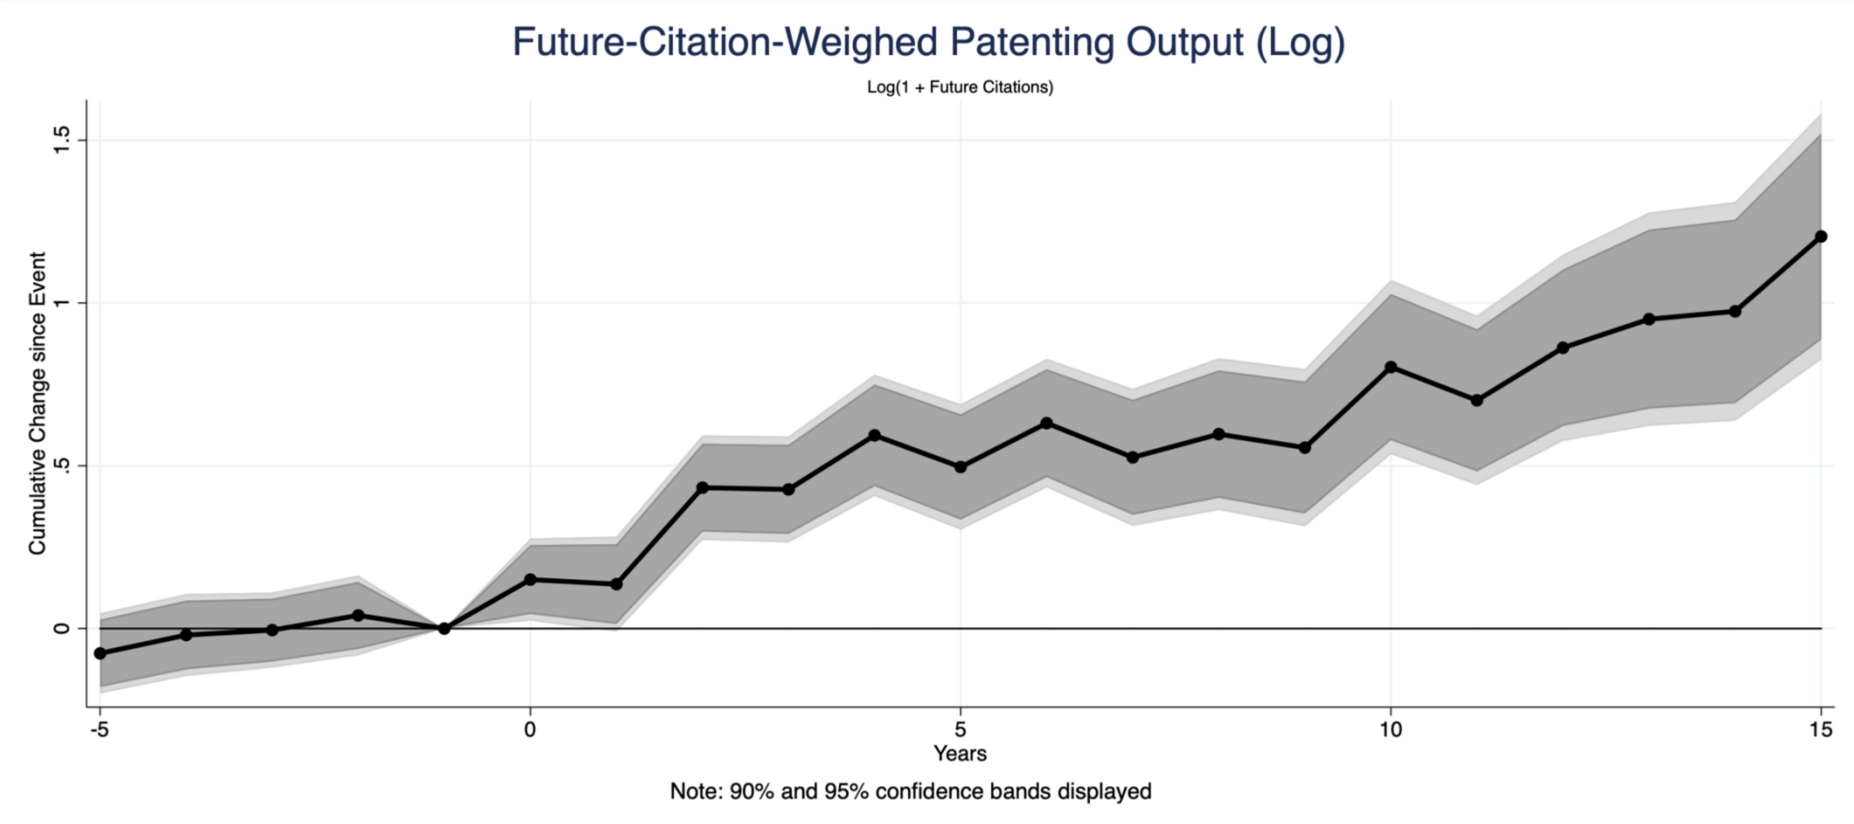
Future-Citation-Weighed Patenting Output (Log)
Log(1 + Future Citations)
Cumulative Change since Event
1.5
1
.5
0
-5
0
5
10
15
Years
Note: 90% and 95% confidence bands displayed

Figure A3

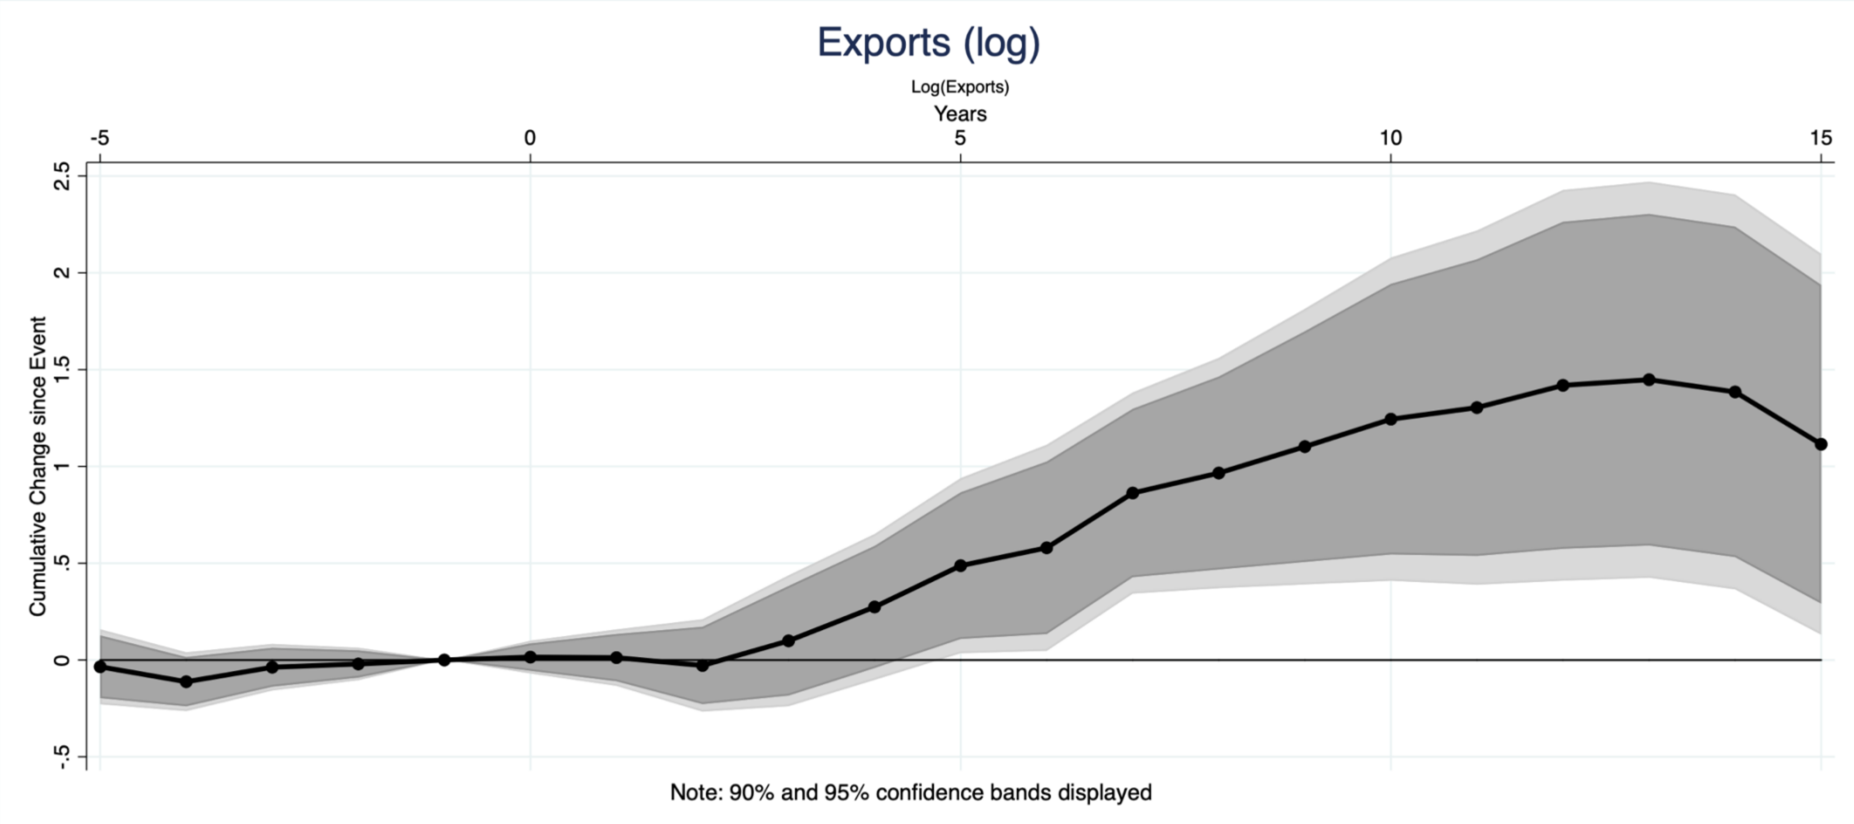

Exports (log)
Log(Exports)
Years
-5
0
5
10
15
Cumulative Change since Event
2.5
2
1.5
1
.5
0
-.5
Note: 90% and 95% confidence bands displayed

Figure A4

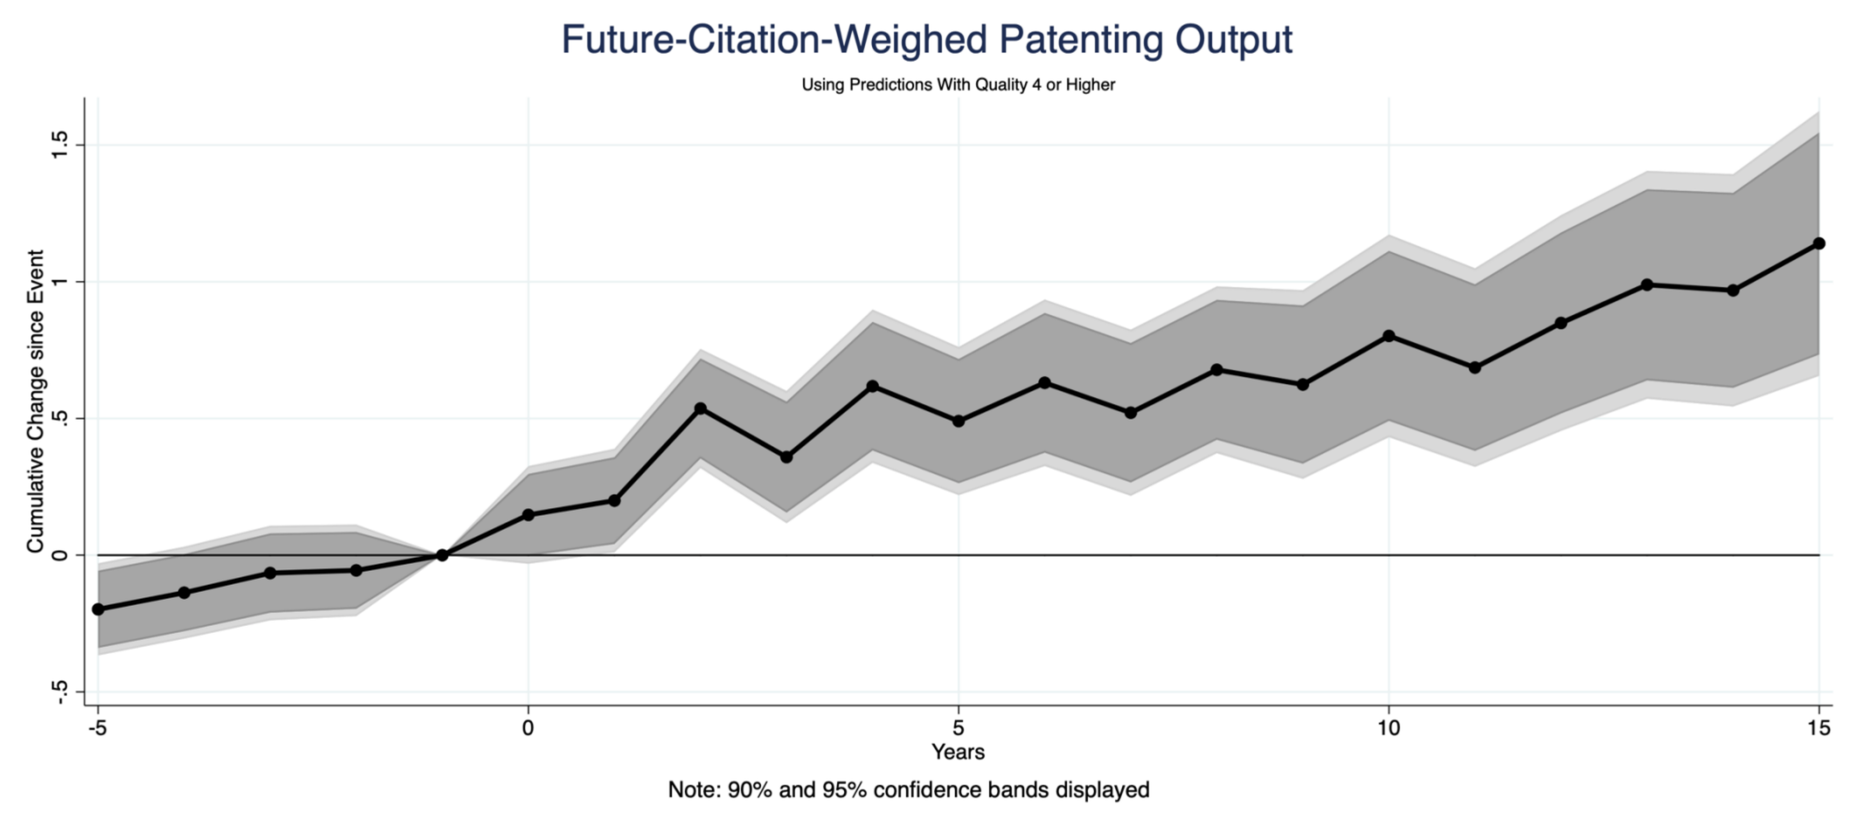
Future-Citation-Weighed Patenting Output
Using Predictions With Quality 4 or Higher
Cumulative Change since Event
-.5
0
.5
1
1.5
-5
0
5
10
15
Years
Note: 90% and 95% confidence bands displayed

Figure A5

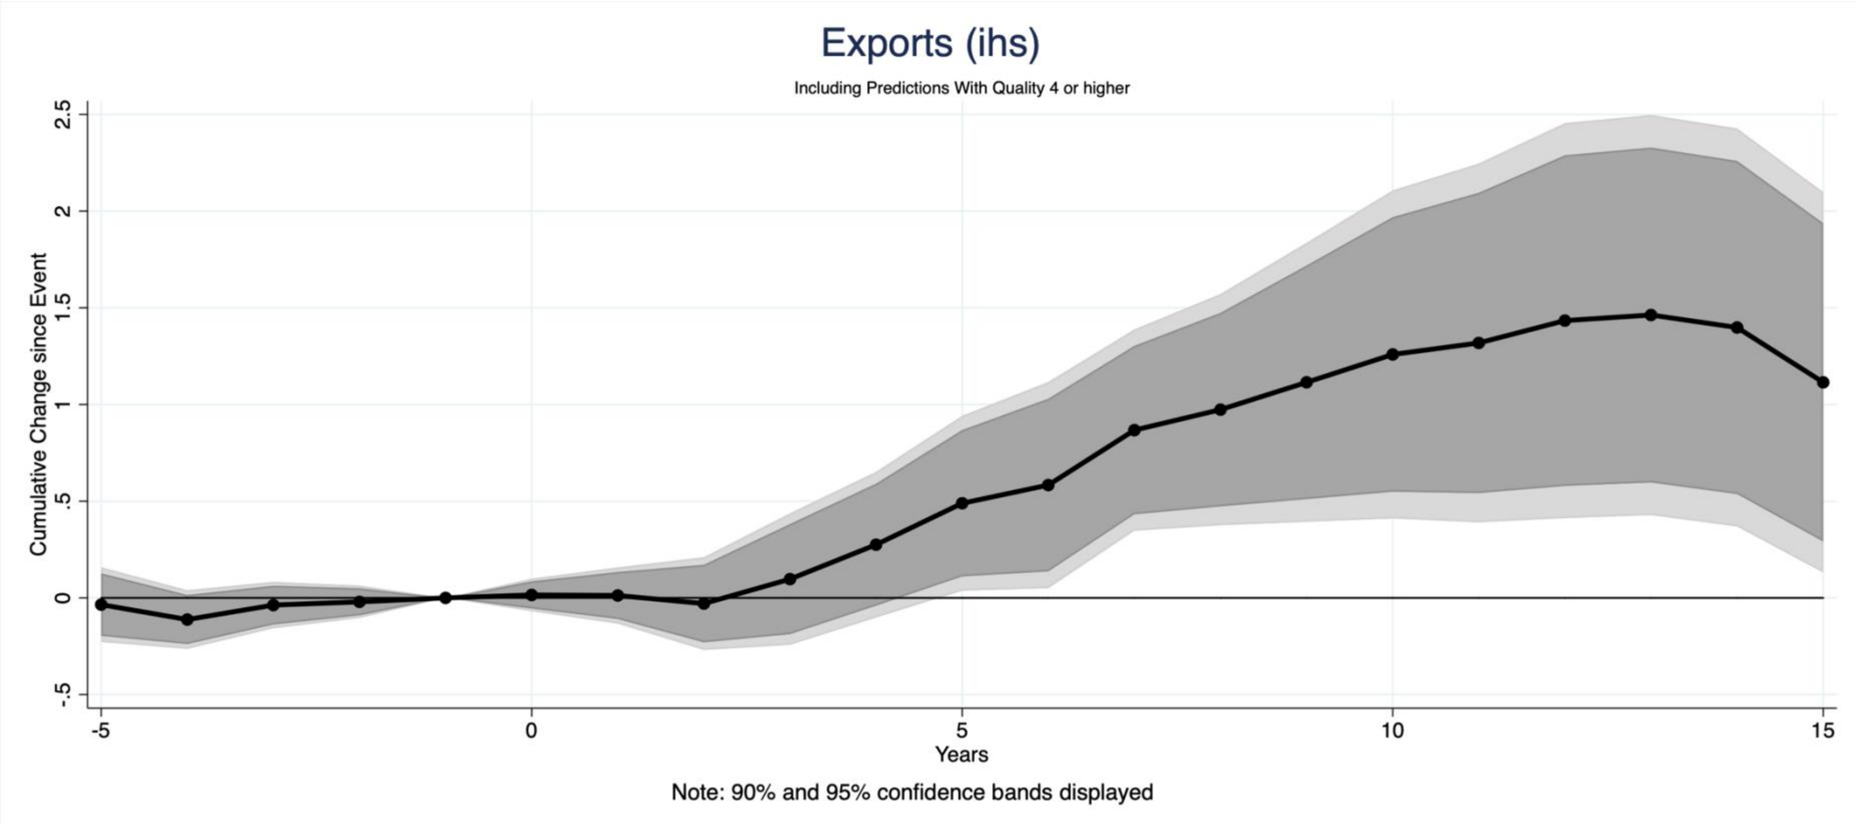

Exports (ihs)
Including Predictions With Quality 4 or higher
Cumulative Change since Event
2.5
2
1.5
1
.5
0
-.5
-5
0
5
10
15
Years
Note: 90% and 95% confidence bands displayed

Figure A6

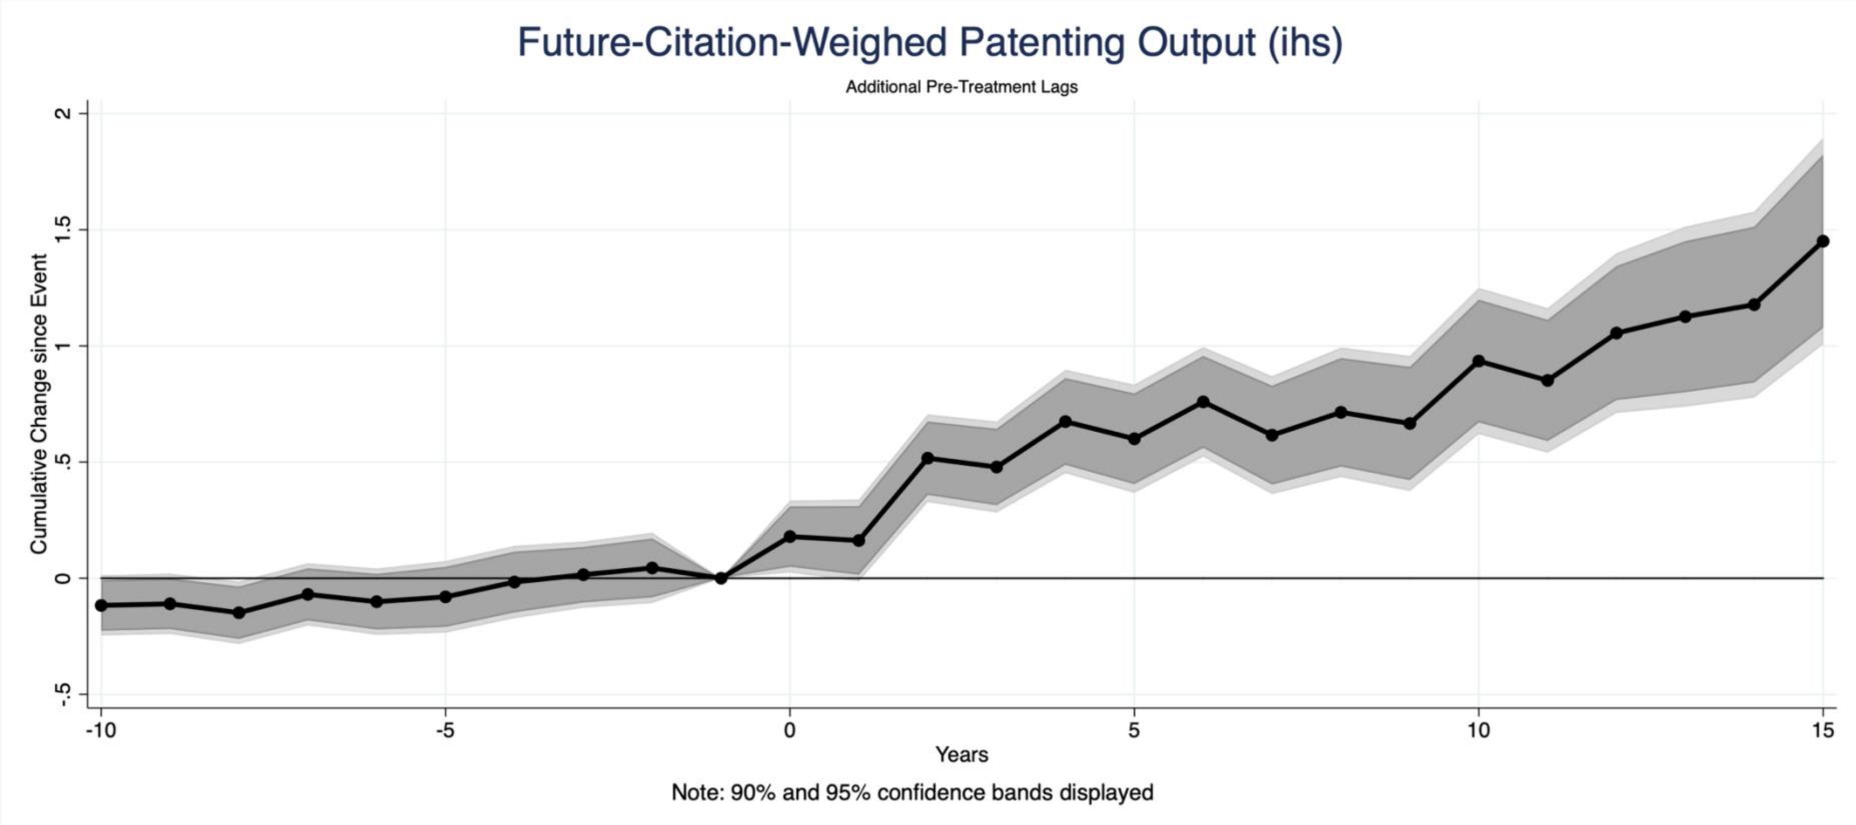
Future-Citation-Weighed Patenting Output (ihs)
Additional Pre-Treatment Lags
Cumulative Change since Event
-.5
0
.5
1
1.5
2
-10
-5
0
5
10
15
Years
Note: 90% and 95% confidence bands displayed

Figure A7

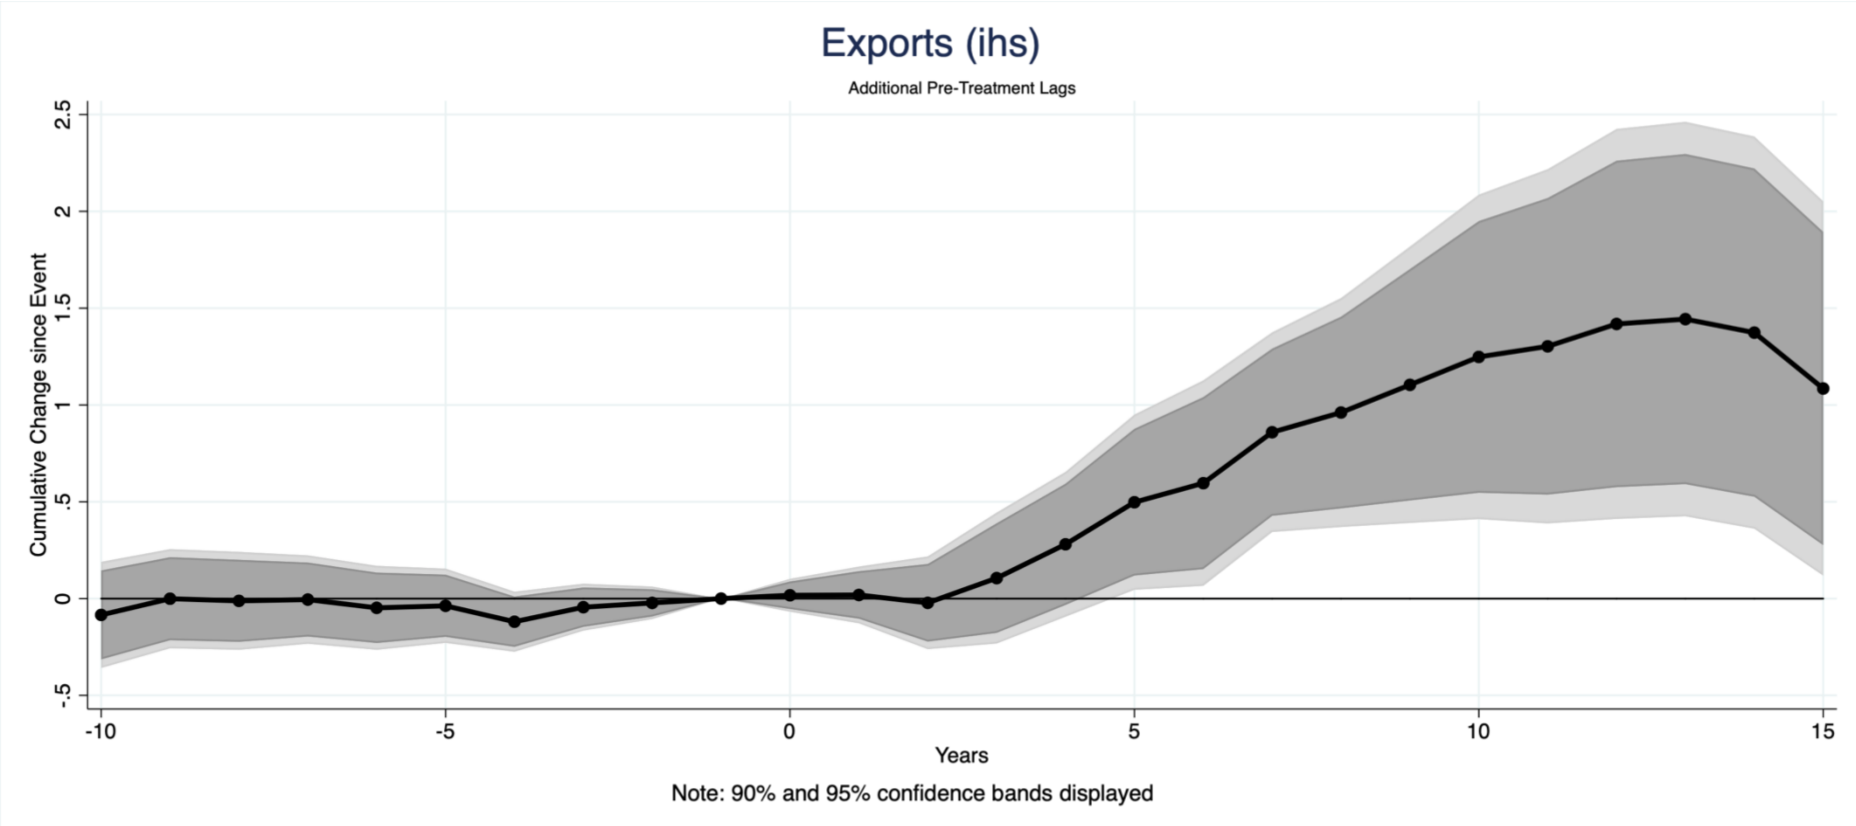

Exports (ihs)
Additional Pre-Treatment Lags
Cumulative Change since Event
2.5
2
1.5
1
.5
0
-.5
-10
-5
0
5
10
15
Years
Note: 90% and 95% confidence bands displayed