Bureaucrats and the Korean Export Miracle *

Philipp Barteska
LSE

Jay Euijung Lee
Stockholm University

Job Market Paper

November 24, 2023

Please click here for the latest version

Abstract

What makes an industrial policy successful? This paper finds that the effect of an industrial policy changes tremendously with the implementing bureaucrat. We study Korean bureaucrats who promote exports on appointments to 87 countries between 1965, when Korea was one of the world's poorest countries, and 2001. We exploit the rotation of bureaucrats between countries to show that individual bureaucrats matter greatly in boosting exports. Moving from a bureaucrat at the 20th percentile to the median is associated with a 40% increase in exports. This effect is comparable to that of opening an office, implying that this industrial policy has no effect under a 20th percentile bureaucrat. We exploit differential import demand growth to study a mechanism via which better bureaucrats increase exports: transmitting information about market conditions. Under better bureaucrats Korean exports increase more with a product's import demand. Finally, we investigate whether experience can bridge the gaps between bureaucrats. We isolate quasi-random variation in experience exploiting a product's import demand growth during the bureaucrat's first appointment. In subsequent appointments exports increase in products with greater bureaucrat experience. This highlights learning-by-doing as a channel to build bureaucratic capacity. However, the differences between bureaucrats are larger than the effect of experience, suggesting selecting good bureaucrats may be more important than training them.

JEL classification: L52, F13, D73, O11, O12

Key Words: Industrial Policy, Bureaucracy, Economic Development, Managers, Government, Political Economy, Export Promotion, Trade Policy

* Barteska: LSE (email: p.barteska@lse.ac.uk, website: https://www.philippbarteska.com);
Lee: Stockholm University, (email: jay-euijung.lee@su.se).

This paper benefited from comments from Tim Besley, Mike Callen, Jonas Hjort, Nathan Lane, Rocco Macchiavello, Torsten Persson, Daniel Sturm, James Robinson, Jose Vasquez, Noam Yuchtman. We would like to thank audiences at the London School of Economics, Warwick, the IFS, NEUDC (HKS), the Yale-IGC firms, trade, development conference. Philipp Barteska would especially like to thank his advisors: Oriana Bandiera, Robin Burgess, and Gharad Bryan. This study was made possible with the excellent research assistance of Minjeong Kang, Sunae Choi, Gippeum Lee. We would also like to thank the staff at the KOTRA archives, in particular Hyunsun Cho. Financial support from STICERD and the Hayek programme is gratefully acknowledged.

1 Introduction

State and bureaucratic capacity are strongly associated with economic development (Besley, Burgess, Khan, and Xu (2022) – BBKX). Less is known about *how* bureaucratic capacity causes economic growth. Explanations of Asia's growth miracles suggest one channel: bureaucracies are central to industrial policy success (Juhász, Lane, and Rodrik (2023) – JLR)[1]. Understanding to what extent the effect of industrial policy depends on bureaucratic capacity is crucial in determining what lessons low- and middle-income countries can draw from development success stories such as South Korea.

In this paper, we make two contributions. First, we provide evidence that the effect of an industrial policy on economic development crucially depends on bureaucratic capacity. Second, we show that learning-by-doing can build bureaucratic capacity.

Investigating whether capacity impacts the effect of industrial policy has been difficult because doing so requires a setting that satisfies the following conditions: First, we need variation in bureaucratic capacity while holding constant the policy. This condition may be satisfied if a national policy is implemented decentrally across locations. Second, this capacity needs to vary while holding constant the location, whose economic conditions may directly impact the policy's effect. Such variation may occur when the bureaucrats implementing a policy move between locations. Third, enough bureaucrats need to move so that locations and bureaucrats form large connected sets, ideally one connected set containing all locations and bureaucrats. Fourth, the mapping from bureaucrats to the policy's effect is one-to-one if there is no multi-tasking: This is satisfied if each bureaucrat only works on this policy, and the policy's outcome is measurable in each location – ideally, this outcome is closely linked to economic growth.

To satisfy these conditions we pick an appropriate context: South Korean overseas export promotion. First, this policy was implemented decentrally in 87 different destination countries. Second, bureaucrats rotate between countries every three years, giving us poten-

[1]Narrative accounts of the rapid economic growth in East Asia emphasize the positive role of industrial policy and the development of state capacity for carrying out complex policies, in particular in South Korea (Amsden, 1989; Wade, 1990; Evans, 1995; Woo-Cumings, 1999). At the presence of market failures, such as production externalities, agglomeration failures, and public provision of production inputs, the state needs to intervene for firm growth by enacting industrial policy (JLR).

tially exogenous variation in the implementing capacity within location. Third, the largest connected set includes 86 of 87 countries due to the frequent movement of bureaucrats and our long period of study (1965-2001). Fourth, in each country the policy has a sole target: exports to that country, an important development outcome[2]. This setting also is of substantial intrinsic interest: Korea may be the most prominent example of a low-income country to reach high income. Exports were a key target of its policies and Korea's growth in exports is a particularly remarkable phenomenon to be explained.

We first estimate the effect of opening an overseas office, conducting an event-study estimation which uses the offices' staggered roll-out. We estimate an increase in exports of 38% compared to a never-treated control group after ten years. Assuming an elasticity of trade to distance of -1, a similar increase in exports would occur if the distance between London and Seoul was reduced to that between Mumbai and Seoul. We are able to rule out the two most plausible alternative interpretations for the event-study findings: (1) Countries do not experience an increase in demand, measured as non-Korean exports, after an office opens. (2) The scope for strategic timing of office openings is limited as pre-determined variables explain the year in which a country's office opens.

Second, we show that the effect of this policy differs depending on the bureaucrat assigned to a country. We use a movers design in a two-way fixed effects framework that exploits the regular rotation of bureaucrats to offices. (1) Moving from the 20th percentile bureaucrat to the median increases exports by 39%, after applying a standard shrinkage procedure to the estimated bureaucrat effects. (2) This indicates that an office opening has no effect under a 20th percentile bureaucrat. We obtain similar-sized estimates of the effect on exports when moving from the median to the 90th percentile (38%) or when increasing the bureaucrat effect by one standard deviation (37%). The latter estimate is obtained via a variance decomposition that uses a leave-out estimator to correct for a limited mobility bias under

[2]Exporting is important for economic growth and development more broadly. For evidence highlighting the effect of exports on development outcomes at the firm-level, see Atkin, Khandelwal, and Osman (2017). For evidence at the macro-level, see Hausmann, Hwang, and Rodrik (2007); Atkin, Costinot, and Fukui (2021) For support that demand-side factors may be decisive in economic development, see Goldberg and Reed (2020). Exporting remains central to many sectoral industrial policies. Lederman, Olarreaga, and Payton (2010) report that more than 100 countries have an export promotion agency comparable to the Korean one. For evidence on the centrality of export promotion to sectoral industrial policy, see Juhász, Lane, Oehlsen, and Pérez (2022).

heteroskedasticity (Kline, Saggio, and Sølvsten, 2020). We are able follow this approach due to the unusual interconnectedness of our data. We conduct numerous diagnostic checks, including event-study estimates, to alleviate concerns regarding our identifying assumptions.

Third, we provide evidence that bureaucrat experience increases Korean exports. This points to learning-by-doing as a channel to build bureaucratic capacity. But it also highlights a novel channel for path dependence in this capacity. We isolate quasi-random variation in bureaucrat experience: a product's import demand growth during the bureaucrat's first appointment. Event-study estimates around a switch in bureaucrats indicate that exports increase by 3.0% when the quasi-random component of product-specific experience increases. In isolating this quasi-random component of experience, we address three main sources of endogeneity in correlations of bureaucrat experience and exports. The increase in exports is equivalent to reducing the distance between London and Seoul to that between Frankfurt and Seoul.

We investigate a mechanism for the increases in exports due to both greater bureaucrat fixed effects and bureaucrat experience. Relaying information about market conditions is among the overseas offices' main tasks. We show that, upon the appointment of a better bureaucrat, exports go up partly due to an increased elasticity of Korean exports to market conditions – e.g., exports of a product to a country increase more strongly when there is increased import demand. This suggests that better bureaucrats more effectively relay information about market conditions.

We foremost shed light on the oft-hypothesized but under-researched link between state capacity and industrial policy (IP) (Juhász, Lane, and Rodrik, 2023). We study bureaucrats who implement an IP. By showing how important this one determinant of implementation is to IP success, we inform current policy debates on the circumstances required for successful IP (Juhász, 2018; Liu, 2019; Lane, 2022; Shim and Choi, 2022; Choi and Levchenko, 2021). How to learn from Korean IP matters as IP is widespread across developing countries and developed countries, with export promotion (EP) often forming an important component of sectoral IP (Juhász et al., 2022). Studying how bureaucrats affects export promotion provides a link between research on state capacity and research on firm performance in developing countries, in particular as these relate to demand-side shocks and EP more specifi-

cally (among many others: Atkin, Khandelwal, and Osman (2017); Alfaro-Ureña, Manelici, and Vasquez (2022), reviewed by Atkin and Donaldson (2022), Atkin et al. (2022); on EP, see Munch and Schaur (2018); Volpe Martincus and Carballo (2008, 2010, 2012)). Our paper studies how bureaucratic capacity shapes the effect of a policy which may alleviate demand-side constraints that may hamper development in many countries (Goldberg and Reed, 2020).

By more closely linking bureaucrats to an outcome important to economic growth, we contribute to understanding the bureaucratic determinants of economic growth (BBKX). We methodologically relate to research that finds substantial effects of managers and individual workers on the performance of organizations (Fenizia, 2022; Best, Hjort, and Szakonyi, 2023; Otero and Muñoz, 2022; Metcalfe, Sollaci, and Syverson, 2023) as well as effects of teachers on student test scores (e.g., Chetty, Friedman, and Rockoff, 2014a,b) by applying methods from the labor literature on worker and firm heterogeneity (Abowd, Kramarz, and Margolis, 1999; Abowd, Creecy, and Kramarz, 2002; Card, Cardoso, and Kline, 2016; Card, Heining, and Kline, 2013; Kline, Saggio, and Sølvsten, 2020). We contribute to this literature by linking individual bureaucrats to industrial policy and exports. Moreover, we shed light on a previously understudied mechanism for increasing state and organizational capacity by showing that bureaucrats gain capacity via learning-by-doing. As our bureaucrats are managers, this finding is informative about managers in other organizations.

The rest of the paper proceeds as follows. Sections 2 describes the institutional background. Section 3 introduces the data. Section 4 discusses the effect of office openings. Section 5 shows how much industrial policy depends on individual bureaucrats. Section 6 focuses on experience as one factor determining differential effectiveness between bureaucrats. Section 7 concludes.

2 Institutional Background

Our study commences at a time when South Korea was one of the world's poorest countries. During our period of study, Korea's real GDP per capita increased from $1,304 (1961) to

$25,421 (2001).[3] In 1961, the average income in Korea was below most countries in Sub-Saharan Africa.[4] In 2001, Korea's average income was above Portugal's. This growth is prominently attributed to a well-functioning, activist state that conducted successful industrial policies.[5] On the other hand, the Korean state was described as aid-dependent and corrupt until at least the mid-1960s (Kim and Vogel, 2011).[6] This makes Korea an interesting case for understanding the role of state capacity in economic development broadly and the implementation of industrial policy.

Figure 1 highlights Korea's growth of exports per capita between 1952 and 2001. Exports per capita in 1952 were below 2% of the U.S. level with little convergence between 1952 and 1960. From 1960 on, exports increased rapidly, reaching parity with the U.S. before the end of the century. This paper sheds light on this transformative growth in exports, central to narratives of South Korea's broader economic miracle. Export promotion as a prominent area of state activism is highlighted by a representative survey of Korean manufacturers in 1976 (Jones and Il, 1980). These manufacturers reported "foreign marketing" as the policy area where government intervention most markedly improved under President Park Chung-hee (1961-1979), compared to President Syngman Rhee (1948-1960).

2.1 KOTRA: Tasks and Outputs Produced

We study the overseas offices of Korea's Trade Promotion Agency (KOTRA) founded in 1962. At its inception, KOTRA was tasked with "promot[ing] the increases of exports. In order to accomplish this goal, its functions include sales promotion and research, a campaign of public relations and advertising, [and] information service to exporters and importers" (Udell, 1965). Figure 2 displays the number of countries with an overseas KOTRA office

[3]Both in 2017 USD. Relative to the U.S., this corresponds to an increase from 1/15 of real GDP per capita in 1961 to 1/2 in 2001. Data from Penn World Tables.

[4]The countries with higher GDP per capita in 1961 in Sub-Saharan Africa in order of 2023 population: Nigeria, the Democratic Republic of the Congo, South Africa, Kenya, Ghana, Madagascar, Côte d'Ivoire, Cameroon, Niger, Zambia, Chad, Senegal, Zimbabwe, Guinea, Benin, Togo, Republic of the Congo, the Central African Republic, Liberia, Mauritania, Gambia, Namibia, Gabon, Mauritius, the Comoros, Cape Verde, the Seychelles.

[5]Wade (1990) and Cheng et al. (1998) as cited by BBKX; Amsden (1989); Juhász, Lane, and Rodrik (2023). See also the well-known popular book by Studwell (2013).

[6]Korea's level of state capacity may be highlighted by the lack of continuity in its ministries. Between 1948 and 1960, under President Rhee, the average agriculture minister lasted just 9 months. The average commerce minister lasted 13 months (Haggard, Kim, and Moon, 1991).

over time. Offices opened in 28 countries until 1970, 68 until 1980, and 82 until 2000.

The overseas offices contributed to three main functions of KOTRA that were maintained consistently from the early years of the organization's establishment. First, KOTRA's "Investigation/Research" division investigated factors related to export supply and demand: (1) Korea's capability to supply a product for exports and (2) the import demand in the foreign market. The overseas offices produced reports by product and country that were compiled and published by the head office. Second, the overseas offices served a key role in the "Market development" division by helping domestic producers and retailers find new trade partners in new and existing markets. They received export inquiries from domestic companies and import inquiries from foreign ones, which got published in KOTRA's Daily Market Newspaper. Business transactions were then mediated between the inquirers and respondents. Third, the overseas offices helped the "Trade fair" division with the organization of a Korean pavilion at international trade fairs, which were viewed as a means to produce great export results within short periods of time by allowing exporters to engage in direct conversations with local buyers. To assist with this, the overseas offices coordinate logistics. They also recruit, select, and brief exporters who will represent their products at the fairs. At the same time, they disseminate information about these exporters and their products to attract potential buyers to the Korean pavilion or individual firms. The bureaucrats did this by running ads, sending letters and making phone calls to promising exporters and foreign buyers, and reaching out to trade associations. The selected domestic companies often produced goods with newly trending styles and designs that matched the marketability of the venues of the fairs.

Each of these three functions are aptly captured by our data on KOTRA office activity described in section 3.3. The data include market reports investigating export capability and import demand, importer requests, and sales and attendance of firms at KOTRA-organized trade fair pavilions.

Compared to other bureaucracies, KOTRA's overseas offices have a large degree of discretion regarding how to carry out the task of promoting exports. For this reason, this paper's main results focus on KOTRA's ultimate outcome of interest: exports. Clearly, it is difficult to centrally plan whether exports to a particular destination will benefit more from market

reports or networking with potential importers, and whether networking should happen via attending fairs, phone calls, or some other channel. Instead, such a goal relies on the bureaucrats' knowledge, which may be both tacit and local, and requires substantial improvisation. So rather than having a centrally mandated list of tasks to fulfill (as in Bandiera, Best, Khan, and Prat (2021); Fenizia (2022); Best, Hjort, and Szakonyi (2023)), KOTRA office directors are somewhat more like the proverbial "man on the spot" charged with the running of an entire geographic region in the Indian Administrative Service (Bertrand, Burgess, Chawla, and Xu, 2020) or the British colonial administration (Lugard (1926), as cited by Xu (2018)). However, compared to someone in charge of an entire region, KOTRA bureaucrats have a more narrowly defined target that can largely be summarized into the measure of exports during their appointments. The primary performance measure, as assessed by the head office, is whether export targets are met. This makes studying KOTRA bureaucrats much less susceptible to the multi-tasking problem faced by studies evaluating the effectiveness of most bureaucrats with regional responsibilities. Moreover, other than existing studies, the outcome targeted by overseas offices, exports to the country, is an outcome of direct importance for economic growth and development.

2.2 KOTRA: Assignment to Overseas Offices

Official rules do not dictate which bureaucrat gets assigned to which office. The assignment system falls under the discretion of the HR Team at the head office. According to interviews we conducted with current and former KOTRA employees, however, there is a general understanding that several factors come into play. The most important factor is language skills; a Spanish speaker is deemed more likely to get sent to a Hispanophone country. Second, if one has worked previously at an office in a junior position, they might get assigned to the same office as a director in the future. Third, given the 2.5- to 3.5-year rotation schedule, a bureaucrat can only be sent to offices that become vacant at the right time. Fourth, an officer who got posted to an undesirable location, such as a small, low income country far from Korea, might be compensated by getting posted to a desirable location next. Lastly, connections with KOTRA executives might matter for assignments to desirable locations.

Over the entire time period from 1962 to 2001, KOTRA operated 138 overseas offices in

87 countries[7], with the most important or geographically largest countries having multiple offices in different cities[8]. Most analysis will focus on the main country offices as outcomes are available at the country level. The regular nature of these directors' appointments is highlighted by the fact that both the modal and median appointment duration is 36 months – three years. Appendix figure A.1 plots the distribution of appointment durations. Between appointments, managers return to Korea, typically at KOTRA's headquarters in Seoul and sometimes at regional offices. The timing of their re-appointment is also largely pre-determined: The median duration for the gap between appointments is 29 months, the modal gap is 30 months. Appendix figure A.2 plots the distribution of gaps between appointments.

2.3 KOTRA and Korea's Largest Scale IP

One reason for studying EP is the narrative of Korea's development as being export-driven, as well as EP's prominent role in Korean IP. Korea's largest scale IP, the Heavy and Chemical Industries drive (HCI), commenced in early 1973 and ended in October 1979..

To show the connection between EP and HCI, we linked about 45,000 of the reports written by KOTRA's overseas offices between 1965 and 2001 to the products or sectors discussed by each report. When discussing whether a product was treated by HCI, we follow Lane (2022), who included those "listed in the enforcement decrees and national sectoral acts underlying HCI". HCI's six broadly defined target sectors included steel, nonferrous metals, shipbuilding, machinery, electronics, and petrochemicals.

Appendix figure A.3 displays how the targeting of KOTRA's activity changed over time. Before the HCI drive, only 15-25% of product-specific reports discuss HCI products. During the HCI drive, this share increases rapidly, reaching close to half of all reports in the late 1970s. After the HCI drive, the share of reports targeting these sectors remains relatively constant. This supports the view that export promotion was used as part of Korea's overall industrial policy.

[7]For example, by 1977, KOTRA had 79 overseas offices, of which 64 were the respective country's head office.

[8]In Canada, a geographically large country, KOTRA has offices in Vancouver and Toronto for most of our study period.

3 Data

Our main analyses uses data on bureaucrat appointments to explain Korean exports. This is complemented with additional data regarding the three main functions of KOTRA's overseas activities.

3.1 Bureaucrat Appointments

The most relevant source regarding bureaucrat appointments comes from contemporaneous reports in major Korean newspaper on appointments to KOTRA's overseas offices. These have the advantage of denoting the precise date of the announcement. In most years, there were two main dates at which appointments were announced, usually in January and July. The actual start most frequently occurred in April and October. Further, this information is usually reported in three major newspapers (*Dong Ah Ilbo*, *Choson Ilbo*, and *Kyonghyang Sinmun*), implying that there are almost no rounds of announcements that we fully missed. For almost all rounds of announcements we were able to corroborate the information using at least two of these sources.

The data on appointments is further complemented and corroborated using a variety of KOTRA publications on the director in charge of an office at a given point in time. We obtained and digitized the names of bureaucrats in (i) monthly publications aimed at non-Korean importers (1966-1971), (ii) a directory of KOTRA's network including all of its overseas bureaucrats (1977, 1991-1994, 1998-2000), (iii) KOTRA's reports on trade fairs (1969, 1971-1997), and (iv) a full directory of all overseas office directors using the Korean Business Directory, published by the Korean Chamber of Commerce and Industry.

Overall, we are able to identify 138 offices that existed between 1962 and 2001, located in 87 distinct countries. We identify 475 unique directors and 974 unique appointments of directors to offices. Table 1 provides further descriptive statistics on directors and appointments.

Directors are identified using their names, which requires us to avoid two types of errors. First, we may code two bureaucrats as the same one, e.g, it may be that bureaucrats share names. A priori, this could have been a problem as 45% of bureaucrats in our sample share

the last names *Kim*, *Lee*, and *Park*.[9] However, this is remedied by a great diversity in first names.[10] After a plethora of checks, it appears very unlikely that any bureaucrats in our data share the exact same full name. More challenging in practice, we had to determine whether slightly different names truly corresponded to distinct bureaucrats. This task was complicated as over time our sources move from Chinese to Korean characters to render the bureaucrats' names. In addition, in the few cases where names are given using romanizations, inconsistent romanization is used, e.g. *yul* and *ryul*. We resolved this task in four steps: Identify wrongly spelled or digitized names by (1) matching very unusual names to more common ones, (2) harmonizing the rendering of certain syllables, e.g. *yul* and *ryul*, (3) identifying offices with likely mistakes, e.g. the director's name flips back and forth. (4) Re-creating the career of each bureaucrat and assessing patterns of overlap or missing years. Following these steps meticulously allowed us to create a consistent panel of unique bureaucrats covering all offices and all years.

3.2 Exports

Our main measure of exports comes from Feenstra and Romalis (2014) who create consistent measures of bilateral trade flows, based on UN Comtrade data, at the year and 4-digit product level starting in 1962 and covering the entire period, up to 2001. Examples of these 4-digit products are given by "Rails of iron or steel", "Aircraft, heavier than air", and "Fur clothing".

In addition to these country×product×year export data, we obtained and digitized firm-level export data for the years 1968 to 1977 from KOTRA's archival publications. These data contain observations at the firm-country-product-year level.

3.3 Bureaucrat Output

We complement the data on exports with measures of concrete bureaucrat activity digitized from KOTRA documents.

[9] Moreover, the top 15 last names account for 76% of bureaucrats.

[10] Only twenty first names occur more than once. Only two first names occur three times in our data (*Dae-gyun* and *Won-kyung*).

First, we extract data on KOTRA's activity as a provider of "information service" such as market reports and transmission of importer requests to potential importers. We extract the market reports and importer requests from around 7,936 daily publications covering almost every weekday from 1965 to 2001. Of the 80,000 market reports, we are able to link 45,000 to both a 2-digit product and a country. The remaining reports are either not product-specific or do not discuss specific countries. Of the 200,000 inquiries, we are able to link 170,000 to both a 4-digit product, a country, and a specific office.

Second, we observe attendance and sales during trade fairs where a Korean representation was organized by KOTRA. This data covers 893 trade fairs attended by KOTRA between 1969 and 1997, including 192 events where KOTRA's responsible for a fair changes from one year to the next. On average, the Korean representation was composed of 2-3 KOTRA bureaucrats, usually headed by the local office director, and 15 Korean exporting firms. Overall, the data contains 34,000 encounters between a KOTRA bureaucrat and a Korean firm, i.e., bureaucrat and firm attend the same trade fair. Our data hence allows us to observe firms' fair attendance often including their sales deals at the fair, as well as certain firm characteristics, at least the firm's history in attending other KOTRA facilitated fairs and the bureaucrats the firm encountered at those fairs.

4 The Effect of Office Opening on Exports

This section of the paper uses the staggered roll-out of each country's first office to identify the causal effect of opening an export promotion office on Korean exports to this country. This allows us to discuss the average effect of export promotion offices, a policy-relevant variable. More importantly for this paper's main question, the effect of an office provides a natural benchmark against which to compare the variation in exports due to individual bureaucrats. Later sections find that the difference in fixed effects between a bureaucrat at the 20th percentile and the median is of a similar magnitude as the effect of opening an office. This suggests an office with a bureaucrat at the 20th percentile has no effect on exports. Our setting is exceptional in enabling this natural comparison between a policy's average

effect and the variation introduced due to implementation by individual bureaucrats[11]. This is because (1) We observe a sufficient number of office openings and (2) Korean exports to a country constitute a well-defined variable even in absence of an export promotion office.

Figure 2 displays the staggered roll-out of offices: There were seven countries with offices in 1965. This increased to 67 by 1980. Figure 3 indicates the countries which had one or multiple offices in 1981, at the end of the main roll-out. Using this roll-out, we estimate a 38% increase in exports 9-11 years after the first office opening. Assuming an elasticity of trade to distance of -1 (Anderson, 2011; Head and Mayer, 2014), this is equivalent to reducing the distance between London and Seoul to the distance between Mumbai and Seoul.

4.1 Identification: Effect of Office Opening on Exports

To estimate the effect of an EP office, the ideal experiment would randomly allocate a fully-developed office to some countries and not to others. As this is not feasible, the analysis here will use the staggered roll-out of offices to countries.

$$y_{cpt} = \lambda_{pt} + \gamma_{cp} + X^T_{cpt} + \sum_{k \neq -1} \theta_k D^k_{ct} + \epsilon_{cpt} \tag{1}$$

As a first step, we estimate the specification given by equation (1). λ_{pt} indicates product-year fixed effects, γ_{cp} indicates country fixed effects that may differ at the product-level. D^k_{ct} are dummies equal to 1 if year t is k years after the first office opened in country c. θ_k corresponds to the effect of the office after k periods. We compare countries with an opening event in the years 1964 to 1981 (ever-treated) to the never-treated: countries which never receive an office, i.e. no office until 1992.[12] Equation (1) corresponds to a Diff-in-Diff specification because we subtract out the difference that exists between a treated country and all never-treated countries in period -1, the year before the office opens. For countries that do receive an office between 1964 and 1981, we include two years prior to the office opening and eleven years after the office opening. The earliest start year for a treated

[11] Settings from related papers do not lend themselves to obtaining such a benchmark (Fenizia, 2022; Best, Hjort, and Szakonyi, 2023; Otero and Muñoz, 2022; Metcalfe, Sollaci, and Syverson, 2023).

[12] The original, very fast-moving, rollout of offices to countries stopped in 1981. This is why we focus on events until this date.

country's event horizon is 1962. The latest end year is 1992. Hence, for countries that do not receive an office, we include all years between 1962 and 1992. In the appendix, we report estimates when omitting events in 1964 and 1965 or 1964 to 1966, allowing for four and five pre-periods, respectively.

We rely on two main assumptions to interpret $\hat{\theta}_k$ as estimating the causal effect of the office opening after k periods. (1) Parallel trends: We assume that counterfactual trends - in absence of an office opening - do not differ in periods $g + k$ with $k > 0$ between those treated in year g and the never-treated. Persistent level differences between the treatment and control group do not constitute a violation of this assumption. Alleviating concerns about parallel trends violations, we find little indication of differential pre-trends and no "effect" on non-Korean exports to a destination. We further show that the offices' rollout largely followed pre-determined variables related to a gravity equation, making it less likely that KOTRA timed office openings with counterfactual trends in exports to a country.

(2) No spillovers, i.e. an office affects exports only to the country in which it is located. More technically this is the stable unit treatment value assumption (SUTVA), as one unit's treatment value may not depend on other units' treatment. This would be violated if firms simply redirect exports from one country to another that gains a KOTRA office. While still constituting an effect due to the office opening, this would lead to an overestimate of the causal effect of treatment on Korean exports to the respective country as the untreated suffer from a negative spillover effect. Such a violation appears less plausible for the main specification using a never-treated control group. As we find similar estimates using a not-yet-treated control group (Callaway and Sant'Anna, 2021), SUTVA violations of this type may not be a first-order concern. Alfaro-Ureña, Castro-Vincenzi, Fanelli, and Morales (2023) – ACFM – provide further reason to believe SUTVA violations do not upwardly bias our estimates . ACFM assume that exporting to one country never decreases a firm's exports to another country. Under this assumption, they find that exports to different countries are complements. Hence, violations of SUTVA could lead to underestimates.

A third assumption – no anticipation – is required for identification. This would be violated if office openings have a causal effect, at $k < 0$.

Our main specification uses the inverse hyperbolic sine of Korean exports as the outcome

variable and does not include a control variable (X_{cpt}). Later specifications will control for non-Korean exports, use the inverse hyperbolic sine of non-Korean exports as a placebo outcome, and focus on the extensive margin of exporting - i.e. the number of products exported to a country.

4.2 Results: Effect of Office Opening on Exports

Figure 4 reports the estimated effects of the first overseas EP office in a destination country around the year of the office opening. We cannot reject the hypothesis that there are parallel pre-trends. The coefficient in period -2 is economically small and not statistically distinct from 0. This assuages concerns that the parallel trends assumption is violated. Figure 4 shows that the opening of an export promotion office is associated with an increase in Korean exports to that destination. While the estimates in nearly all post periods allow us to reject the null-hypothesis of no effect, the point estimates themselves are somewhat imprecisely estimated. The estimates increase over time, suggesting that the entire effect of an office opening only materialize over time. The point estimates flatten off five years after the office opening. The average point estimate in years 9-11 is 0.321, suggesting exports are 38%[13] higher relative to the control group.

Assuming an elasticity of -1 of trade to distance, an office opening has an effect similar to reducing the distance between London and Seoul to the existing distance between Mumbai and Seoul. Alternatively, an office opening makes a country with a fixed effect at the 25th percentile (Ecuador) as attractive as a country at the 50th percentile (Greece).[14] At the same time a country at the 50th percentile (Greece) becomes as attractive as a country at the 75th percentile (Spain) due to an office opening.

4.3 Robustness, no increase in import demand, roll-out pre-determined

First, we consider that export promotion offices may be opened strategically in years when a destination country experiences increases in import demand. We address this concern in two ways. First, we re-estimate equation (1) while controlling for non-Korean exports

[13] $lim_{x\to\infty} sinh(x + 0.321)/sinh(x) = 1.379$

[14] Percentiles are calculated for those countries that ever have an office.

to a country (also transformed as the inverse hyperbolic sine). Appendix figure A.4 shows that the estimates from this specification are largely unchanged compared to the baseline. Second, instead of Korean exports we use non-Korean exports as the dependent variable. The coefficients from this regression are reported in appendix figure A.4. It shows that opening export promotion offices does not coincide with statistically significant effects regarding this placebo outcome. Appendix figure A.4 also reports estimates when restricting attention to events that happen in 1966 (1967) or later. We do so to allow for the estimation of more pre-treatment coefficients. Panel (c) suggests quite parallel trends between periods -4 and -2 with a moderate uptick in period -1. Panel (d) more convincingly finds a parallel pre-trend. The uptick in period -1 is discussed further below when we check the sensitivity of our estimates to parallel trend violations following Rambachan and Roth (2023).

Second, appendix figure A.5 uses a "not-yet-treated" control group instead of the "never treated" used by our main estimation strategy. These figures report estimates following the estimator proposed by Callaway and Sant'Anna (2021), which allows for consistent estimates in cases where our main TWFE approach fails. We obtain estimates of very similar magnitude and precision to our main estimation strategy. However panel (a.i) finds coefficients below 0 that are statistically significant, albeit small, in periods -4 to -2. This leads us to investigate the sensitivity to parallel trends violations. We do so for our main estimate using the "not-yet-treated" control, reported in panel (a.i), and an estimate that treats period -1 as the first treated period. This is reported in panel (b.i). The latter would suggest that KOTRA has an effect on exports in the year before opening an office. This appears plausible as setting up an office already requires resources dedicated to a country that may have a direct or indirect effect on exports. Panels (a.ii) and (b.ii) show that our estimates of the effect on exports ten years after an office opening remain statistically significant when allowing for parallel trends violations up to 1 times (1.5 times with one period of anticipation) the largest pre-treatment violation of parallel trends. Panels (a.iii) and (b.iii) show that our estimates remain significant when allowing for slope changes of 0.15% ($\approx 0.4\%$) across consecutive periods (Rambachan and Roth, 2023).

We also find sizable effects on the extensive margin. Appendix figure A.6 finds a 5% increase in the likelihood of a product being exported to a destination country point 10

years after an office opening. These estimates are very similar using our main estimation strategy - whether or not restricting attention to openings from 1967 - as well as the "not-yet-treated" control group following Callaway and Sant'Anna (2021). While the former approach suggests there are no differential pre-trends, the latter approach again suggests there may be an effect in the year before an office opens up. Hence, panel (d) report estimates allowing for one period of anticipation.

Third, we show that the year in which a country's first office opened was largely pre-determined by time-invariant factors. As long as the effect of these factors is also time-invariant, they are absorbed in γ_c. Even if the effect of these time-invariant variables is not stable over time, the pre-determined order of the roll-out makes it unlikely that office openings are timed to coincide with counterfactual increases of exports, whether strategically or coincidentally, rendering violations of the parallel trends assumption less plausible. To predict office openings, we use insights from a gravity equation. Apart from the U.S., the first office openings took place in Taiwan, Thailand, Japan, Singapore, Indonesia, and South Vietnam – among the geographically closest non-communist countries. Within Europe, distance from Korea is relatively stable, so the main predictor for office openings from a gravity equation would be the size of each destination's market. To rule out that openings were timed based on counterfactual increases in Korean exports, we use 1962 non-Korean exports to a country to predict the year when a country would get its first office. We do so for all European countries where an initial office opened during the main roll-out of overseas offices, between 1962 and 1981. Figure 5 shows that a country's pre-determined market size predicts most of the variation in office openings. Appendix table B.1 shows that true and predicted opening years often coincide exactly.

4.4 Extension: Office Openings Immediately Increase Activity

We analyze how country-specific activities change around the opening of an office by re-estimating equation 1. Instead of exports, we aim to explain three measures of KOTRA activity, each transformed using the inverse hyperbolic sine. (1) The number of reports about a country, (2) the number of product-specific reports – which may be more specific or informative, (3) the number of inquiries for trade related to the country. Figure A.7 reports

results for these three outcomes. For each outcome, the coefficients stabilize after a couple of years at around 1. This translates into roughly multiplying by 2.7[15] the annual number of reports and inquiries – from 8 to 21 and from 26 to 70.

5 Bureaucrats as Drivers of Korean Exports

This section shows that the effect of an office differs substantially depending on the assigned bureaucrat. It finds the following two results: (1) Moving from the 20th percentile bureaucrat to the median increases exports by 39%, after applying a standard shrinkage procedure to the estimated bureaucrat effects. This effect is roughly the same as opening an office for the first time, i.e. as moving London as close to Seoul as Mumbai is. (2) This implies that the industrial policy under study is ineffective under a 20th percentile bureaucrat. We conduct a number of diagnostic checks that assuage concerns that our estimates of bureaucrat ability are not causal.

5.1 Identification: Estimating Bureaucrat Effects

We adapt the AKM framework to study how much bureaucrats matter in explaining Korean exports (Abowd, Kramarz, and Margolis, 1999; Abowd, Creecy, and Kramarz, 2002; Bertrand and Schoar, 2003; Fenizia, 2022; Best, Hjort, and Szakonyi, 2023). We exploit the rotation of office directors across countries to estimate bureaucrat and country (office) fixed effects. We follow two main approaches to correct for the fact that raw fixed effects contain measurement error. First, we follow a shrinkage approach to correct for the sampling error (Best, Hjort, and Szakonyi, 2023; Chetty, Friedman, and Rockoff, 2014a). Second, we explicitly correct for the limited mobility bias (Kline, Saggio, and Sølvsten, 2020).

Table 1 describes the structure of our sample. Column 1 reports the statistics for the full sample of offices. Column 2 restricts attention to the main country offices which will be used to explain country-level outcomes, mainly exports to a destination country. The sample for the country-level analysis contains 397 directors and 86 countries. 194 directors move across

[15] $lim_{x\to\infty} sinh(x+1)/sinh(x) = 2.718$

countries, and 82 countries experience multiple directors[16] The remaining 4 countries do not contribute to the estimation of the bureaucrat effects. All 86 countries and 397 bureaucrats in sample (2) are part of the same connected set. This is because our data covers four decades with the average office experiencing nine different office directors.

We model the inverse hyperbolic sine of exports, henceforth "exports", associated with country c, product p, year t, and the bureaucrat assigned to that country-year - $b(c,t)$. Exports are explained by the sum of a product-year component (λ_{pt}), a bureaucrat component ($\theta_{b(c,t)}$), a country component (γ_c), and an error term (ϵ_{cpt}).[17] As in other parts of the paper, we aim to explain exports at the product-level. This avoids that results for a country-year are driven by a couple of dominant export products, and increases statistical power.

$$\text{exports}_{cpt} = \lambda_{pt} + \gamma_c + \theta_{b(c,t)} + \epsilon_{cpt} \tag{2}$$

Our preferred estimation of the leave-one-out connected set to obtain the raw fixed effects (Kline et al., 2020). This is the sample of countries and bureaucrats that remains connected even when one bureaucrat-country pair. The leave-one-out connected set contains 75 countries and 93% of appointments is because most offices in our data experience a large number of appointments.

Identification requires that director mobility is as-good-as-random, conditional on product-year and country fixed effects. In other words, bureaucrat assignments need to be uncorrelated with underlying trends in exports. On the other hand, this orthogonality condition allows for director assignment to offices on the basis of the permanent component of country exports α_c or the permanent component of director ability $\theta_{b(c,t)}$. That is, sorting of better bureaucrats to destinations with greater time-invariant Korean exports, e.g. larger countries, would not violate the identifying assumptions.

To alleviate concerns that the identifying assumption is violated, we conduct a number of checks in section 5.3. Here we discuss factors influencing bureaucrat appointments. These al-

[16]This is slightly larger than the 184 movers in the balanced analysis sample of Fenizia (2022). Compared to Fenizia (2022), our power is enhanced because most of our countries and bureaucrats are part of the same connected set, even a leave-one-out connected set.

[17]To account for the fact that it takes time for a new director to influence exports, we code each country-year as being headed by the bureaucrat in office until March that year. This means, we attribute effects to a bureaucrat for up to nine months after their successor has been appointed.

leviate concerns that appointments are strategically timed to coincide with increasing export trends.

In qualitative interviews, KOTRA employees mentione two factors determining appointments. (1) Bureaucrats are more likely to be appointed to a country when they speak the local language. (2) Bureaucrats prefer being appointed to high-income, English-speaking countries. Because these preferences are largely homogeneous between bureaucrats, KOTRA's HR manages discontent by rotating bureaucrats between low- and high-desirability appointments. In most cases both a country's language and its income relative to other countries do not change much over time. So the above-mentioned factors suggest appointments may be correlated with country fixed effects. They also provide a reason why export trends are less likely to predict bureaucrat appointments.

A further constraint on bureaucrat appointments is due to their three-yearly appointment schedule. As highlighted in appendix figure A.1, whether there is a new appointment to country c in year t is largely a function of the timing of the previous bureaucrat's appointment to c. Moreover, whether bureaucrat b is available to be appointed to country c at this time is largely a function of the timing of bureaucrat b's previous appointment.

Lastly, one may be concerned that despite the above-mentioned constraints, better bureaucrats are strategically appointed to country-years with high import demand. It seems plausible that time-invariant country characteristics are more important than trends in import demands: The time-varying demand-shocks that make Portugal a more important export destination than the UK would have to be very large.

Overall, these factors suggests that it is somewhat implausible that KOTRA is strategically timing appointments to coincide with increasing export trends.

Note that equation (2) corresponds to an inverse hyperbolic sine-linear regression specification. This approximates a log-linear regression equation, but allows for the data containing zeroes.[18] Equation (2) implies the assumtion that the inverse hyperbolic sine of Korean exports is linear in bureaucrat and country effects. Section 5.3 presents results in support of this ihs-linear specification. Section 5.6 shows how the fixed effects translate into extensive

[18] We explore robustness to the inverse hyperbolic sine transformation in section 5.6. We find that bureaucrat fixed effects are predicted of changes in both the extensive and intensive margin.

and intensive margin changes to exports.

$$(\text{exports}|pt)_{cpt} = \text{exports}_{cpt} - \hat{\lambda}_{pt} = \theta_{b(c,t)} + \gamma_c + \epsilon_{cpt} \tag{3}$$

For the two bias corrections, we simplify equation (2) by removing the effect of product-year dummies from the value of exports to obtain $(\text{exports}|pt)_{cpt}$ according to equation (3), where $\hat{\lambda}_{pt}$ is estimated from equation (2). This is necessary for explicitly correcting the limited mobility bias (Kline et al., 2020) and we chose to do likewise for the shrinkage correction.[19] $\hat{\lambda}_{pt}$ is likely important.[20] However, these fixed effects are common to all countries in the same year and hence unrelated to individual offices or bureaucrats.

Our preferred approach shrinks the raw fixed effects under an approach that bootstraps the estimation of equation (3) to distinguish the true, signal variance in bureaucrats' and the variance of their sampling error (Chetty, Friedman, and Rockoff, 2014a; Best, Hjort, and Szakonyi, 2023). This has the advantage of yielding shrunk fixed effects for each bureaucrat, hence allowing us to compare different parts of the distribution, e.g. the 20th and 50th percentile. To obtain the bootstrapped samples our preferred approach draws appointments from the set of all appointments. [21]

$$\text{Var}[(\text{exports}|pt)_{cpt}] = \text{Var}(\theta_{b(c,t)}) + \text{Var}(\gamma_c) + 2\text{Cov}(\theta_{b(c,t)}, \gamma_c) + \text{Var}(\epsilon_{cpt}) \tag{4}$$

As a second approach, we follow Kline, Saggio, and Sølvsten (2020) to obtain a variance decomposition that directly corrects for the limited mobility bias that arises in two-way fixed-effects specification when the switches across groups occur too infrequently. As variation in residualized exports within spells is uninformative in the estimation of the bureaucrat or country fixed effects, we take the spell-level averages of the residualized exports as the total

[19] This follows Chetty, Friedman, and Rockoff (2014b) who explain that to remove the effect of pt without biasing the bureaucrat effects θ and country effects γ, $\hat{\lambda}_{pt}$ needs to be estimated using only within-bureaucrat and within-country variation.

[20] $\hat{\lambda}_{pt}$ captures macroeconomic shocks, but also long-run changes in Korea's industrial structure. E.g., $\hat{\lambda}_{cars,1965}$ is very small compared to $\hat{\lambda}_{cars,1995}$. Table B.2 highlights the importance of these factors as year-product fixed effects explain 35.5% of the variation in exports.

[21] Alternative approaches yield similar or less conservative shrinkage factors. These include (ii) drawing countries from the set of all countries, (iii) drawing years from the set of all years, (iv) drawing country-year-product observations from the set of all country-year-products observations.

variation.[22,23] We use the computational algorithm of Bonhomme, Holzheu, Lamadon, Manresa, Mogstad, and Setzler (2023) for implementation. Although unreported, the Andrews, Gill, Schank, and Upward (2008) correction method that assumes homoskedasticity delivers quantitatively very similar results.[24] We report the variance decomposition according to equation (4).

5.2 Main Results: Bureaucrats Are Crucial to Policy Success

Figure 6 reports the cumulative distribution function of bureaucrat fixed effects obtained from estimating equation (3). The difference in fixed effects between the 20th percentile bureaucrat shrunk by the ratio of signal to sample variation implies a difference in exports of 39%.[25] As this difference is similar to the effect of an office opening, an office causes an increase in exports only to the extent that the bureaucrat in charge is better than the 20th percentile. Moving from the median bureaucrat to the 90th percentile has a similar effect.

Next, we report the results from the variance decomposition of equation (4). Columns (1)-(2) of Table 2 report our preferred variance-decomposition results. The variation explained by bureaucrat fixed effects implies a standard deviation of the bureaucrat fixed effects of 0.31, implying a difference in exports of 37%.

Hence, increasing bureaucrat ability by one standard deviation amounts to roughly the effect of moving London as close to Seoul as Mumbai actually is - similar to opening an office or moving from the 20th percentile to the median or from the median to the 90th percentile.

Column (1)-(2) also highlight that bureaucrats are about 1/7 as important as countries. The negative correlation between bureaucrat and country fixed effects suggests that better

[22] In fact, the two-way fixed-effects estimation is performed on the data that is already collapsed at the spell level. The bureaucrat and country fixed effects estimated on this collapsed data are perfectly correlated to those that are estimated on the uncollapsed, raw data.

[23] The variance of the raw (i.e. country×product×year-level) residualized exports is also reported in Table 2 for reference.

[24] While the Kline, Saggio, and Sølvsten (2020) correction method can only be performed on the leave-one-out connected set which covers 75 countries and 380 bureaucrats, the Andrews, Gill, Schank, and Upward (2008) correction method can also be performed on the largest connected set covering 86 countries and 397 bureaucrats. The Andrews, Gill, Schank, and Upward (2008) correction method delivers extremely similar results for either measure of connectedness.

[25] We bootstrap to obtain the sample variance in each bureaucrat fixed effect. Our preferred approach bootstraps over appointments, as our data can be thought of as a random sample of all feasible combinations of bureaucrat-country matches.

bureaucrats work in smaller countries. Overall, after taking out time-trends, bureaucrat and country fixed effects explain 88% of the spell-level variation in exports.

Moreover, we perform a "placebo check" on the validity of the variance decomposition exercise when bureaucrat fixed effects should *not* have any explanatory power. Columns (5)-(6) show the results when bureaucrats are randomly shuffled to countries while preserving the number of different appointments for each bureaucrat. Both the variation in bureaucrat fixed effects, as well as the covariance between bureaucrat and country fixed effects, go to zero, as they should.

To allay concerns that the fixed effects of single-appointment bureaucrats may suffer from aggravated overfitting[26] and therefore magnify the variation in bureaucrat fixed effects, we also report in columns (3)-(4) the variance decomposition results excluding them. The share of total variation in residualized exports explained by bureaucrats does drop to around 8%. It should also be noted that the Bonhomme, Holzheu, Lamadon, Manresa, Mogstad, and Setzler (2023) algorithm is designed to handle an abundance of individuals with one spell only in the sample. On the other hand, section 5.5 suggests that the lowest ability bureaucrats are endogenously not re-appointed. Hence, the somewhat smaller variation in ability among re-appointed bureaucrats is an interesting results. The fact that the variation in bureaucrat fixed effects is not any larger when including the single-appointment bureaucrats (columns (5)-(6)) than when excluding them (columns (7)-(8)) in the randomly shuffled data also supports the reliability of the preferred decomposition results of columns (1)-(2) that includes single-appointment bureaucrats.

5.3 Diagnostics: No evidence for pre-trends or misspecification. Bureaucrat effects are predictive out of sample.

To check the validity of the fixed effects obtained from estimating equation (4), we next obtain event study estimates which show that switches in bureaucrat effectiveness correspond to sharp changes in exports without corresponding differential pre-trends. Hence, the bureaucrat fixed effects are likely to correspond to each bureaucrat's causal effect on

[26]For a single-appointment bureaucrat, their fixed effect value equals the residualized export value to the country they were appointed to during their appointment, less the country fixed effect value of that country.

exports. Given the lack of pre-trends and the finding of sharp jumps around new appointments, it seems implausible that the fixed effects are driven by spurious correlations between bureaucrat appointments and time trends.

$$y_{ept} = \eta_{ep} + \lambda_{pt} + \sum_{k \neq -2} \Big(\alpha_k + \boldsymbol{\beta_k} \ \mathbf{1}\{\Delta\hat{\theta}_e \text{ in top tercile}\} + \boldsymbol{\delta_k} \ \mathbf{1}\{\Delta\hat{\theta}_e \text{ in middle tercile}\} \Big) \mathbf{1}\{t = T + k\} + \epsilon_{ept} \tag{5}$$

First, we statistically test the whether the change in bureaucrat fixed effects is predicted by differential pre-trends prior to their appointment. To that end, we estimate equation (5), which explains exports as a time-varying function of the change in fixed effect for this office due to the switch from the old to the new bureaucrat. Following the literature (Fenizia, 2022; Otero and Muñoz, 2022), we divide the events into terciles depending on the change in bureaucrat fixed effects due to them. β_k (δ_k) are the effect in event time k of a change in the top (middle) tercile relative to one in the bottom tercile. e indicates the event. It is uniquely defined by the country, c, and the year of the event, T, defined as the first full year that the new bureaucrat is appointed to country c. Equation (5) obtains the event-study estimates while controlling for trends using product-year fixed effects (λ_{pt}) and for pre-event levels of exports using event-product fixed effects (η_{ep}). In obtaining the event study estimates, we normalize by the last full year in which the *old* bureaucrat was in charge: $T - 2$.

Figure 7 reports the effect of a switch in bureaucrats depending on the change in bureaucrat fixed effects it implies. It shows that top (middle) tercile transitions are not predicted by differential pre-trends compared to a bottom tercile transition. They do, however, imply a jump in exports by 30% (11%) upon the appointment of the new bureaucrat.

$$y_{ept} = \eta_{ep} + \lambda_{pt} + \sum_{k \neq -2} \left(\alpha_k + \boldsymbol{\beta_k} \hat{\theta}_e^{\text{new}} + \boldsymbol{\delta_k} \hat{\theta}_e^{\text{old}} \right) \mathbf{1}\{t = T + k\} + \epsilon_{ept} \tag{6}$$

Next, we test whether this jump in exports is driven by the fixed effect of the newly appointed bureaucrat or by the bureaucrat that moves away. To that end, we estimate equation (6), which explains exports as a time-varying function of the fixed effects of the new bureaucrat ($\hat{\theta}_e^{\text{new}} = \hat{\theta}_{b(c,T)}$) and the old bureaucrat ($\hat{\theta}_e^{\text{old}} = \hat{\theta}_{b(c,T-1)}$). Other than distinguishing

between $\hat{\theta}_e^{\text{new}}$ and $\hat{\theta}_e^{\text{old}}$, this specification follows equation (5).

Figure 8 plots the event-study estimates (β_k and δ_k) obtained from equation (6). It shows that exports change sharply in the direction of the ability of the incoming bureaucrat and symmetrically against the direction of the outgoing bureaucrat's ability. Pre-trends are not statistically distinct from 0 and economically very small.

The fact that *losing* a bureaucrat has a symmetric effect to *gaining* a bureaucrat further alleviates concerns about endogenous movement of bureaucrats for two reasons: (1) The three-yearly rotation provides a strong rationale that *losing* a bureaucrat is exogenous (to bureaucrat ability and underlying export trends). So it is reassuring that the fixed effects strongly predict the drop in exports due to losing a bureaucrat. (2) Equation (3) implies the assumption that the effect of bureaucrats on exports is static – it is fully materialized every year of the bureaucrat's appointment and absent in every other year. The symmetry reported in figure 8 allays concerns that this is unrealistic. It may be surprising that there is a strong drop in exports upon the appointment of an ineffective bureaucrat. However, this is only relative to Korean exports to other countries. Given that Korean exports were growing at 20% annually, it appears plausible that losing a good bureaucrat means exports drop *relative to this trend*.

Next, appendix figure A.8 shows time trends in de-trended exports around the years when an office experiences a change in the director, closely following Card, Heining, and Kline (2013), Card, Cardoso, and Kline (2016) and Best, Hjort, and Szakonyi (2023). To further explore the dynamics of exports around switches in bureaucrats, it classifies switches into terciles of effectiveness of the new and old bureaucrat. These are obtained from average de-trended exports of a product during a bureaucrat's appointments, i.e. bureaucrat fixed effects after residualizing exports by product-country and product-year fixed effects.

Figure A.8 corroborates the main takeaways from figure 8. First, in the pre-periods, exports are highest when the outgoing bureaucrat is in the top tercile and lowest when the old bureaucrat is in the bottom tercile. Second, in the post-period, the effect of the outgoing bureaucrat's tercile becomes less important, the effect of the new bureaucrat's tercile becomes dominant. In year 1, exports are lowest when the new bureaucrat is in the bottom tercile. They are highest when the new bureaucrat is in the top tercile. Third,

exports change sharply, and in the expected direction, precisely when a destination switches to a less or more effective bureaucrat. Exports increase the most upon a switch to the highest tercile and (relatively) decrease the most upon a switch to the lowest tercile. The figure shows little sign that exports are increasing in organizations that subsequently switch to a better bureaucrat, and vice versa. This suggests that drift in effectiveness and switches are uncorrelated.

Next, we explore the assumption of additive separability between bureaucrat and country effects. A violation of additive separability would result in residuals with high absolute values for certain kinds of bureaucrat-country pairs. Following the literature (Fenizia, 2022; Otero and Muñoz, 2022), we divide our observations based on the quartile of the estimated manager fixed effect and the quartile of the estimated country fixed effect. Figure 9 shows that mean residuals are small for each combination of bureaucrat and country quartiles – between -0.05 and 0.05 in absolute value – with no clear pattern, allaying concerns about the assumption of additive separability.

We further corroborate the additive separability assumption. Table B.2 reports the variance explained by the different levels of fixed-effects when estimating equation (2), which identifies the causal effect of the two sets of fixed effects under the assumptions discussed above. While informative, these are subject to some of the criticisms addressed by the Kline, Saggio, and Sølvsten (2020) and Andrews, Gill, Schank, and Upward (2008) bias correction reported in table 2.

First, we discuss the explanatory power of bureaucrat fixed effects over and above country-product fixed effects. Adding bureaucrat fixed effects increases R^2 by 0.018, about 18.6% of the increase in explanatory power from adding country fixed effects. The absolute increase in R^2 is smaller than other recent papers, studying managers of organizations that process insurance claims (increase in R^2 of 0.11, Fenizia), or hospital CEOs (0.09, Otero and Muñoz). Relative to the explanatory power increase from adding country or organization fixed effects, the increase in R^2 due to bureaucrats is slightly smaller than other recent papers studying bureaucrats who run organizations that process insurance claims (23.4%, Fenizia) and public hospitals (28.0%, Otero and Muñoz).

Next, we assess how much the effect of a bureaucrat differs across their appointments.

For this we compare the explanatory power when including appointment fixed effects (column 4) compared to column (3) which assumes that bureaucrat and country effects are time-invariant. The increase in explanatory power from this is negligible, suggesting that bureaucrats' effects are relatively stable across appointments, which provides some support that the productivity of a bureaucrat-country match is well approximated by the linear combination of the bureaucrat fixed effect and the country fixed effect.

A further set of checks assesses whether the fixed effects we estimate are also predictive out of sample. We find that this is the case, further allaying concerns about overfitting.

The most natural and conservative way to do this, in our setting, is to only use *other countries* to estimate the fixed effects. This means to estimate the fixed effects of bureaucrats appointed to the UK, we obtain their fixed effects on a data set using all country-years, except the UK. This comes at a cost. For a bureaucrat with n appointments, the out-of-sample FE are estimated on $n-1$ appointments. This means, only for about half our sample are out-of-sample FE defined. Half of the remaining bureaucrats has a total of two appointments, meaning their out-of-sample FE are estimated from only one appointment. Nevertheless, the interconnectedness of our data makes it possible to estimate such out-of-sample fixed effects. When estimating FE while leaving out one country, we always retain one very large connected set, as 75 countries in our data are part of the same leave-one-out connected set.

Figure 10 displays a binned scatterplot of residual exports and in-sample as well as out-of-sample fixed effects. By construction, the slope for the in-sample fixed effects equals 1. More interestingly, out-of-sample we get a coefficient of 0.52. This is very close to the coefficient found by Metcalfe, Sollaci, and Syverson (2023), who study managers of retail stores. However, their approach splits the sample into two periods pre-Covid and during Covid. That means they may still estimate a bureaucrat's fixed effect from the same appointment in which they then try to predict performance which is never the case for us.

Further, figure A.9 replicates figure 8 using out-of-sample, i.e. *other country*, fixed effects. Around a switch in bureaucrats it shows that new and old ability still statistically significantly predict exports in the expected way even when ability is estimated only using other countries.

Overall, this section provides support to the interpretation that bureaucrat fixed effects

identify the causal impact of an individual bureaucrat on exports. Given the lack of pre-trends and the finding of sharp jumps around gaining *and losing* bureaucrats, it seems implausible that the fixed effects are driven by correlations between bureaucrat appointments and underlying export trends. Hence, this section provides evidence in support of the findings from the variance decomposition.

5.4 Mechanism: Good Bureaucrats Increase Responsiveness of Exports to Market Conditions

We next investigate whether part of the increase in exports upon the appointment of a high ability bureaucrat is due to an increased elasticity to market conditions. We show that upon the switch to a more effective bureaucrat, Korean exports' elasticity to market conditions increases sharply. Our findings suggest that most of the effect of high ability bureaucrats comes from more effectively exploiting market conditions, e.g., by relaying information about destination market demand.

$$
\begin{aligned}
y_{ecpt} = \eta_{ep} + \lambda_{pt} + \psi_d^0 \text{demand}_{cpt} + \psi_s^0 \text{supply}_{cpt} + \psi_{d,new}^0 \text{demand}_{cpt} \times \hat{\theta}_e^{\text{new}} + \\
\psi_{s,new}^0 \text{supply}_{cpt} \times \hat{\theta}_e^{\text{new}} + \psi_{d,old}^0 \text{demand}_{cpt} \times \hat{\theta}_e^{\text{old}} + \psi_{s,old}^0 \text{supply}_{cpt} \times \hat{\theta}_e^{\text{old}} + \\
\sum_{k \neq -2} \Big[\alpha_k + \psi_{dk} \text{demand}_{cpt} + \psi_{sk} \text{supply}_{cpt} + \boldsymbol{\beta_k}\ \hat{\theta}_e^{\text{new}} + \boldsymbol{\delta_k}\ \hat{\theta}_e^{\text{old}} + \\
\boldsymbol{\beta_k^{demand}}\ \text{demand}_{cpt} \times \hat{\theta}_e^{\text{new}} + \boldsymbol{\beta_k^{supply}}\ \text{supply}_{cpt} \times \hat{\theta}_e^{\text{new}} + \\
\boldsymbol{\delta_k^{demand}}\ \text{demand}_{cpt} \times \hat{\theta}_e^{\text{old}} + \boldsymbol{\delta_k^{supply}}\ \text{supply}_{cpt} \times \hat{\theta}_e^{\text{old}} \Big] \mathbf{1}\{t = T + k\} + \epsilon_{ecpt}
\end{aligned}
\tag{7}
$$

We estimate equation (7), which explains changes in exports around a new appointment. This estimating equation includes all the components from equation (6). In addition, it includes main effects and interactions of "demand" and "supply". "Demand" is the short-hand for other countries' exports of the same product to the same destination. "Supply" is the short-hand for Korean exports of the same product to other destinations. ψ_d^0 and ψ_s^0 estimate the elasticity to market conditions in the pre-period. $\psi_{d,new}^0$, $\psi_{s,new}^0$, $\psi_{d,old}^0$, $\psi_{s,old}^0$ allow for differences in the pre-period based on the ability of the new or old bureaucrat. The new parameters of interest are β_k^{demand} , β_k^{supply} , δ_k^{demand}, δ_k^{supply} which give the difference

in elasticity to market conditions relative to the last full year the old bureaucrat was in the country that is due to the estimated ability of the new or old bureaucrat.

Figure 11 plots the estimates of β_k, β_k^{demand}, β_k^{supply}, δ_k, δ_k^{demand}, and δ_k^{supply} for each event year. We find a sharp increase in the elasticity of Korean exports to market conditions. The elasticities to market conditions increase by around 5 pp when the bureaucrat ability increases by one standard deviation. This implies an increase in the reaction of Korean exports to market conditions by around 20% (from a base of around 25%).

As above, the reaction of Korean exports to market conditions decreases symmetrically in the old bureaucrats fixed effects. As the three-yearly rotation provides more reasons to believe losing a bureaucrat is exogenous, this further suggests these effects are not driven by strategic appointments.

The point estimates for the effect of incoming and outgoing ability due to the change in bureaucrat mostly remain statistically significant. They are, however, reduced to about 1/10 of their size in figure 8, suggesting that much of the effect of high ability bureaucrats is due to the increased elasticity of Korean exports to market conditions, e.g. by relaying information about local conditions (demand) and identifying opportunities based on market developments common to Korean exporters across destination markets (supply).

Figure 11 also is informative about pre-trends. The absolute values in the pre-period are never statistically significant at the five percent level and much smaller in absolute values than the estimates in the post-period.

Overall, this section provides additional support that more effective bureaucrat causally impact exports. It does so by highlighting a mechanism via which this takes place. Switching to a more effective bureaucrats causes a sharp increase in the elasticity of Korean exports to market conditions. Losing an effective bureaucrat causes a sharp decrease of similar magnitude.

5.5 Extension: Performance in 1st Office Predicts Career Length

This section finds that residualized exports during a bureaucrat's first appointment, part of their estimated fixed effects, are predictive of bureaucrats' careers. Figure 12 reports the probability density function of residualized exports, splitting the sample by the number

of appointments a bureaucrat has over their career. This distribution has a substantially fatter left tail for bureaucrats with only one career appointment. While far from causal, this suggests KOTRA may decide not to re-appoint bureaucrats with a very low ability.

We next regress a bureaucrat's number of appointments on residualized exports during their first appointment, part of a bureaucrat's fixed effect used in the preceding parts of section 5. By including fixed effects for the year of a bureaucrat's first appointment we rule out various omitted variables biases, most prominently: (1) The number of appointments depends highly on a bureaucrat's tenure at KOTRA. (2) Bureaucrats may differ systematically by their first year of appointment as an overseas office director. Within year of first appointment, we find a positive significant effect of residualized exports during a bureaucrat's first appointment on number of appointments of 0.240 (standard error: 0.112). This effect is robust to alternative specifications. We find a positive significant effect of 0.430 (standard error: 0.109) when regressing on a dummy that indicates residualized exports above the 25th percentile.

Overall, we find that residualized exports during a bureaucrat's first appointment are associated with a greater number of subsequent appointments as director of an overseas office. Allaying concerns that this may be due to differences in bureaucrat cohorts or bureaucrat tenure, this effect holds among bureaucrats whose first appointment began in the same year.

5.6 Extension: Bureaucrats affect extensive and intensive margin

This section unpacks the effect on exports into the extensive and intensive margin. We find that bureaucrat effects cause increases both along the extensive and the intensive margin.

Appendix figure A.10 reports the event study estimates of bureaucrat effects on a dummy indicating whether Korean exports exceeded 0. There is no indication of differential pre-trends. In event years 0 and 1, new ability increases the likelihood of positive exports in a given product by 5-7%, a sizable effect. Old ability decreases it by the same amount, suggesting that losing bureaucrat ability has symmetric effects to gaining such ability. This finding suggests that at least half the extensive margin effect of office openings is lost when a bureaucrat at the 20th percentile is appointed instead of one at the median.

Appendix figure A.11 reports the estimates using only the sample of products with ex-

tensive margin changes. For this sample, new ability increases exports by an even larger 22-31% while old ability decreases it by a similar amount.

Appendix table B.3 shows that the number of products with extensive margin changes remains roughly constant across decades. So the extensive margin response remains similarly important over time. However, appendix table B.3 shows an increase over time in the number of products for which only the intensive margin matters.[27] Appendix figure A.12 replicates figure 8 using data on only these products. As expected, the estimates become noisier. However, pre-trends remain absent. The point estimates go in the expected direction and are quantitatively similar to figure 8. Due to the decreased statistical power, only the coefficients on the old bureaucrat's effect remain statistically significant.

6 The Effect of Bureaucrat Experience

We saw in section 5 that a large share of the variation in Korean exports can be attributed to the directors of overseas export promotion offices. This raises the question whether the capacity of these bureaucrats can be built.

We isolate quasi-random variation in a bureaucrat's exposure to different products to estimate the causal effect of product-specific experience. We find that exports of a product increase by 3.0% if this switch implies an increase in product-specific experience, about the same as moving London as close to Seoul as Frankfurt is.

This is the first evidence regarding learning-by-doing as a channel for increasing bureaucratic capacity. It complements the existing literature on bureaucracy which has focused on selection and incentives. It further contributes to the literature by showing that bureaucrat ability is not uni-dimensional, but differs across products, or dimensions of the policy space. One bureaucrat may outperform their peers in promoting exports of textiles, while being outperformed in promoting exports of cars.

[27] The omitted category are products without any exports throughout the event horizon.

6.1 Identification: Quasi-Random Component of Experience

$$y_{ept} = \eta_{ep} + \lambda_{T(e),pt} + \tau_{et} + \sum_{k \neq -2} \beta_k \text{ increase}_{ep} \, \mathbf{1}\{t = T + k\} + \epsilon_{ecpt} \tag{8}$$

This section discusses our strategy to identify the causal effect of product-specific experience β. We do so by estimating equation (8), a reduced form event-study which aims to identify the causal effect on exports from a switch in bureaucrats that implies an increase in the quasi-random component of experience. For the new and old bureaucrat, this is given by $\text{instrument}_{b(c,T(e)),p}$ and $\text{instrument}_{b(c,T(e)-1),p}$. $T(e)$ indicates the year of the event e, defined as the first full year in which the new bureaucrat is in charge. As before, t indicates the observation year, c the country, p the 4-digit product, and $b(c,t)$ the bureaucrat assigned to country c in year t.

Equation (8) includes event $\times$ year fixed effects (τ_{et}). As each et has a unique bureaucrat, τ_{et} absorbs bureaucrat fixed effects and any experience that translates across products. We estimate the effect of experience within bureaucrat. Being able to include this level of fixed effects, highlights one advantage of our strategy based on product-specific experience. The explanatory variable of interest is increase_{ep}, a dummy indicating whether the new bureaucrat's experience in p is greater than the old bureaucrat's experience.

$$\text{experience}_{b(c,t),pt} = \sum_{k=0}^{2} \text{exports}_{C_1(b),p,T_1(b)+k} \tag{9}$$

We conceptualize a bureaucrat's experience as the exports to which a bureaucrat was exposed at the time of their first appointment, given by equation (9). $T_1(b)$ and $C_1(b)$ indicate the year and country of bureaucrat b's first appointment. As in the remainder of the paper, *exports* always refer to the inverse hyperbolic sine of exports.

Our multi-dimensional measure of experience allows us to address three main sources of endogeneity to identify β_k. First, a bureaucrat's first appointment may be endogenous if they are strategically appointed based on existing exports to that destination. We rule out that experience is due to such strategic appointment, by subtracting lagged exports from our measure of experience. Doing so means that our measure of experience is not due to differences in exports (of product p) between countries that existed in the three years prior

to a bureaucrat's first appointment.

$$\text{instrument}_{b(c,t),pt} = \sum_{k=0}^{2} \widehat{\text{exports}}_{C_1(b),p,T_1(b)+k} - \sum_{k=-3}^{-1} \widehat{\text{exports}}_{C_1(b),p,T_1(b)+k} \tag{10}$$

$$\widehat{\text{exports}}_{cpt} = \text{exports}_{cpt}^{\text{non-Korean}} \frac{\text{exports}_{-c,pt}}{\text{exports}_{-c,pt}^{\text{non-Korean}}} \tag{11}$$

Second, to avoid that our measure of experience is endogenous to bureaucrat' actions during their first appointment, we instrument for experience as described in equation (10). This instrument follows the same form as equation (9) but replaces actual exports with predicted exports, calculated according to equation (11).

To calculate predicted Korean exports, we use contemporaneous non-Korean exports to the same product-country. This captures a country's overall import demand. To make this measure of import demand relevant to Korea, this is normalized by the relative scale of Korean to non-Korean exports of the same product to *other* countries in the same year.

Third, bureaucrats' later appointments may be correlated with their experience gained during their first appointment. However, this is problematic only if bureaucrat appointments are endogenous to our instrumented experience, given by the differences across products in the change in non-Korean exports during a bureaucrat's first appointment. We follow two separate approaches to rule out such remaining correlation between experience and ϵ_{ecpt}, e.g. due to strategic appointment. The main approach we follow is to allow for product-year fixed effects $\lambda_{T(e),pt}$ that are differential by year of event – $T(e)$. The first concern this addresses is given by a mechanical relationship between our measures of experience and exports due to secular changes in Korea's exports of certain products over time. If a bureaucrat is first appointed in 1968, they gain more experience regarding the type of products that Korea was exporting in 1968 (e.g. textiles, not cars). This bureaucrat is more likely to be re-appointed in 1973 – when Korea still exported more textiles than cars – rather than 1993 – when cars had become much more important than textiles. This type of correlation is avoided by including year-product FE. Year-product FE further avoid spurious correlations due to the fact that Korean exports in later years are larger for any product or the fact that textiles always make up a larger share of Korean exports than do maize or crude oil.

The second concern we rule out is that bureaucrats may be re-appointed to countries with high demand for the products in which they are experienced. If product-specific experience matters, it seems intuitive this would be taken into account for re-appointments – even though our qualitative research suggests that product-specific experience is not a factor when deciding bureaucrat appointments. Further, for this to be an identification concern, re-appointment would need to take into account our instrumented measure of experience. As long as our instrumented measure of experience does not predict reappointments, a correlation between other components of a bureaucrat's experience and their re-appointments would not constitute a violation of our identifying assumptions. We include event-product FE to avoid attributing any effects to demand for a product that is time-invariant during the event horizon. They further avoid spurious correlations due to the fact that Korean exports to bigger importers are larger across all products.

To causally identify the effect of product-specific experience on exports, we again rely on a parallel trends assumptions and a SUTVA. The SUTVA is the same as regards office opening: The bureaucrats only affect exports to their country of appointment. The PTA requires that exports in the country to which a bureaucrat would be appointed follow parallel trends, whether or not $\Delta\text{instrument}_{b(c,T(e)),b(c,T(e)-1),p} > 0$.

Our event-study specification (8) allows us to investigate pre-trends which are informative about the plausibility of the PTA.

6.2 Results: Experience Increases Exports Moderately

Figure 13 plots the event-study estimates (β_k) obtained from equation (8). It shows that pre-trends are not statistically distinct from 0 and economically very small. Exports increase sharply in those products where the change in experience due to the switch in bureaucrats exceeds 0, i.e. the new bureaucrat is more experienced than the outgoing bureaucrat. Combining the post-periods, the point estimate is 0.0300 (0.0147). This means exports increase by 3% in products in which a bureaucrat is experienced relative to those products in which the bureaucrat is not experienced. As our estimates are within bureaucrat, the results are informative about shifts in the composition of exports rather than an average effect.

6.3 Robustness: Effect of Experience Stays Similar

This result remains robust when excluding small changes in experience. We do this by comparing the top and bottom tercile of changes in experience, and excluding the middle tercile. We find a similar-sized, slightly larger effect on exports. The same is true when comparing the top vs the bottom quartile. The effect also exists when excluding very large changes. We report the coefficients from comparing a third to second quartile change. This last comparison gives slightly smaller effects.

Appendix figure A.13 reports coefficients from a similar regression that distinguishes between changes in experience in the 1st, 2nd, 3rd, and 4th quartile. Our comparison category are products with a 1st quartile change in experience due to the switch in bureaucrat. We find that a 2nd quartile change only barely increases exports. A 3rd quartile change increases exports by 2.5%, while a 4th quartile change causes an increase of 5%. The lack of pre-trends allays concerns about the causality of these estimates. The ordering of these effect sizes further raise our confidence that our measure of experience affects exports in the expected manner.

6.4 Mechanism: Experience Increases Responsiveness of Exports to Market Conditions

This section considers variation within appointment × product. We do so foremost to explore our central mechanism: Do bureaucrats with greater experience increase the responsiveness of exports to market conditions?

In addition, by exploring variation within bureaucrat by product, we further allay concerns about our previous results on the effects of experience. While somewhat implausible given the rigid appointment lengths of KOTRA bureaucrats, appointments of a bureaucrat experienced in a product could still take place exactly when Korean exports of said products to this country would have increased sharply. This would need to be an unanticipated increase, given the absence of pre-trends from figure 13.

$$\begin{aligned}\text{exports}_{cpt,b(c,t)} = \eta_{ep} + \lambda_{T(e),pt} + \tau_{et} + \psi_d^0 \text{demand}_{cpt} + \psi_s^0 \text{supply}_{cpt} + \\ \psi_{d,\text{increase}}^0 \text{demand}_{cpt} \times \text{increase}_{ep} + \psi_{s,\text{increase}}^0 \text{supply}_{cpt} \times \text{increase}_{ep} + \\ \sum_{k \neq -2} \Big[\boldsymbol{\beta_k} \text{ increase}_{ep} + \psi_{dk} \text{demand}_{cpt} + \boldsymbol{\beta_k^{demand}} \text{ demand}_{cpt} \times \text{increase}_{ep} + \\ \psi_{sk} \text{supply}_{cpt} + \boldsymbol{\beta_k^{supply}} \text{ supply}_{cpt} \times \text{increase}_{ep} \Big] \mathbf{1}\{t = T + k\} + \epsilon_{ecpt}\end{aligned} \tag{12}$$

We estimate equation (12) and show that bureaucrat experience increases the elasticity of Korean exports to demand and supply . Throughout, we assume that other countries' exports are determined exogenously, i.e. Korean exports are not important enough to cause changes in overall exports of any product to a destination.

Figure 14 reports the results from this regression. We find evidence for the same mechanism discussed earlier for the increases in exports caused by bureaucrats with high fixed effects. Bureaucrats with experience regarding a product may increase exports because they are more effective at transmitting information regarding demand shocks to Korean exporters or helping them effectively react to such shocks.

As with bureaucrat fixed effects, we would like to say how much of experience's overall effect is mediated by this increased reactivity to market conditions. However, as highlighted by appendix figure A.14, our estimate of the main effect becomes very noisy in this specification.

7 Conclusion

Our paper closely links individual bureaucrats to exports, a variable important to economic growth and development. We find that offices openings increase exports by 38%. The importance of bureaucrats is illustrated as this effect is entirely offset if the bureaucrat in charge is at the 20th percentile of effectiveness. Moreover, we show that there is some scope for bureaucrats to acquire capacity on the job. Our findings have important implications for debates on industrial policy and the role of state capacity in economic development.

First, the bureaucrats we study engage in the implementation of an industrial policy. Our finding thus imply that implementation matters substantially in determining whether an industrial policy is successful. This adds nuance to the resurgent debate on industrial

policy. As we compare different bureaucrats who implement the same policy, our results highlights one important dimension of "how industrial policy should be carried out" instead of "whether governments should carry out industrial policy" (Juhász, Lane, and Rodrik, 2023). This focus on the "how" is especially pertinent as export promotion is a policy many governments choose to pursue, especially as part of a broader industrial policy (Juhász, Lane, Oehlsen, and Pérez, 2022).

Second, export promotion is quite distinct from the tasks studied by the existing literature on bureaucrats. Our paper highlights that there may be returns to developing countries who do not only build bureaucratic capacity domestically, but use it to support their firms as they navigate global markets. Further, tasks like export promotion that aim to identify and overcome frictions that constrain firm growth may not require Weber's impersonal bureaucrats following standardized processes, but something closer to an "entrepreneurial bureaucrat" (Mazzucato, 2013) making use of tacit knowledge.

Third, we find that our bureaucrats learn to promote exports of certain products when exogenously exposed to them. This suggests a potential path for building state capacity endogenously as bureaucrats acquire capacity as they are exposed to certain opportunities and problems (Hirschmann, 1958). However, it also points to potential path dependence in state capacity. A bureaucracy will be most effective at carrying out familiar tasks. Expanding into policy areas in which the bureaucracy has no (recent) experience builds capacity but is less likely to bring immediate policy success.

Finally, our findings are informative regarding narratives about South Korea's "export-led" growth from one of the poorest countries in 1960. Our findings support this narrative in the sense that the bureaucrats we study are not alleviating supply-side constraints but strongly target the demand for Korean products abroad. In this sense, our study points towards a role for policy aiming to create growth that is led by demand for exports. It should be noted that for such a policy to be effective, bureaucrats have to be well-informed – embedded (Evans, 1995) – regarding domestic firms' export capacity.

Bibliography

Abowd, J. M., R. H. Creecy, and F. Kramarz (2002). Computing person and firm effects using linked longitudinal employer-employee data. Technical Report 2002-06, Center for Economic Studies, U.S. Census Bureau.

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica 67*(2), 251–334.

Alfaro-Ureña, A., J. Castro-Vincenzi, S. Fanelli, and E. Morales (2023). Firm export dynamics in interdependent markets. Working paper.

Alfaro-Ureña, A., I. Manelici, and J. P. Vasquez (2022). The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages. *The Quarterly Journal of Economics 137*(3), 1495–1552.

Amsden, A. H. (1989). *Asia's Next Giant: South Korea and Late Industrialization*. Number 9780195076035 in OUP Catalogue. Oxford University Press.

Anderson, J. (2011). The gravity model. *Annual Review of Economics 3*(1), 133–160.

Andrews, M., L. Gill, T. Schank, and R. Upward (2008). High wage workers and low wage firms: Negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society*.

Atkin, D., L. Boudreau, R. Dix-Carneiro, A. Khandelwal, B. McCaig, P. Medina, A. Morjaria, L. Pascali, B. Rijkers, and M. Startz (2022). International trade. *VoxDevLit 4*.

Atkin, D., A. Costinot, and M. Fukui (2021). Globalization and the ladder of development: Pushed to the top or held at the bottom? Working Paper 29500, National Bureau of Economic Research.

Atkin, D. and D. Donaldson (2022). The role of trade in economic development. Chapter Handbook of International Economics (Volume 5), pp. 1–59.

Atkin, D., A. K. Khandelwal, and A. Osman (2017). Exporting and Firm Performance: Evidence from a Randomized Experiment. *The Quarterly Journal of Economics 132*(2), 551–615.

Bandiera, O., M. Best, A. Khan, and A. Prat (2021). The Allocation of Authority in Organizations: A Field Experiment with Bureaucrats. *The Quarterly Journal of Economics 136*, 2195–2242.

Bertrand, M., R. Burgess, A. Chawla, and G. Xu (2020). The glittering prizes: Career incentives and bureaucrat performance. *Review of Economic Studies 87*(2), 626–655.

Bertrand, M. and A. Schoar (2003, 11). Managing with Style: The Effect of Managers on Firm Policies*. *The Quarterly Journal of Economics 118*(4), 1169–1208.

Besley, T., R. Burgess, A. Khan, and G. Xu (2022). Bureaucracy and development. *Annual Review of Economics 14*(1), 397–424.

Best, M. C., J. Hjort, and D. Szakonyi (2023). Individuals and Organizations as Sources of State Effectiveness. *American Economic Review 113*(8), 2121–2167.

Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2023). How much should we trust estimates of firm effects and worker sorting? *Journal of Labor Economics 41*(2), 291–322.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Card, D., A. R. Cardoso, and P. Kline (2016). Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women. *The Quarterly Journal of Economics 131*(2), 633–686.

Card, D., J. Heining, and P. Kline (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics 128*(3), 967–1015.

Cheng, T.-J., S. Haggard, and D. Kang (1998). Institutions and growth in korea and taiwan: The bureaucracy. *Journal of Development Studies 34*(6), 87–111.

Chetty, R., J. N. Friedman, and J. E. Rockoff (2014a, September). Measuring the impacts of teachers i: Evaluating bias in teacher value-added estimates. *American Economic Review 104*(9), 2593–2632.

Chetty, R., J. N. Friedman, and J. E. Rockoff (2014b, September). Measuring the impacts of teachers ii: Teacher value-added and student outcomes in adulthood. *American Economic Review 104*(9), 2633–79.

Choi, J. and A. A. Levchenko (2021). The long-term effects of industrial policy. Working Paper 29263, National Bureau of Economic Research.

Evans, P. (1995). *Embedded Autonomy: States and Industrial Transformation*. Princeton University Press.

Feenstra, R. and J. Romalis (2014). International prices and endogenous quality. *The Quarterly Journal of Economics 129*(2), 477–527.

Fenizia, A. (2022). Managers and Productivity in the Public Sector. *Econometrica 90*(3), 1063–1084.

Goldberg, P. and T. Reed (2020). Demand-side constraints in development: The role of market size, trade, and (in)equality. NBER Working Papers 27286, National Bureau of Economic Research, Inc.

Haggard, S., B. Kim, and C. Moon (1991). The transition to export-led growth in south korea: 1954-1966. *Journal of Asian Studies 50*(4), 850–873.

Hausmann, R., J. Hwang, and D. Rodrik (2007). What you export matters. *Journal of Economic Growth 12*(1), 1–25.

Head, K. and T. Mayer (2014). Chapter 3 - gravity equations: Workhorse,toolkit, and cookbook. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4 of *Handbook of International Economics*, pp. 131–195. Elsevier.

Hirschmann, A. O. (1958). *The strategy of economic development*. Yale University Press.

Jones, L. P. and S. Il (1980). *The Transition to Export-led Growth in South Korea: 1954-1966*. Studies in the modernization of the Republic of Korea, 1945-1975. Harvard East Asian monographs ; v. 91.

Juhász, R. (2018). Temporary protection and technology adoption: Evidence from the napoleonic blockade. *American Economic Review 108*(11), 3339–76.

Juhász, R., N. Lane, E. Oehlsen, and V. Pérez (2022). The who, what, when, and how of industrial policy: A text-based approach. Working paper.

Juhász, R., N. Lane, and D. Rodrik (2023). The new economics of industrial policy. *Annual Review of Economics*.

Kim, B. and E. F. Vogel (2011). *The Park Chung Hee era: The transformation of South Korea*. Harvard University Press.

Kline, P., R. Saggio, and M. Sølvsten (2020). Leave-Out Estimation of Variance Components. *Econometrica 88*(5), 1859–1898.

Lane, N. (2022). Manufacturing revolutions: Industrial policy and industrialization in south korea. Working paper.

Lederman, D., M. Olarreaga, and L. Payton (2010). Export promotion agencies: Do they work? *Journal of Development Economics 91*(2), 257–265.

Liu, E. (2019). Industrial Policies in Production Networks. *The Quarterly Journal of Economics 134*(4), 1883–1948.

Lugard, F. (1926). *The Dual Mandate in British Tropical Africa*.

Mazzucato, M. (2013). *The entrepreneurial state*. Penguin Books.

Metcalfe, R. D., A. B. Sollaci, and C. Syverson (2023). Managers and productivity in retail. Working Paper 31192, National Bureau of Economic Research.

Munch, J. and G. Schaur (2018). The effect of export promotion on firm-level performance. *American Economic Journal: Economic Policy 10*(1), 357–87.

Otero, C. and P. Muñoz (2022). Managers and public hospital performance. Working paper.

Rambachan, A. and J. Roth (2023, 02). A More Credible Approach to Parallel Trends. *The Review of Economic Studies 90*(5), 2555–2591.

Sant'Anna, P. H. and J. Zhao (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics 219*(1), 101–122.

Shim, Y. and J. Choi (2022). Technology adoption and late industrialization. Working Paper WP033, STEG.

Studwell, J. (2013). *How Asia works: success and failure in the world's most dynamic region*. Grove Press and Grove/Atlantic, Inc.

Volpe Martincus, C. and J. Carballo (2008). Is export promotion effective in developing countries? firm-level evidence on the intensive and the extensive margins of exports. *Journal of International Economics 76*(1), 89–106.

Volpe Martincus, C. and J. Carballo (2010). Beyond the average effects: The distributional impacts of export promotion programs in developing countries. *Journal of Development Economics 92*(2), 201–214.

Volpe Martincus, C. and J. Carballo (2012). Export promotion activities in developing countries: What kind of trade do they promote? *The Journal of International Trade & Economic Development 21*(4), 539–578.

Wade, R. (1990). *Governing the Market: Economic Theory and the Role of the Government in East Asian Industrialization*. Princeton University Press.

Woo-Cumings, M. (Ed.) (1999). *The Developmental State*. Cornell University Press.

Xu, G. (2018). The costs of patronage: Evidence from the british empire. *American Economic Review 108*(11), 3170–98.

Figures and Tables

Figure 1: Growth in Korean Exports

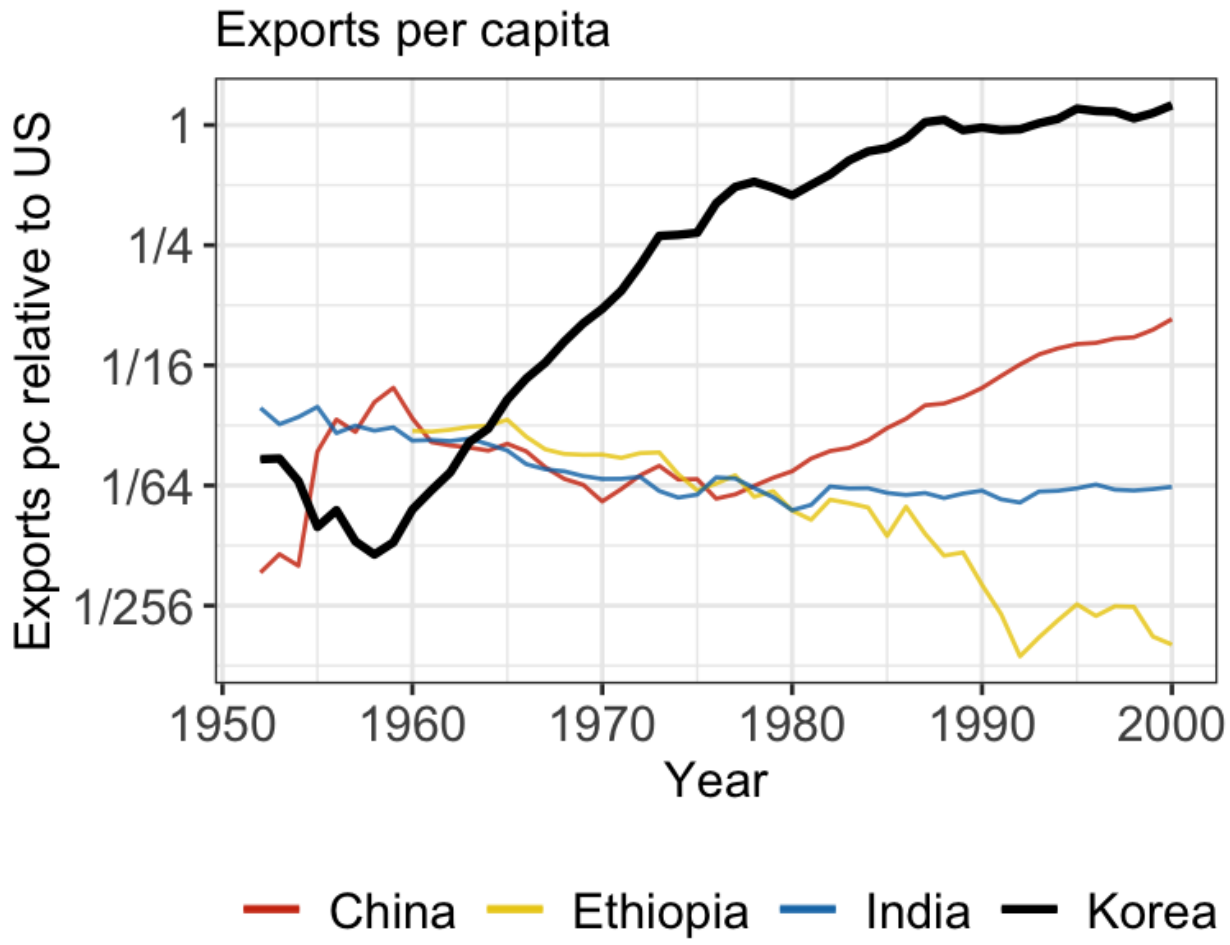

Notes: The figure displays Exports per capita relative to the U.S. the years 1952 to 2000 for Korea and a selected group of other countries. Data on exports and population obtained from International Monetary Fund (2023): Direction of Trade Statistics.

Figure 2: Growth in number of countries with export promotion (EP) offices

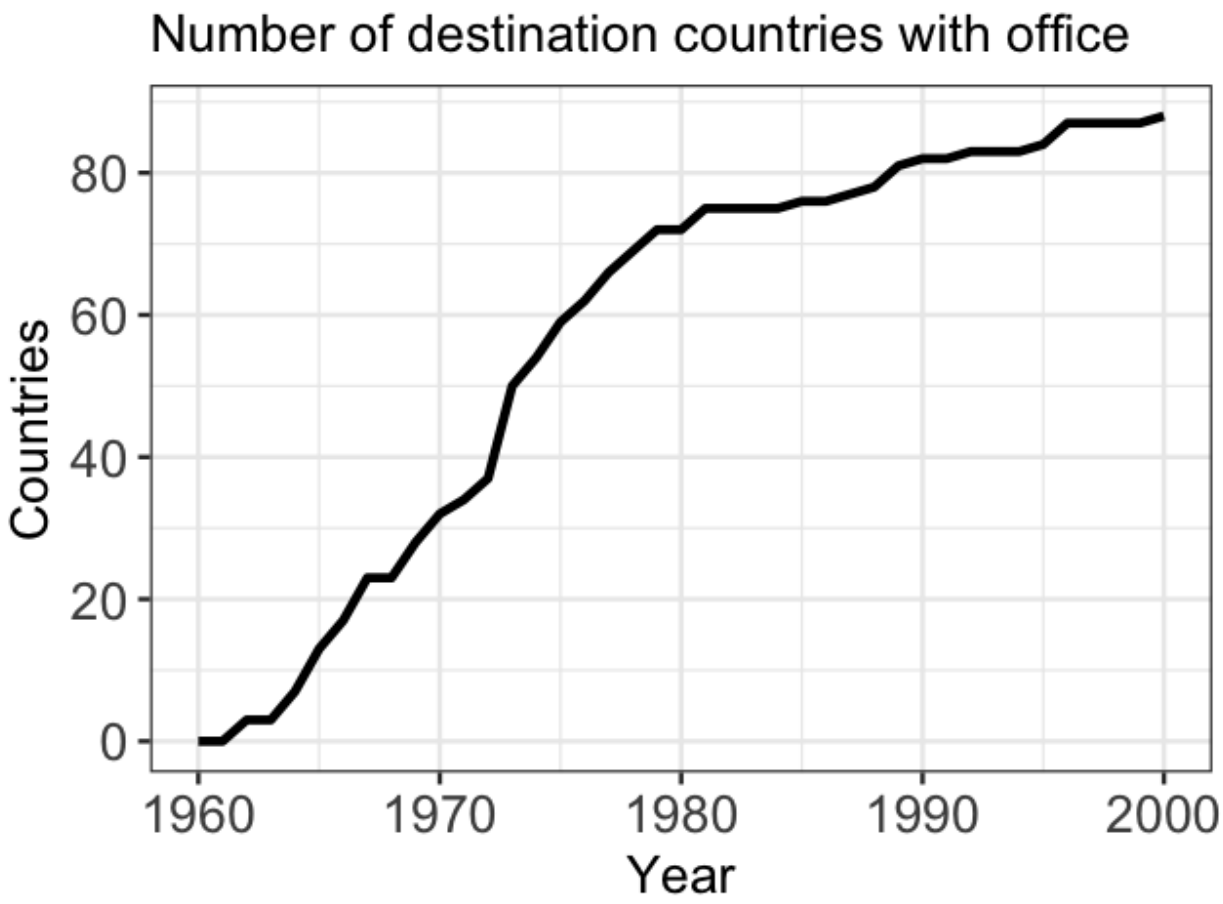

Notes: This figure presents the number of countries with an overseas export promotion office opening up until each year.

Figure 3: The roll-out of KOTRA offices to countries.

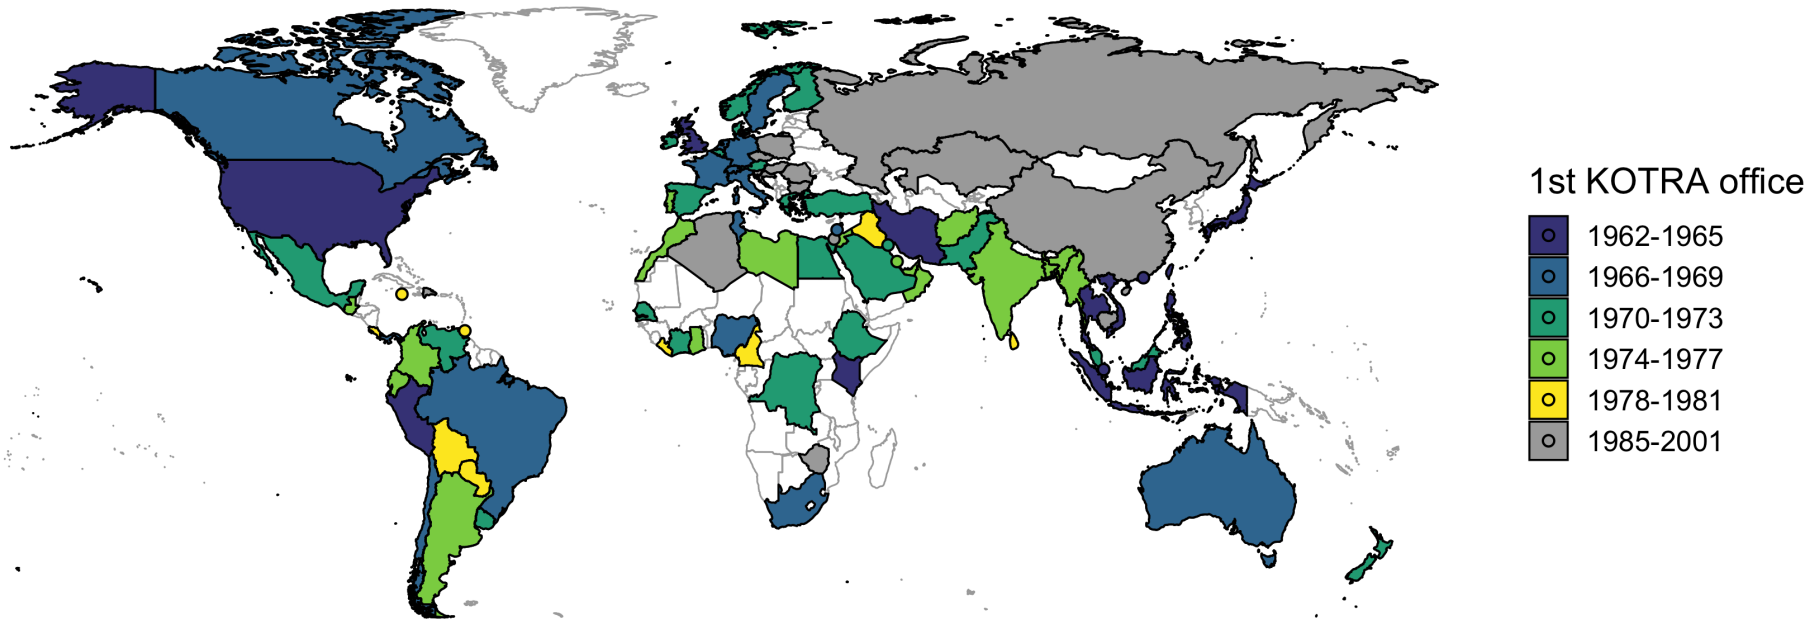

Notes: Colored countries have an office opening between 1962 and 2001. Different colors indicate the year in which the first office opened in a given country. There was a rapid roll-out until 1981 and a plateauing afterwards.

Figure 4: Event-study estimates of the effect of office opening on Korean exports.

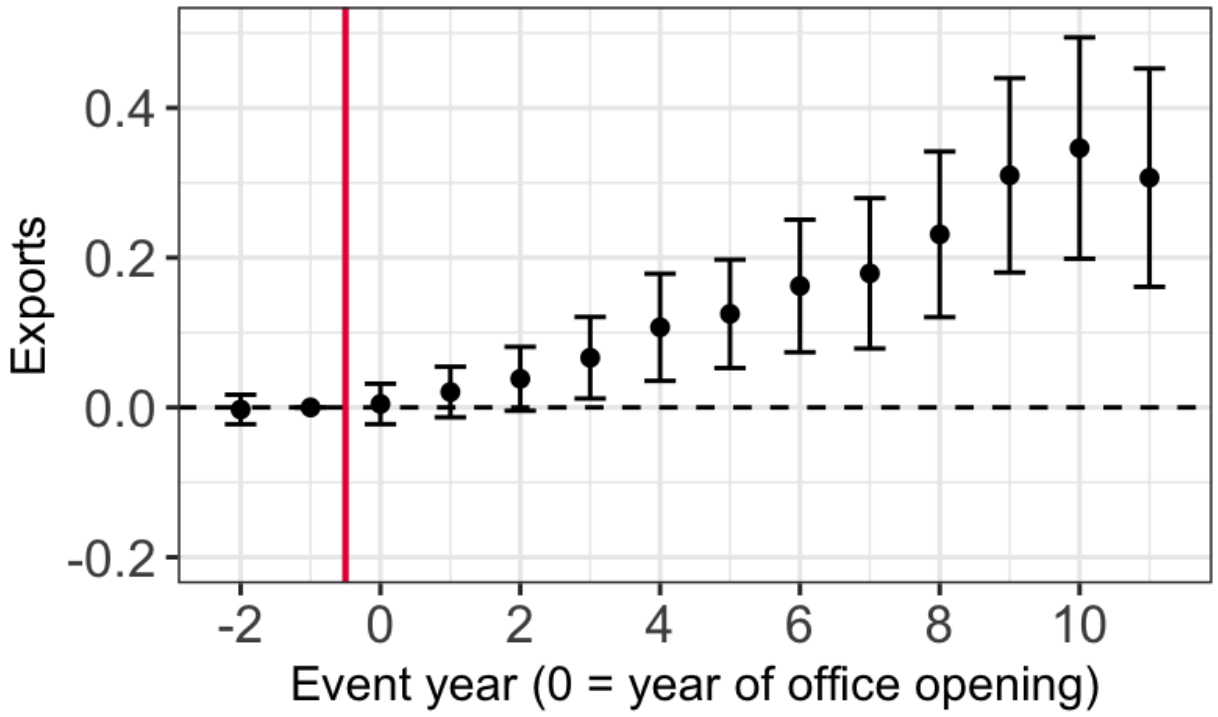

Notes: The outcome variable is the inverse hyperbolic sine of Korean exports to the country-year in question. An observation is at the product-country-year. Point estimates and standard errors are obtained from estimating equation (1). This relies on a never-treated control group. Standard errors clustered at the country-level are reported around each point estimate.

Figure 5: Europe: Openings follow pre-determined market size

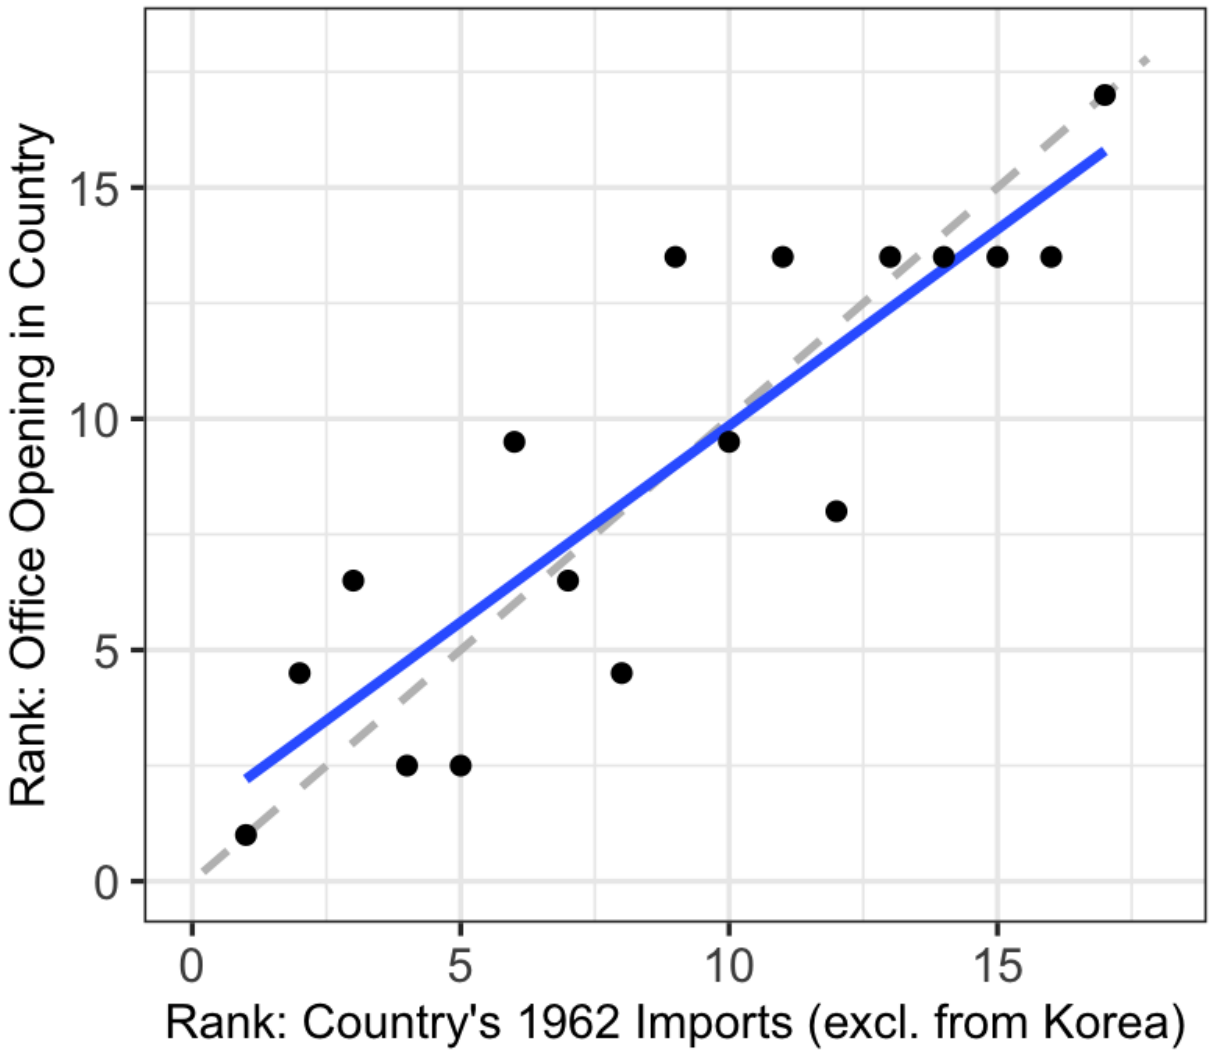

Notes: Each dot corresponds to a European country that received a KOTRA office during the main roll-out of offices (1962-1981). The x-axis gives each country's rank in terms of 1962 imports, excluding imports from Korea. The y-axis gives each country's rank in terms of the order of their office openings. The solid blue line gives the linear fit using 1962 market size to predict the order of office openings. The rank correlation between 1962 imports and office opening is 0.87. The dashed gray line gives the 45-degree line, where the two ranks are exactly equal. This is the case for the UK (rank 1) and Portugal (rank 17). When multiple countries have the same opening year, we assign the average rank to them. For example, Italy and the Netherlands get the second and third offices. As these openings occur in the same year, both have rank 2.5.

Figure 6: CDF of bureaucrat fixed effects.
Moving from 20th to 50th percentile: 0.46. 50th to 90th: 0.42

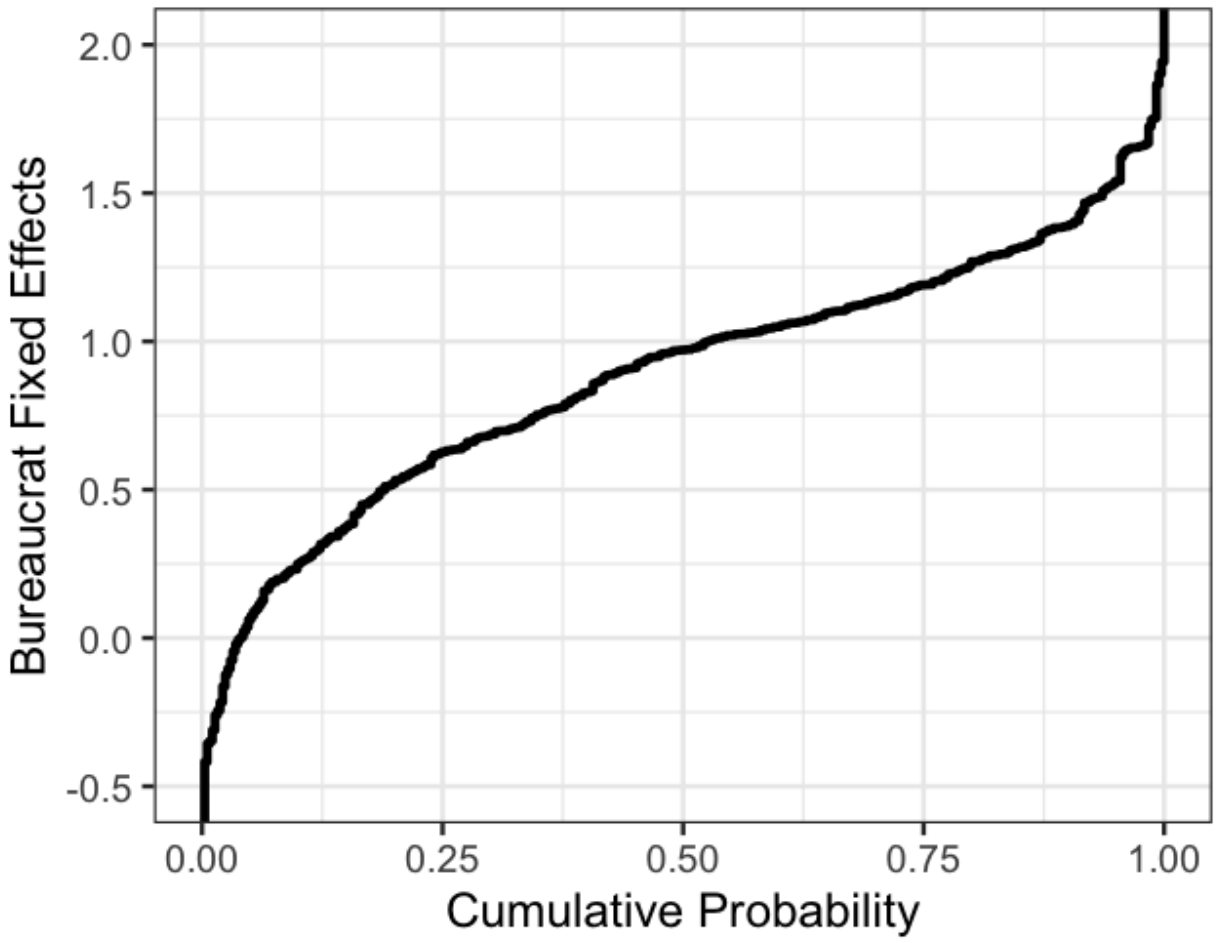

Notes: The figure shows the cumulative density function of bureaucrat fixed effects.

Figure 7: Bureaucrat switches cause jumps in exports without differential pre-trends.

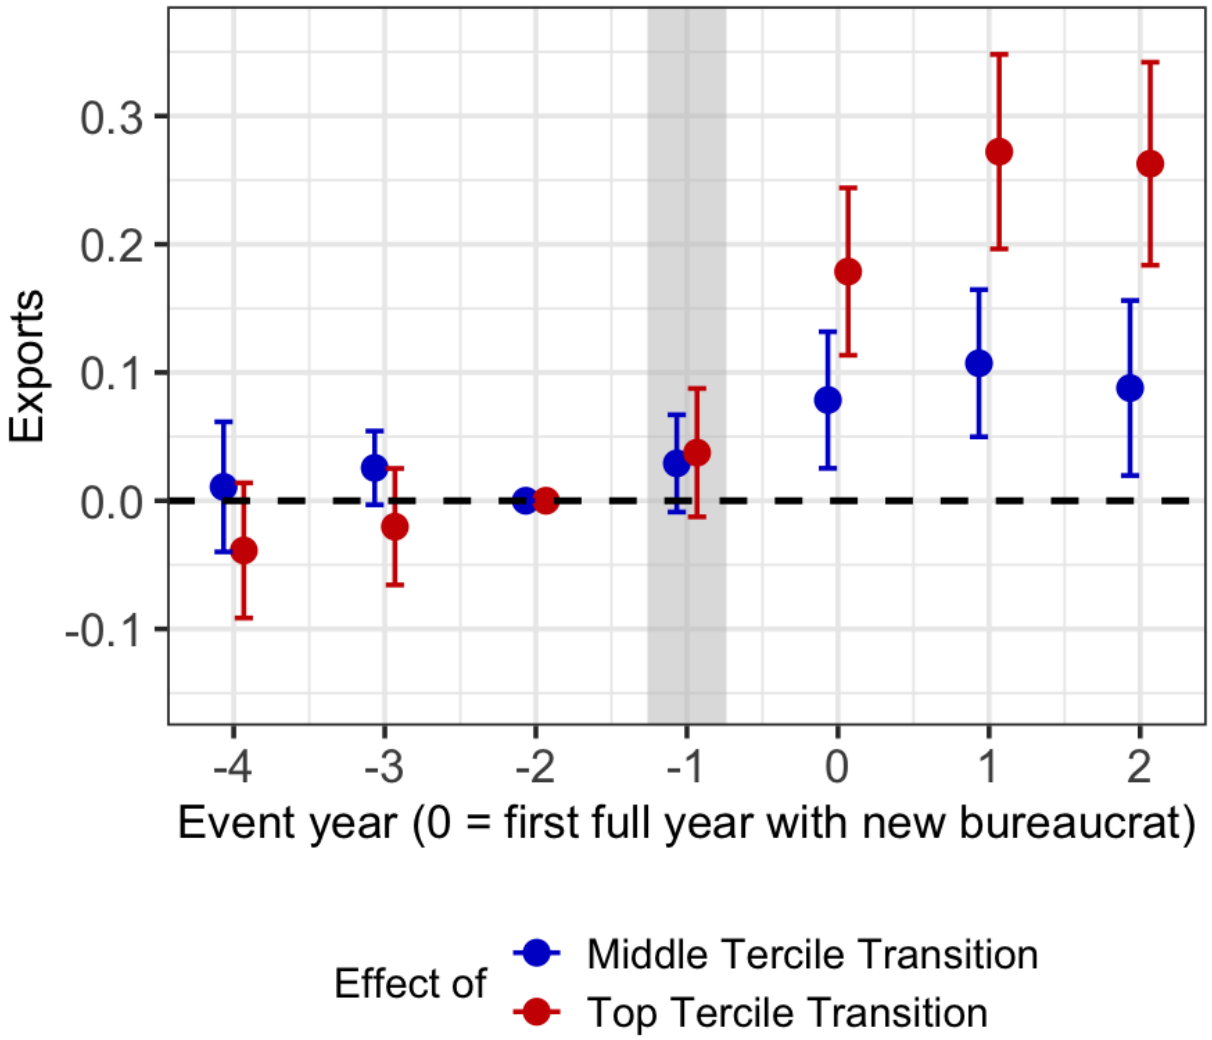

Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects on exports around the time that the bureaucrat heading a country office changes. The estimates are $\hat{\beta}_k$ and $\hat{\delta}k$ obtained from estimating equation (5). The dependent variable is the inverse hyperbolic sine of exports to the country of the switch between bureaucrats. The switch occurs in year -1. Transitions are categorized into terciles depending on the change in fixed effects implied by the switch in bureaucrats in year -1. The omitted category is a transition in the bottom tercile. The omitted year is -2, the last full year with the old bureaucrat.

Figure 8: Event study estimates: Effect on exports due to switch in bureaucrat.
No diff. pre-trends, symmetric jumps up (down) upon *gaining* (*losing*) bureaucrat ability.

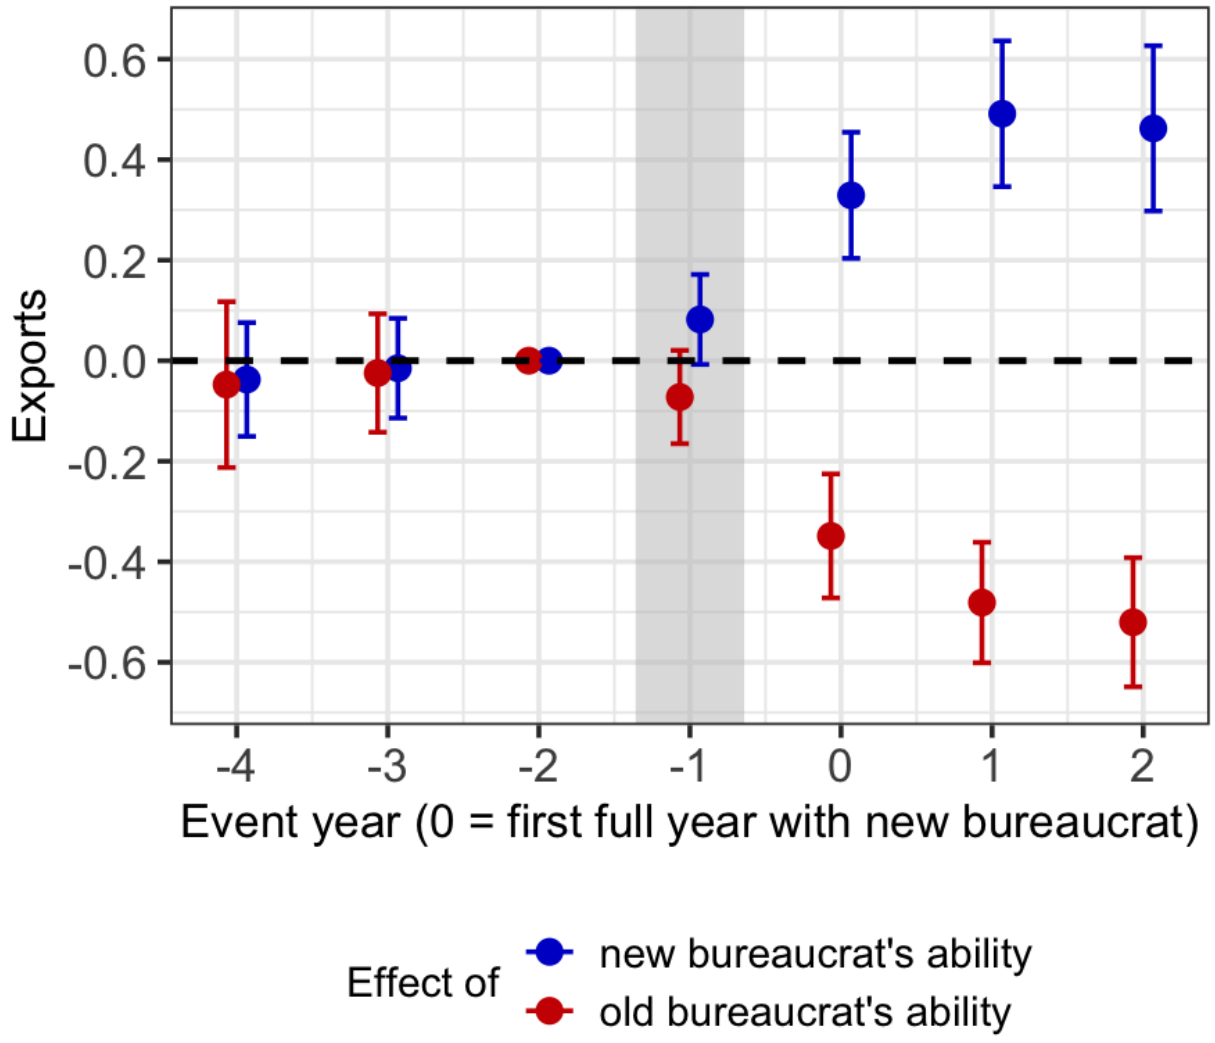

Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects on exports around the time that the bureaucrat heading a country office changes. These estimates are $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (6). The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat heads the country office. Year -2 is the last full year that the old bureaucrat headed the office. The y axis measures the effect of bureaucrat effectiveness on exports. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure 9: Diagnostics: Threat to identification from misspecification.

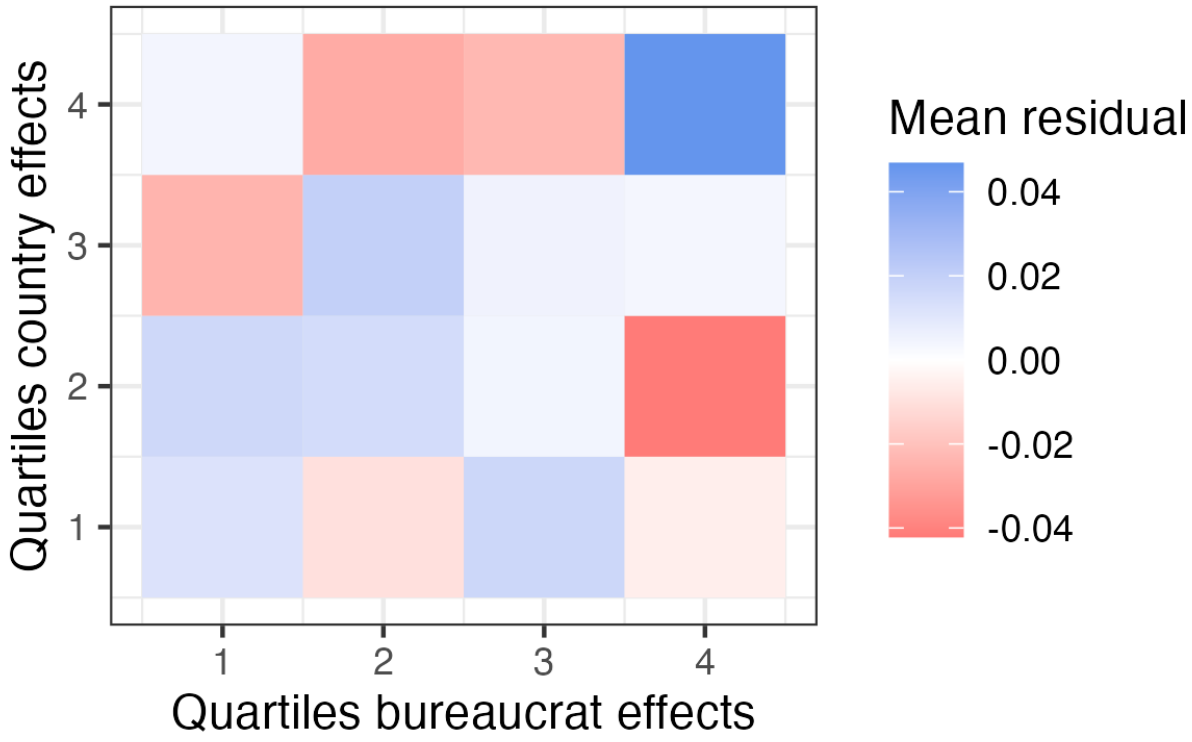

Notes: This figure shows mean residuals from equation 3 with cells defined by quartiles of estimated bureaucrat effect, interacted with quartiles of estimated country effect.

Figure 10: Bureaucrat fixed effects and exports: In and Out of Sample

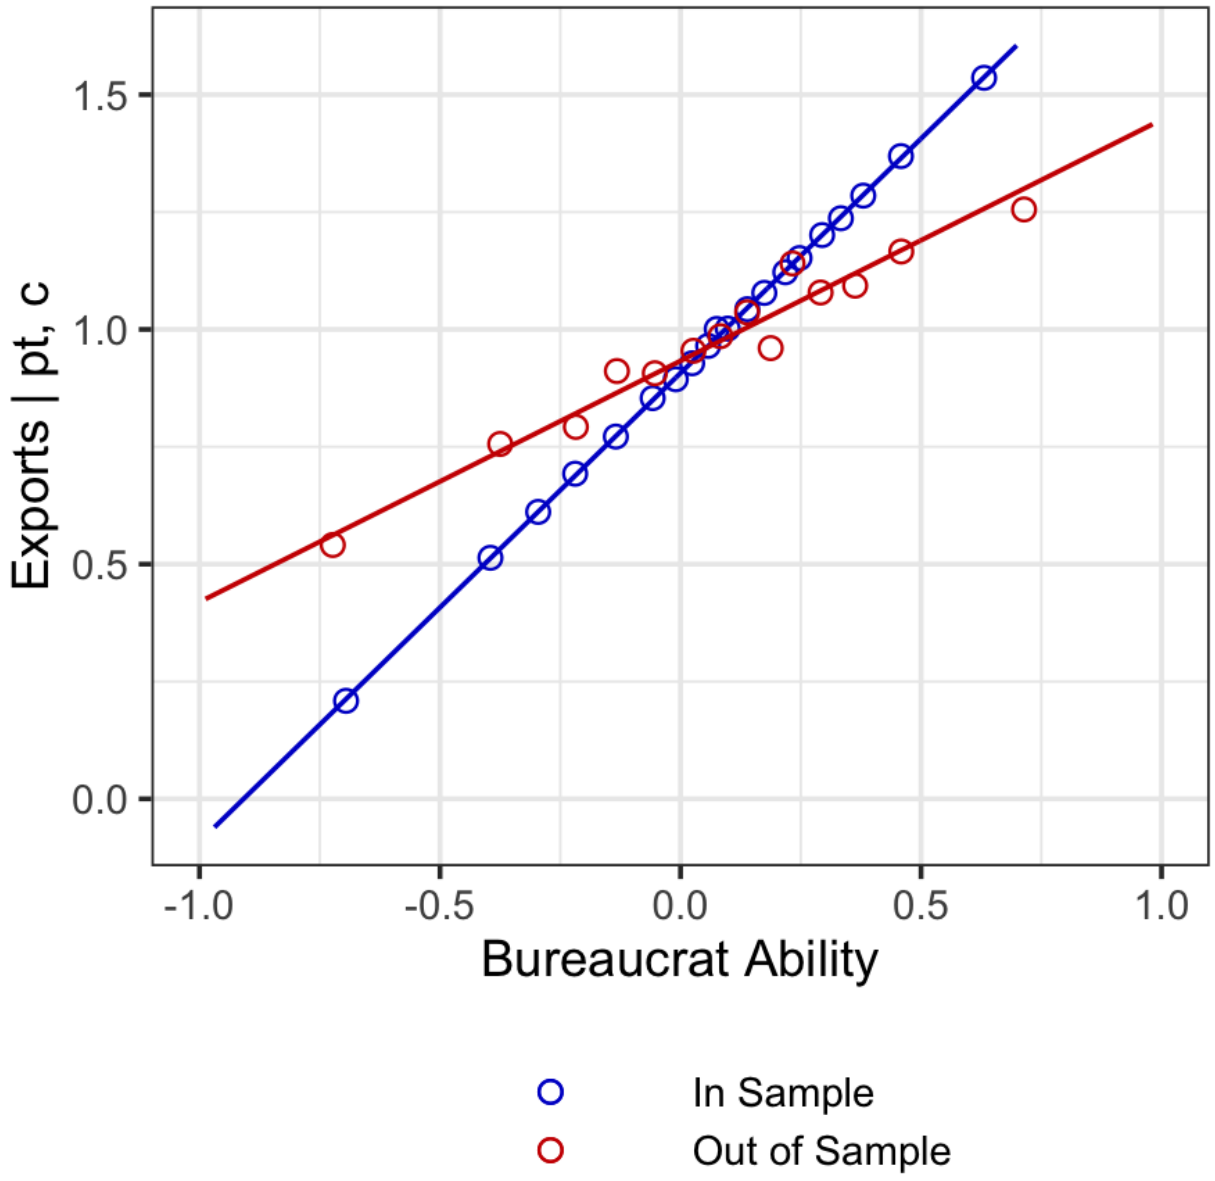

Notes: The figure displays a binned scatterplot. The y-axis shows exports after subtracting product-year fixed effects (pt) and country-year fixed effects. The two above fixed effects, as well as in-sample bureaucrat ability (fixed effects) are estimated using equation (3) and all country-years. Hence, by construction, each in-sample dot lies on a 45-degree line. This also means that in-sample fixed effects translate one-to-one into higher exports. Out-of-sample fixed effects are estimated only using other countries in estimating the fixed effects. This means to predict exports to the UK, we obtain the fixed effects on a data set using all country-years, except the UK. The slope of a regression of residualized exports on these out-of-sample, i.e. *other country*, fixed effects is 0.52.

Figure 11: Event study estimates: Decomposition
Good bureaucrats increase exports where demand (supply) are growing.

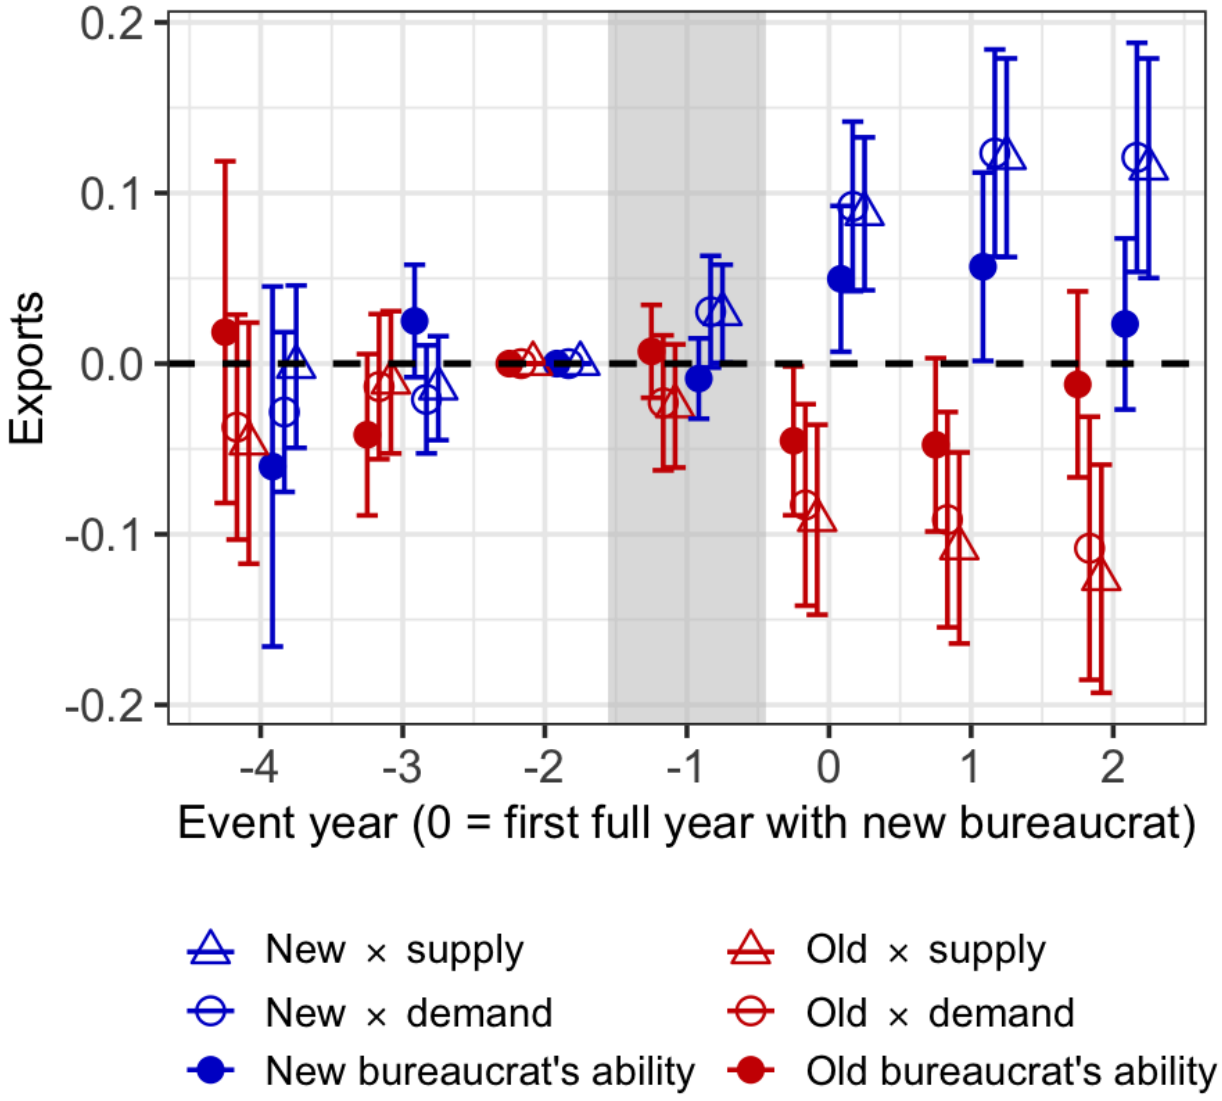

Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects when interacted with two kinds of shocks. The plotted coefficients are estimates of β_k, β_k^{demand}, and β_k^{supply} as well as δ_k, δ_k^{demand}, and δ_k^{supply} obtained from regressions of equation (7). The solid circles give the main effects. The hollow circles give the interaction with exports of the same product to the same destination by other countries (β_k^{demand}, δ_k^{demand}), our proxy for this destination's product-specific demand. The triangles give the interaction with Korean exports of the same product to the other destinations (β_k^{supply}, δ_k^{supply}), our proxy for Korea's product-specific supply. The horizontal axis indicates the years relative to a bureaucrat's appointment. Year 0 is the first full year that the new bureaucrat heads the country office. Year -2 is the last full year that the old bureaucrat headed the office. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure 12: Bureaucrat effect by number of appointments in career.
2+ appointments: Less bureaucrats with negative effects

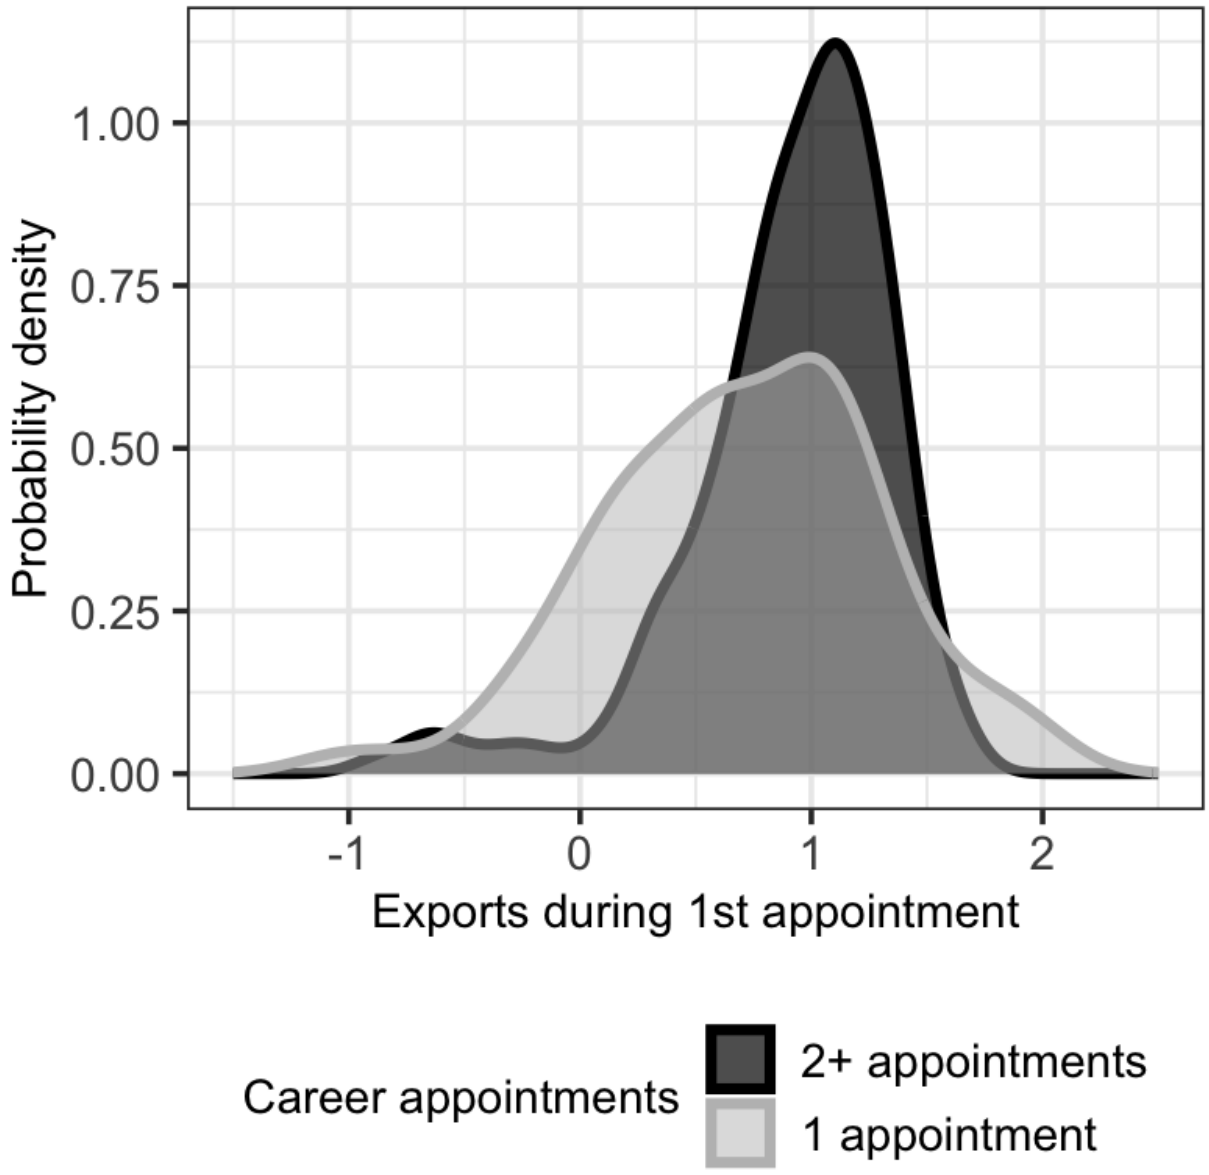

Notes: The figure shows the probability density function of residualized exports during bureaucrats' first appointments. It does so separately for bureaucrats who have 2+ appointments over the course of their career and for bureaucrats who have one career appointment. The distribution of exports under the latter group has a much fatter left tail.

Figure 13: Event study – Effect of increase in quasi-random experience$_p$

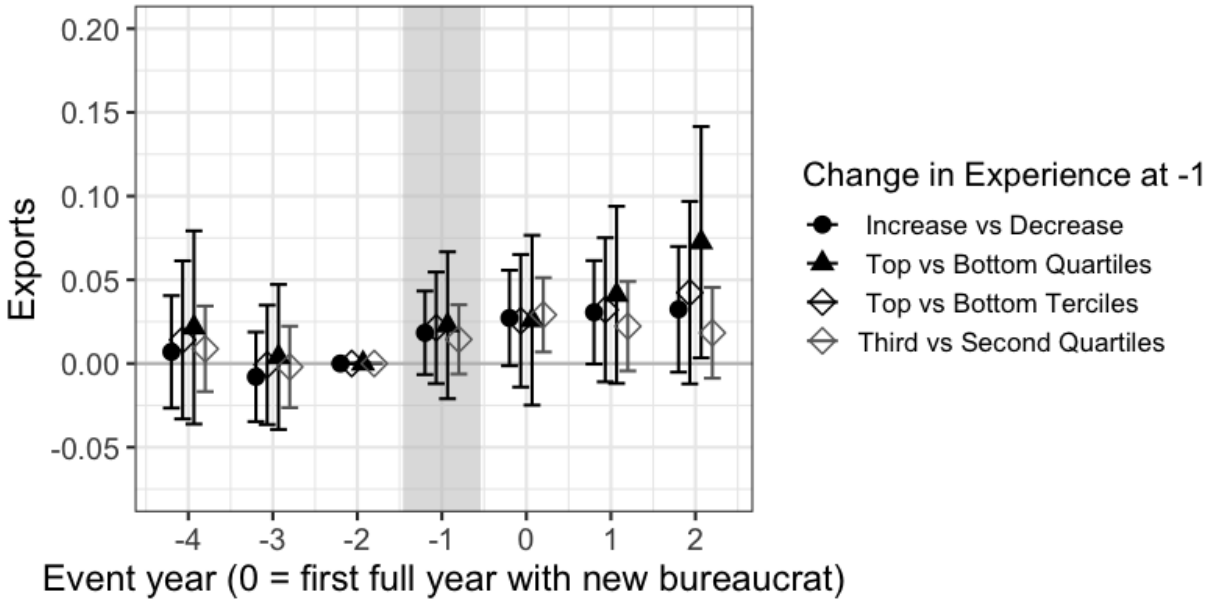

Notes: The figure shows the estimated effect of the change in the quasi-random component of bureaucrat experience on exports around the time that the bureaucrat heading a country office changes. These estimates are $\hat{\beta}_k$ obtained from estimating equation (8). The solid dots indicate the effect of an increase in experience compared to a decrease. The other symbols indicate slight variation on the definition of the change in experience. These results are within event-year, so they compare those products where the change in bureaucrat implies an increase in experience vs those where it implies a decrease. The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat heads the country office. Year -2 is the last full year that the old bureaucrat headed the office. The y axis measures the effect of bureaucrat experience on exports.

Figure 14: Event study – Quasi-random experience$_p$ increases reaction to demand

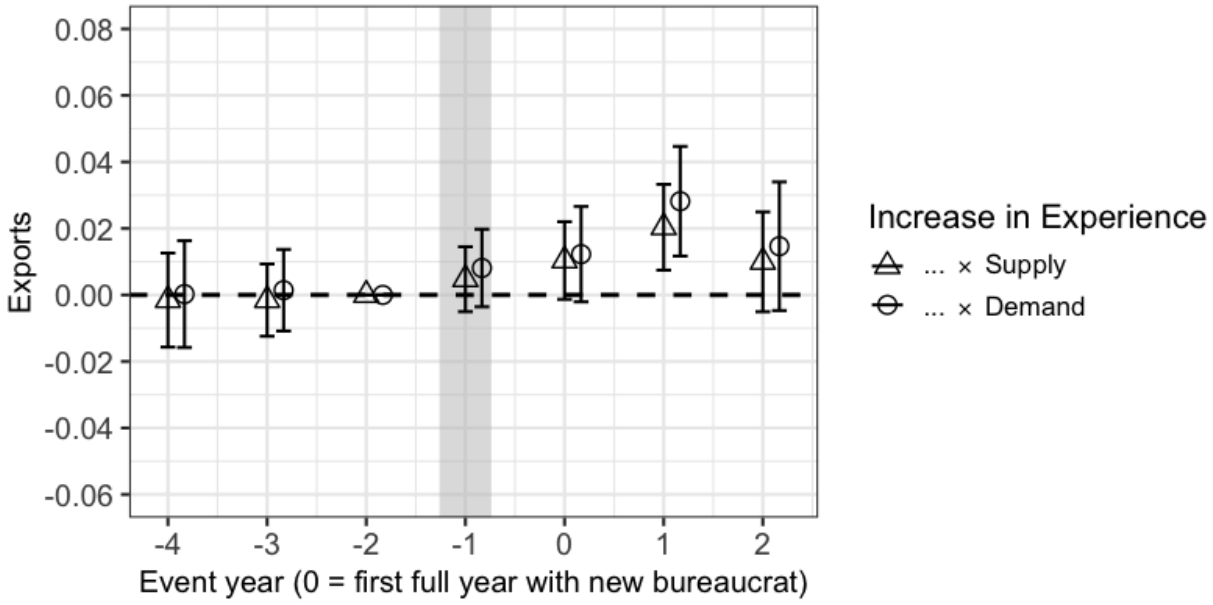

Notes: The figure shows the estimated effect of the change in the quasi-random component of bureaucrat experience when interacted with two kinds of shocks. The plotted coefficients are estimates of β_k^{demand}, and β_k^{supply} (12). The hollow circles give the interaction with exports of the same product to the same destination by other countries (β_k^{demand}), our proxy for this destination's product-specific demand. The triangles give the interaction with Korean exports of the same product to the other destinations (β_k^{supply}), our proxy for Korea's product-specific supply. The horizontal axis indicates the years relative to a bureaucrat's appointment. Year 0 is the first full year that the new bureaucrat heads the country office. Year -2 is the last full year that the old bureaucrat headed the office.

Table 1: Appointments Descriptives.

	Full Sample (1)	Country-Level Analysis (2)
Directors	475	397
Countries/Offices	138	86
Directors $>$ 1 Office over the Sample Period	252	194
Countries $>$ 1 Director over the Sample Period	121	82
Events / Distinct Appointments	974	728
Country-years		2060
Observations		1,772,452

An observation is a product-country-year. A country is included for all the years that it has an office and is linked to a bureaucrat. A product is included for all the years in which Korea exported it to any country. Product refers to 4-digit SITC Rev. 2 codes. All directors and countries in the country-level analysis are part of the same connected set. Restricting the analysis to this connected set is natural as only a single country and a single bureaucrat are outside the largest connected set.

Table 2: Variance decomposition of exports into bureaucrat and country components

		Actual data				Placebo check: Bureaucrats randomly shuffled to countries			
		All bureaucrats		Bureaucrats with ≥ 2 appointments		All bureaucrats		Bureaucrats with ≥ 2 appointments	
		Component (1)	% Share (2)	Component (3)	% Share (4)	Component (5)	% Share (6)	Component (7)	% Share (8)
Var(exports|*pt*), spell-level		0.732	100	0.737	100	0.737	100	0.736	100
Var(bureaucrat)		0.100	13.71	0.056	7.60	0.006	0.77	0.006	0.81
Var(country)		0.722	98.60	0.695	94.29	0.591	80.19	0.589	80.07
Cov(bureaucrat, country)		-0.088	-12.04	-0.045	-6.15	-0.005	-0.67	-0.003	-0.44
Var(bureaucrat+country)		0.646	88.24	0.659	89.45	0.586	79.59	0.588	79.94
Var(exports|*pt*), raw		4.404		4.645		4.360		4.343	
Number of observations		1703465		1222986		1757034.0		1228255.6	
Number of bureaucrats		380		184		389.2		182.7	
by no. of spells in sample:	1	200		4		209.0		2.8	
	2	96		96		99.1		98.3	
	3	56		56		53.8		54.9	
	4	24		24		21.5		21.1	
	5	4		4		5.8		5.7	
Number of countries		75		75		78.7		78.4	

The results of variance decomposition exercise according to equation (4). Columns (1)-(4) use actual data while columns (5)-(8) use data where bureaucrats are randomly shuffled to countries, preserving the number of appointment spells in the data for each bureaucrat. For columns (3), (4), (7), and (8), an initial sample restriction of bureaucrats with at least two appointments is applied. The limited mobility bias correction method follows Kline, Saggio, and Sølvsten (2020) and is implemented via the algorithm of Bonhomme, Holzheu, Lamadon, Manresa, Mogstad, and Setzler (2023) It is possible that there are bureaucrats with only one spell in the sample even when the sample is pre-restricted to bureaucrats with at least two appointments, because some spells drop out when constructing the leave-one-spell-out connected set for the Kline, Saggio, and Sølvsten (2020) method. Since the algorithm is based on numerical approximations of the traces of large matrix inverses, there is a small degree of randomness in the decomposition results. There is also additional randomness in columns (5)-(8) arising from the random shuffling of bureaucrats. Thus, we report the averages of 100 iterations for all columns.

A Appendix Figures

Figure A.1: Distribution of appointment durations.
Median and modal duration: 36 months.

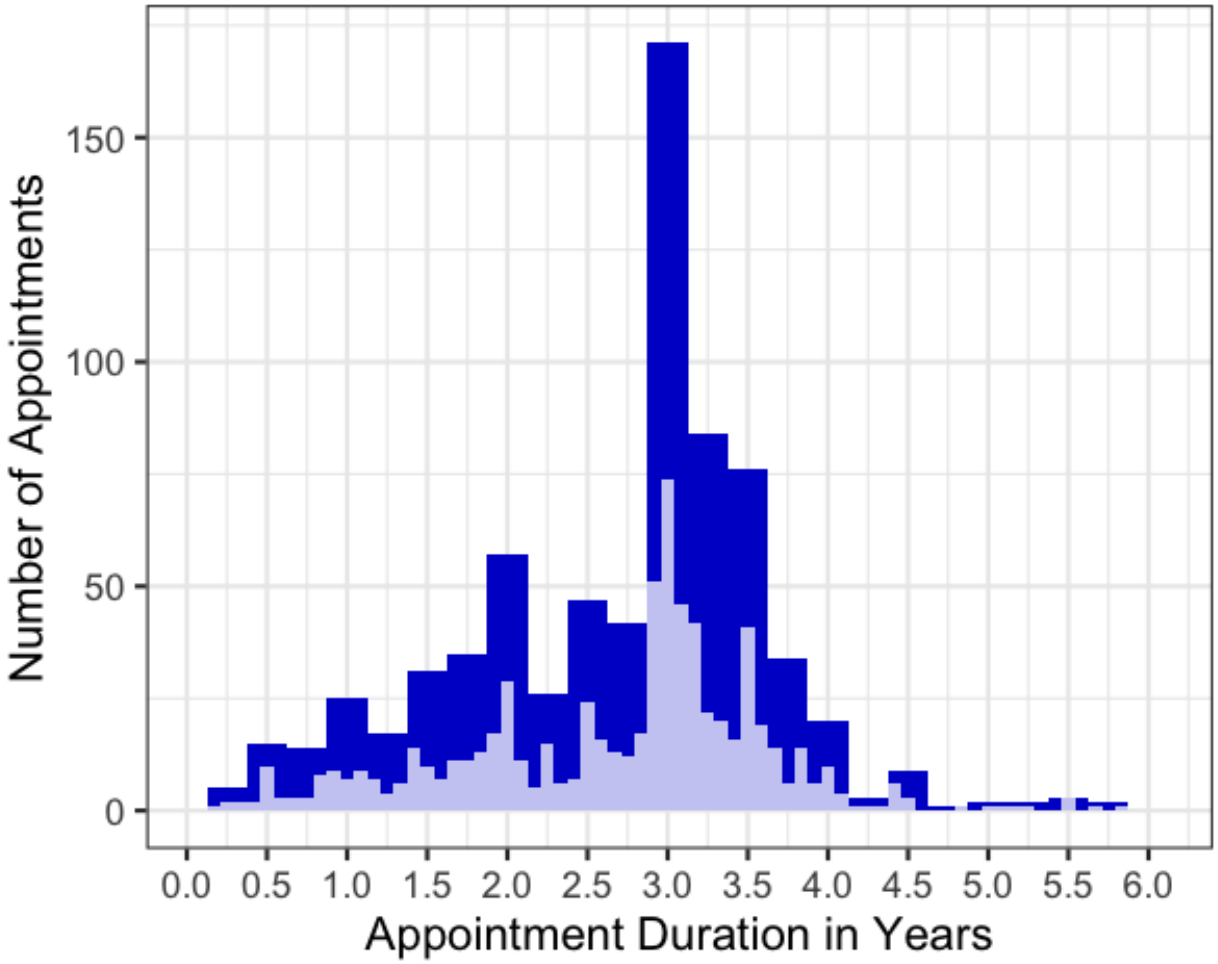

Notes: This figure represents the distribution of appointment durations. The blue bars indicate the number of appointments by quarterly duration whereas the white bars do so for the number of appointments by monthly duration. Hence, as each quarter contains multiple months, the blue bars always (weakly) exceed the white ones. E.g there are 82 appointments that last 3 years and 1 quarter. These are comprised of 42 appointments that last 3 years and 2 months, 21 appointments that last 3 years and 3 months, and 19 appointments that last 3 years and 4 months.

Figure A.2: Distribution of gap lengths.
Median: 29 months. Mode: 30 months.

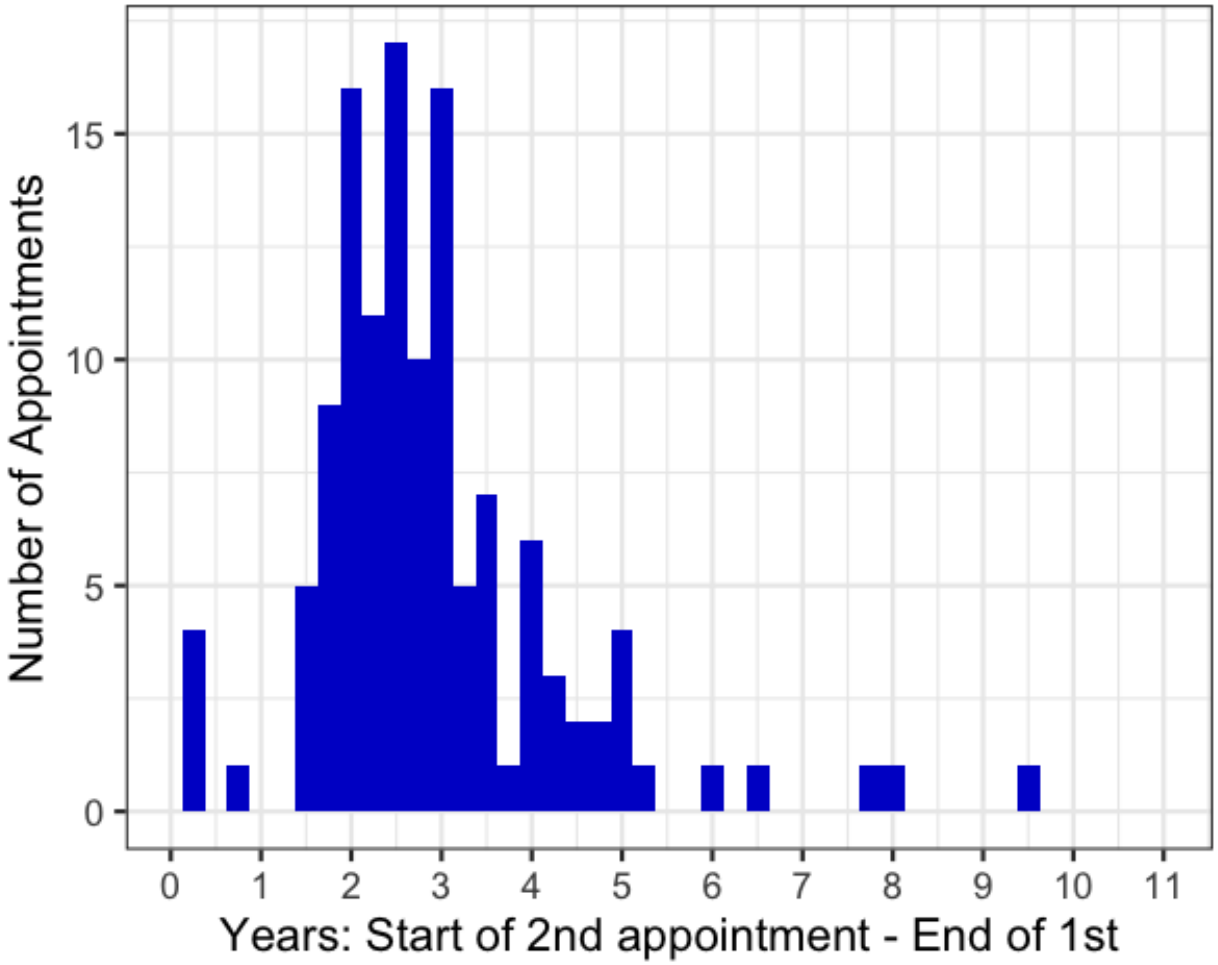

Notes: This figure represents the distribution of the duration of gaps between appointments. The blue bars indicate the number of gaps by quarterly duration.

Figure A.3: Targeting of export promotion activity by product.
Export promotion activity moves in parallel with national industrial policy

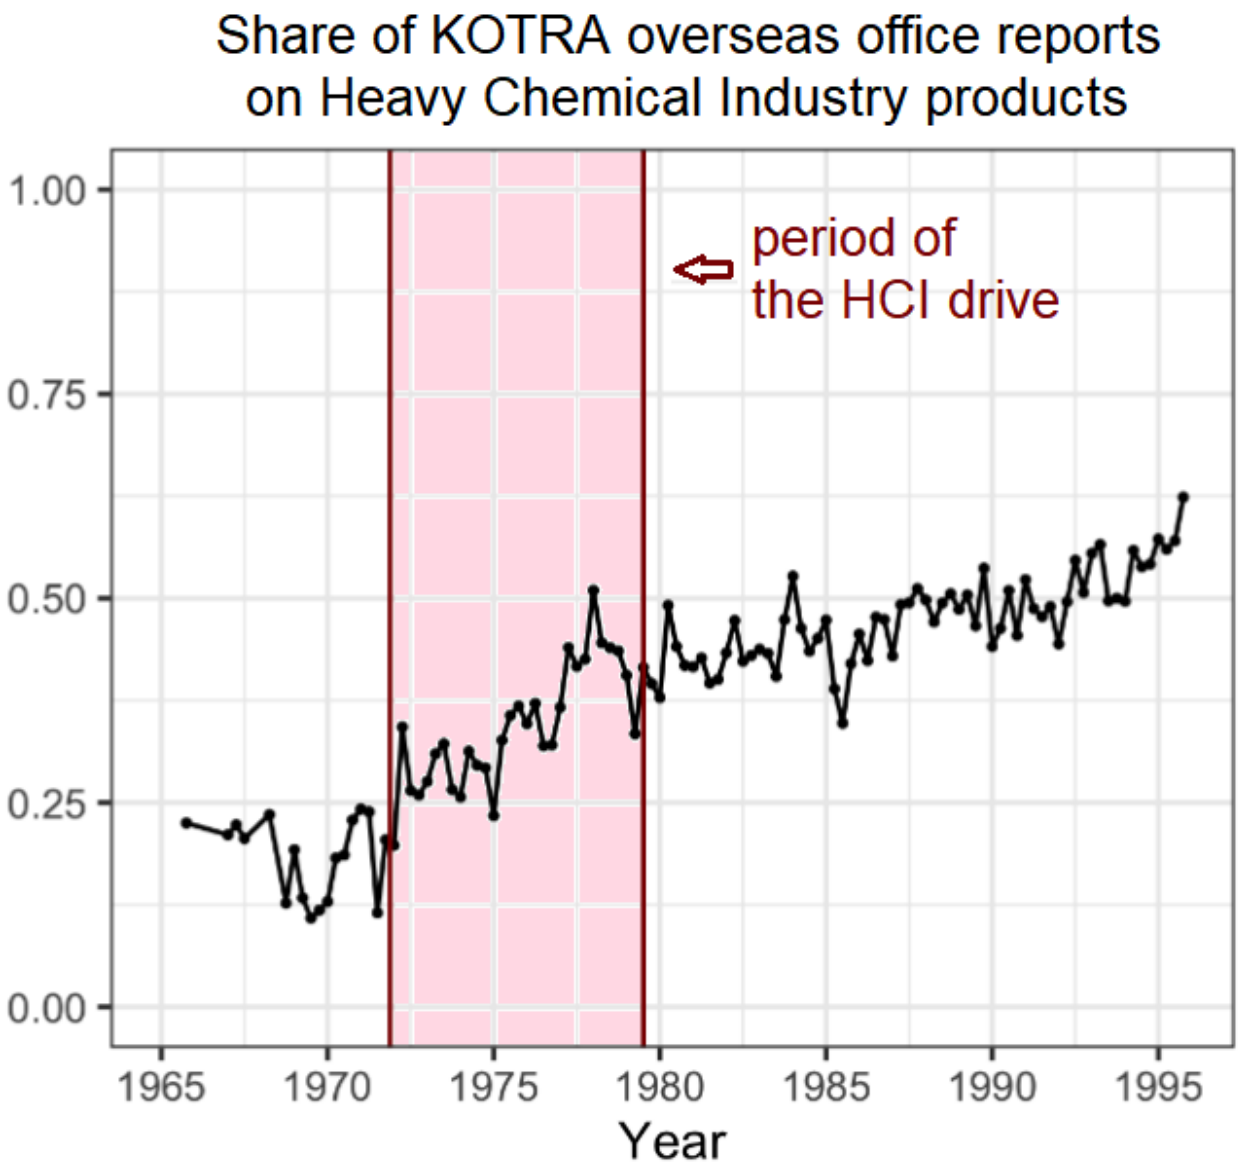

Notes: Targeting of EP activity by product. For each quarter, the y-axis presents the share of overseas office reports that could be linked to an HCI product relative to the number of reports that could be linked to any product.

Figure A.4: Robustness: Controls, sample, placebo

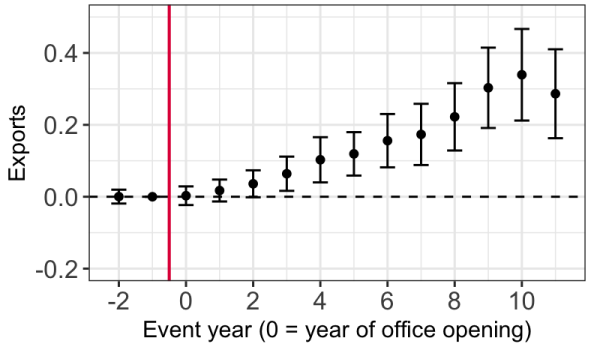

(a) Controlling for non-Korean exports.

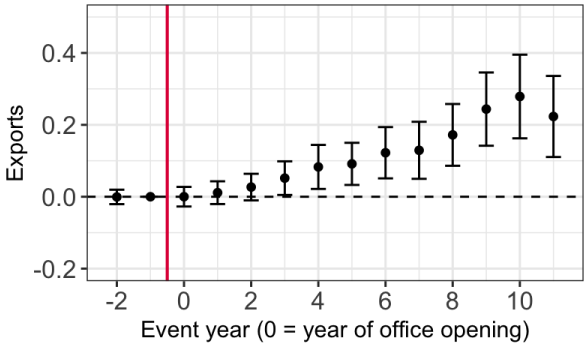

(b) Controlling for non-Korea exp. × year

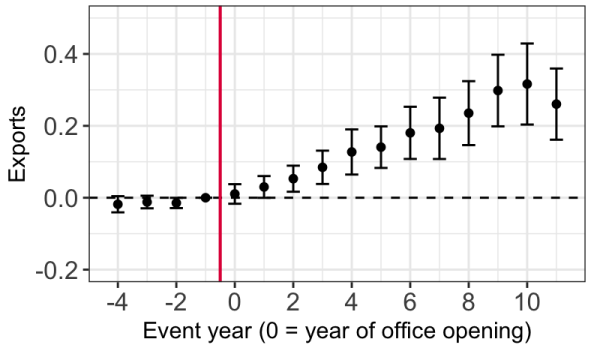

(c) Korean exports as outcome. Openings from 1966. Never-treated as control group.

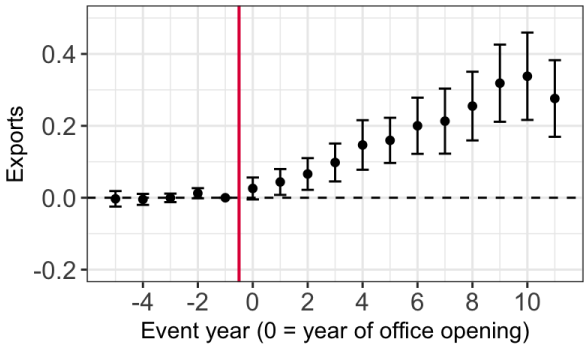

(d) Korean exports as outcome. Openings from 1967. Never-treated as control group.

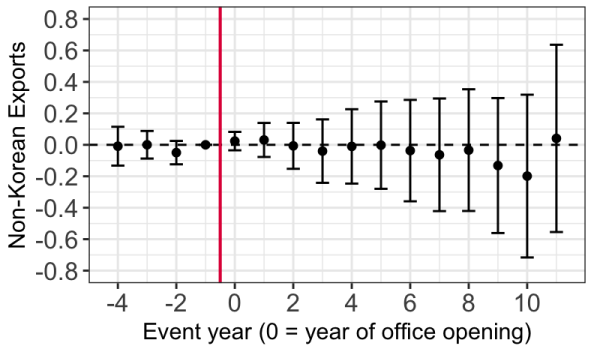

(e) Non-Korean exports as outcome. Openings from 1966. Never-treated as control group.

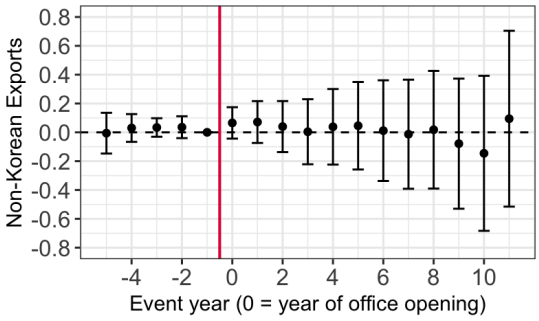

(f) Non-Korean exports as outcome. Openings from 1967. Never-treated as control group.

Notes: For panels (a)-(d), the outcome variable is the inverse hyperbolic sine of Korean exports to the country-year in question. For panels (e) and (f), the outcome is given by the inverse hyperbolic sine of non-Korean exports to the same country-year. An observation is at the product-country-year. Point estimates and standard errors are obtained from estimating equation (1, relying on a never-treated control group. Standard errors clustered at the country-level are reported around each point estimate. A product is included for all the years in which Korea exported it to any country. Product refers to 4-digit SITC Rev. 2 codes.

Figure A.5: Robustness: opening with not-yet-treated control

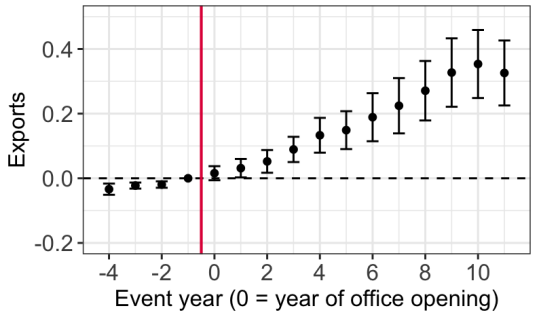

(a.i) CSA estimate, unconditional PTA. 0 periods of anticipation.

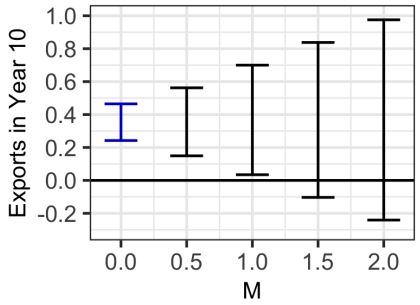

(a.ii) Sensitivity to PTA violation relative to largest pre-treatment violation of PT.

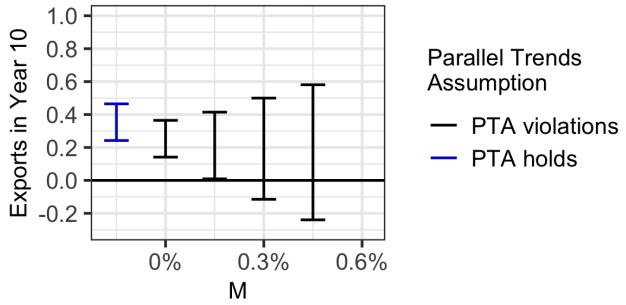

(a.iii) Sensitivity to PTA violations only bounding the extent to which the slope may change across consecutive periods.

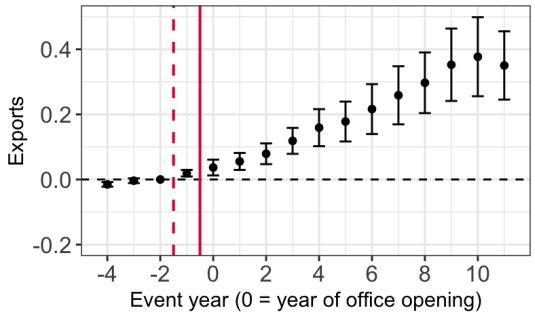

(b.i) CSA estimate, unconditional PTA. 1 periods of anticipation.

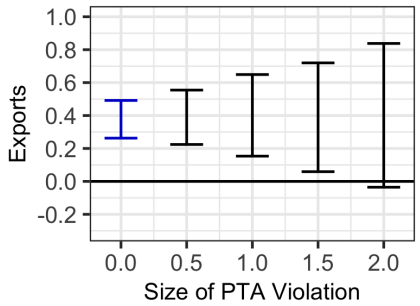

(b.ii) Sensitivity to PTA violation relative to largest pre-treatment violation of PT.

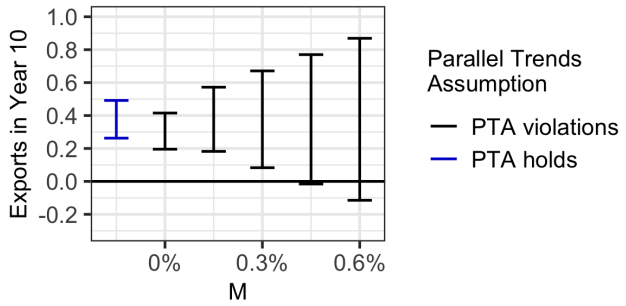

(b.iii) Sensitivity to PTA violations only bounding the extent to which the slope may change across consecutive periods.

Notes: The outcome variable is the inverse hyperbolic sine of Korean exports to the country-year in question. The top panels report results assuming no anticipation. The bottom panel do so assuming one period of anticipation. Point estimates in (a.i) and (b.i), give the aggregation of treatment-group-specific estimates of the average treatment effect (ATT) using a "not-yet-treated" control group and Callaway and Sant'Anna (2021) estimator for Difference-in-Difference settings with staggered roll-out using the doubly-robust estimators form Sant'Anna and Zhao (2020). Bootstrapped standard errors are obtained clustering at the level of the destination country. Panels (a.ii-iii) report the sensitivity of the estimate in (a.i) to violations of the parallel trends assumption Rambachan and Roth (2023) It zooms in on the estimates in year 10. Panels (b.ii-iii) do the same for the estimate in (b.i). The blue bar in each panel corresponds to the 95% confidence interval of the year-10-estimate in the left panel. The black bars represent corresponding 95% confidence intervals when allowing for per-period violations of parallel trends. In panels (a.ii) and (b.ii), we bound the maximum post-treatment violation of parallel trends between consecutive periods by M times the maximum pre-treatment violation of parallel trends. In panels (a.iii) and (b.iii), we impose that the differential trends evolve smoothly over time by bounding the extent to which its slope may change across consecutive periods. Here, M represents the largest allowable change in the slope of an underlying linear trend between two consecutive periods. A product is included for all the years in which Korea exported it to any country. Product refers to 4-digit SITC Rev. 2 codes.

Figure A.6: Extensive Margin Effect of Office Opening

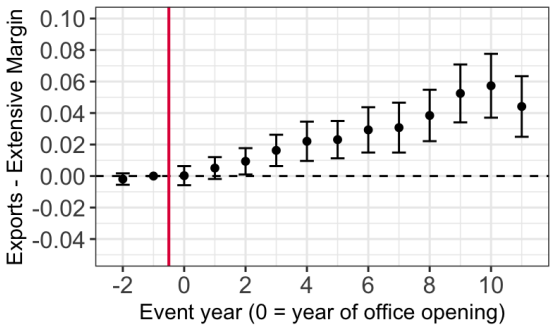

(a) Never-treated as control group. Openings from 1964.

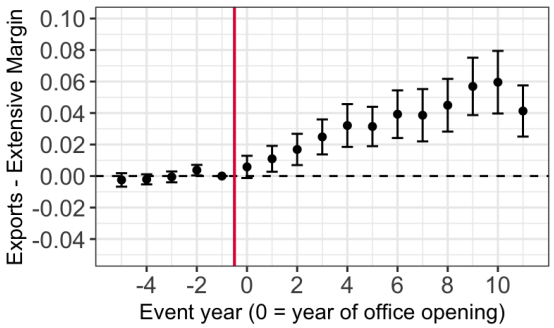

(b) Never-treated as control group. Openings from 1967.

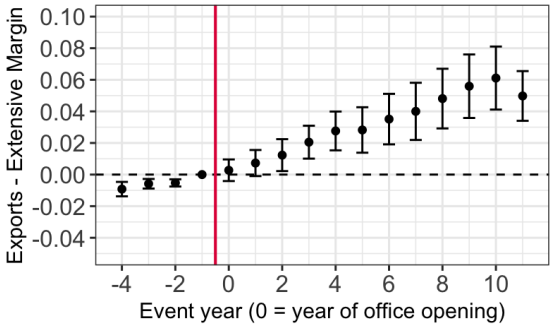

(c) "Not-yet" control. 0 period anticipation.

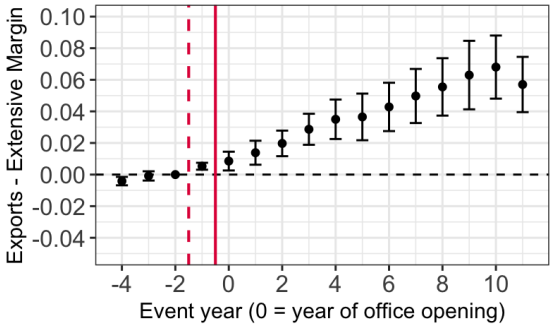

(d) "Not-yet" control. 1 period anticipation.

Notes:

Figure A.7: Event-study estimates of the effect of office opening on KOTRA activity

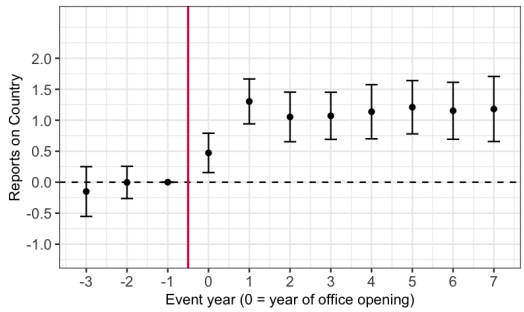

(a) Effect on reports. Never-treated as control.

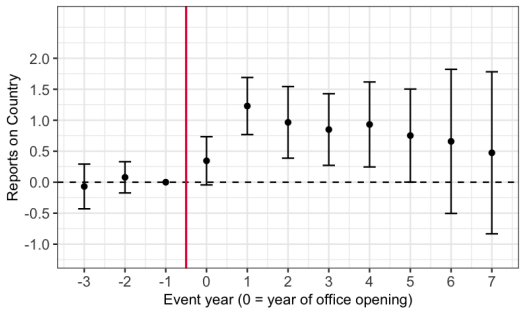

(b) Effect on reports. "Not-yet" as control.

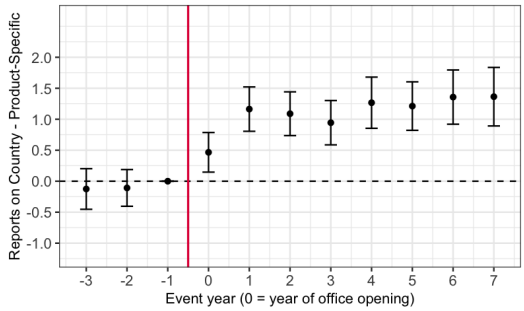

(c) Effect on reports. Never-treated as control.

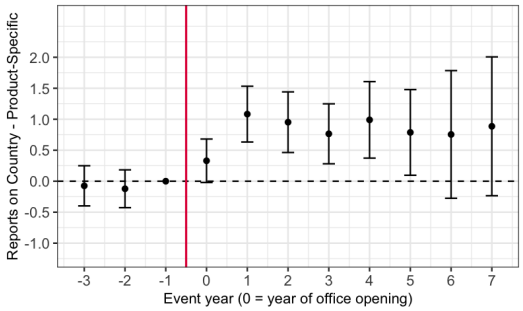

(d) Effect on reports. "Not-yet" as control.

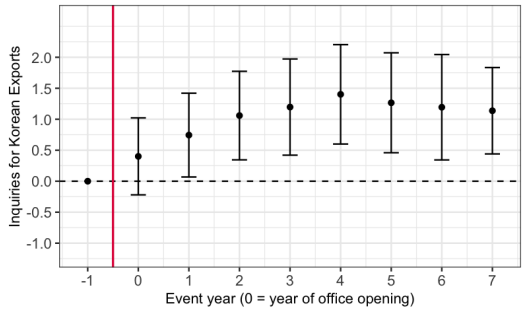

(e) Effect on inquiries. Never-treated as control.

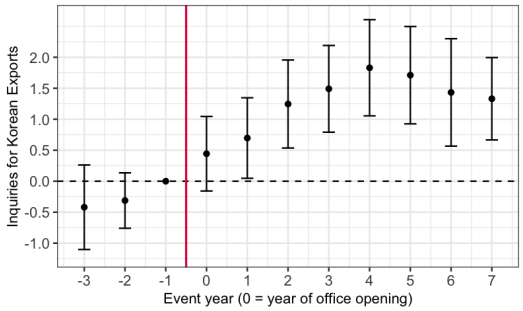

(f) Effect on reports. "Not-yet" as control.

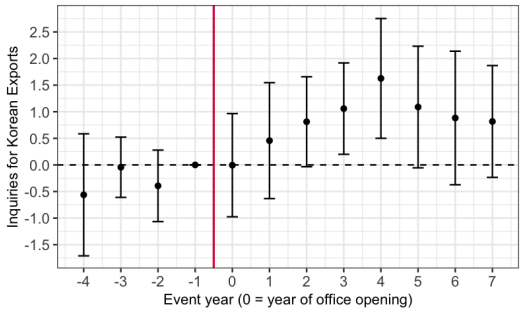

(g) Effect on inquiries. Never-treated as control. Openings from 1978.

Notes: The left panels reports coefficients θ_k from estimating equation (1). The right panels do the same following the approach by Callaway and Sant'Anna (2021). Instead of exports, we aim to explain three measures of KOTRA activity, each transformed using the inverse hyperbolic sine. (1) The number of reports about a country, (2) the number of product-specific reports - which may be more specific or informative, (3) the number of inquiries for trade with the country. Our data on reports covers the years 1965 to 2001. We thus exclude events before 1968 from the analysis in panels (a)-(d). Our data on inquiries covers the years 1974 to 1997. We thus exclude events before 1974 from the analysis in panels (e) and (f). Including events from 1975 comes at the cost of estimating only 1 pre-period in panel (e)

Figure A.8: Event study of Korean exports around switches between bureaucrats. Parallel pre-trends. Discontinuous jump in exports in line with change in bureaucrat ability.

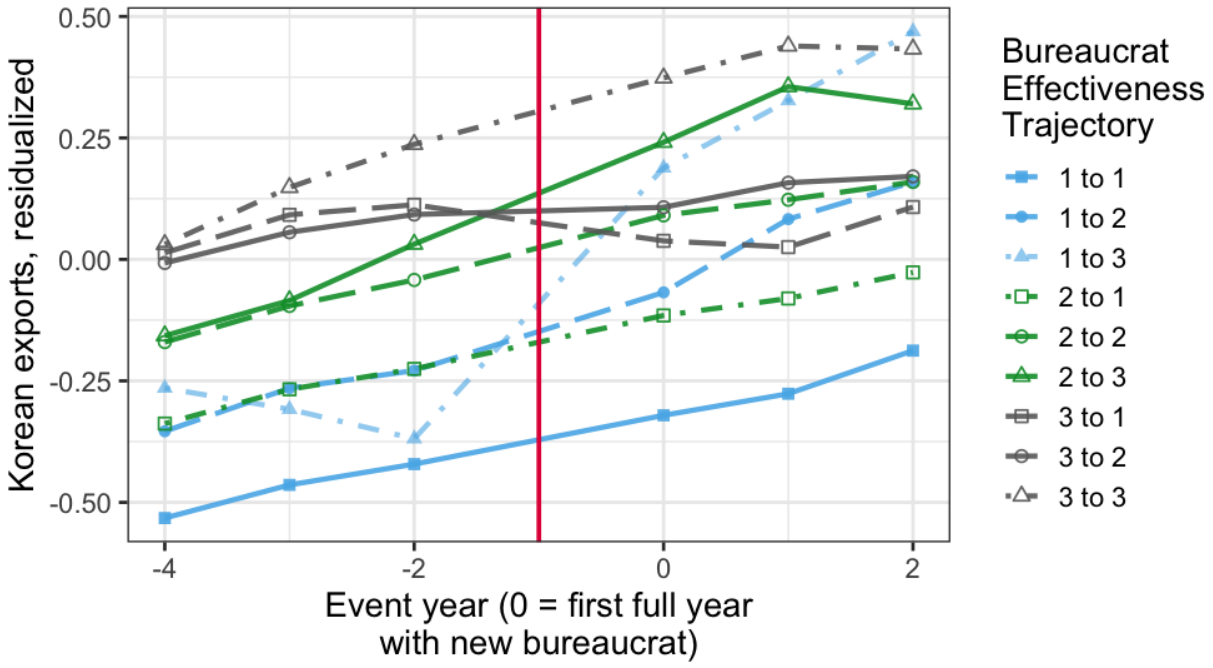

Notes: The figure shows time trends in exports around the time that the director of a country office changes. The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat heads the country office. Year -2 is the last full year that the old bureaucrat headed the office. The y axis measures average residualized exports to a destination of a product. Exports are residualized by regressing product-specific exports to a country on country and product-year fixed effects. Bureaucrats are classified into terciles according to the fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure A.9: Event study estimates: Out-of-sample bureaucrat fixed effects

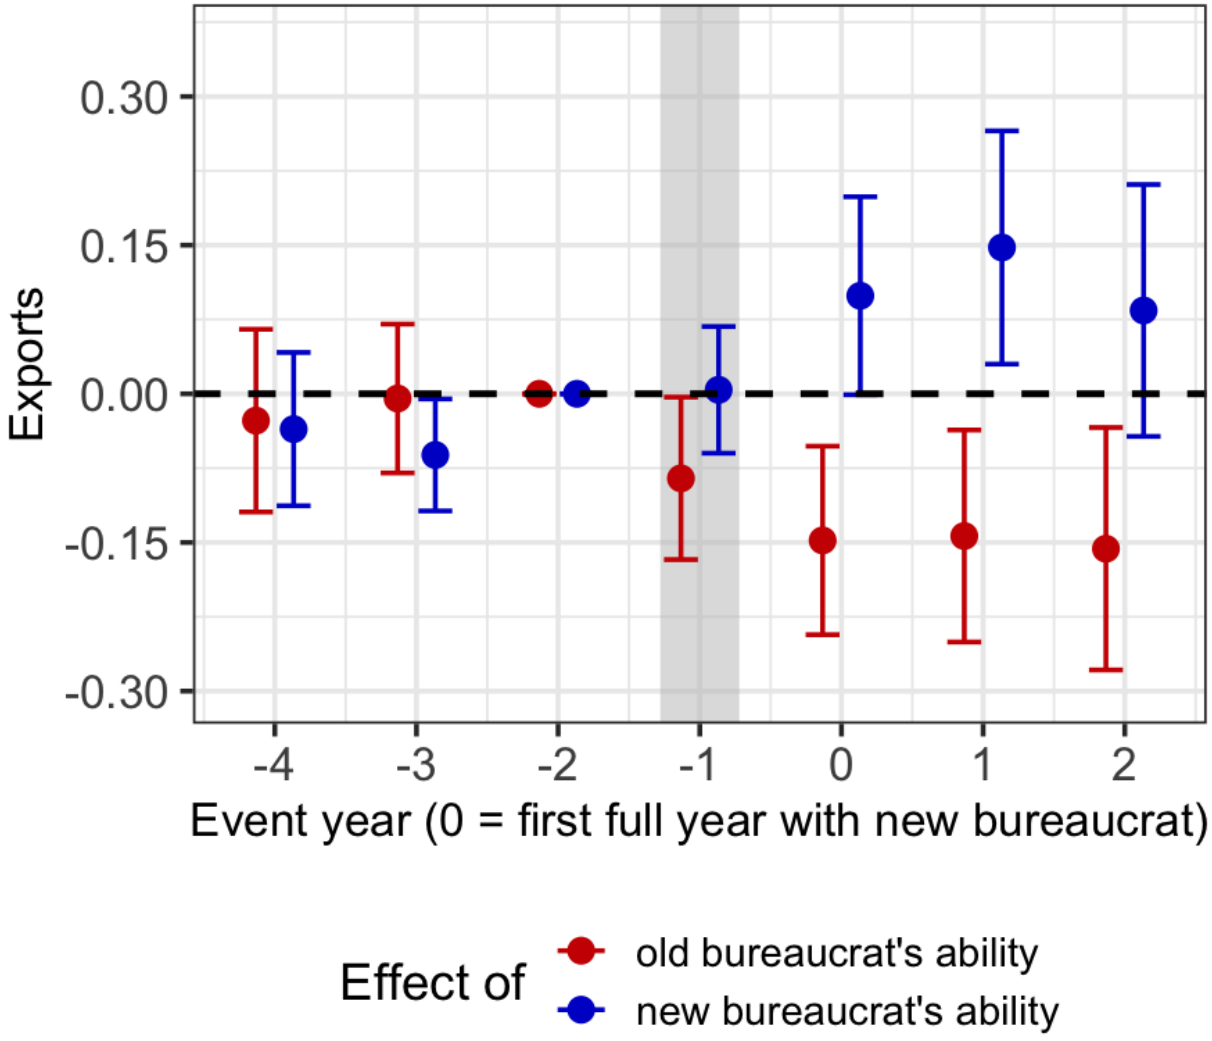

Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects, estimated out of sample, on exports around the time that the director of a country office changes. These estimates are $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (6). As out-of-sample fixed effects are not available for every bureaucrat, to maximize power, we report coefficients from two different models. First, we estimate equation (6) using *out-of-sample* estimates for the outgoing bureaucrat and *in-sample* estimates for the incoming bureaucrat. Second, we estimate equation (6) using *in-sample* estimates for the outgoing bureaucrat and *out-of-sample* estimates for the incoming bureaucrat. For each model, we only report the out-of-sample coefficients, as these are the ones of interest. For each model, the in-sample coefficients are almost symmetric to the out-of-sample ones. The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat heads the country office. Year -2 is the last full year that the old bureaucrat headed the office. The y axis measures the effect of bureaucrat effectiveness on exports. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure A.10: Event study – the extensive margin response to switches between bureaucrats

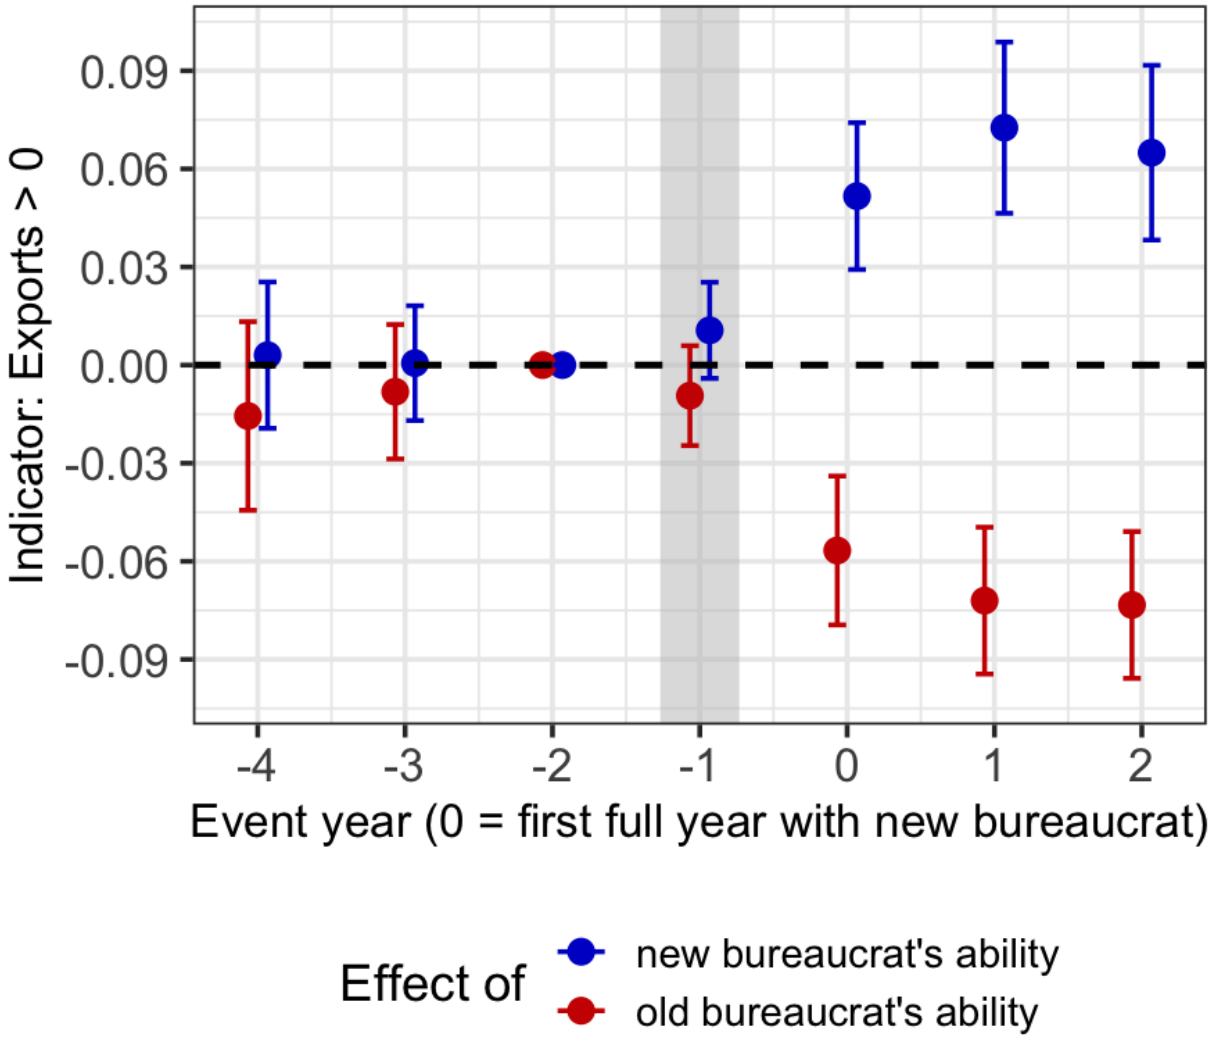

Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects on the likelihood of positive exports in a given product around the time that the director of a country office changes. These estimates are $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (6). Observations are included for a given event-horizon if Korea exports this product to any country for all years in the event horizon. The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat heads the country office. Year -2 is the last full year that the old bureaucrat headed the office. The y axis measures the effect of bureaucrat effectiveness on exports. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure A.11: Large extensive margin response to bureaucrat effects for products with any change in extensive margin during event horizon

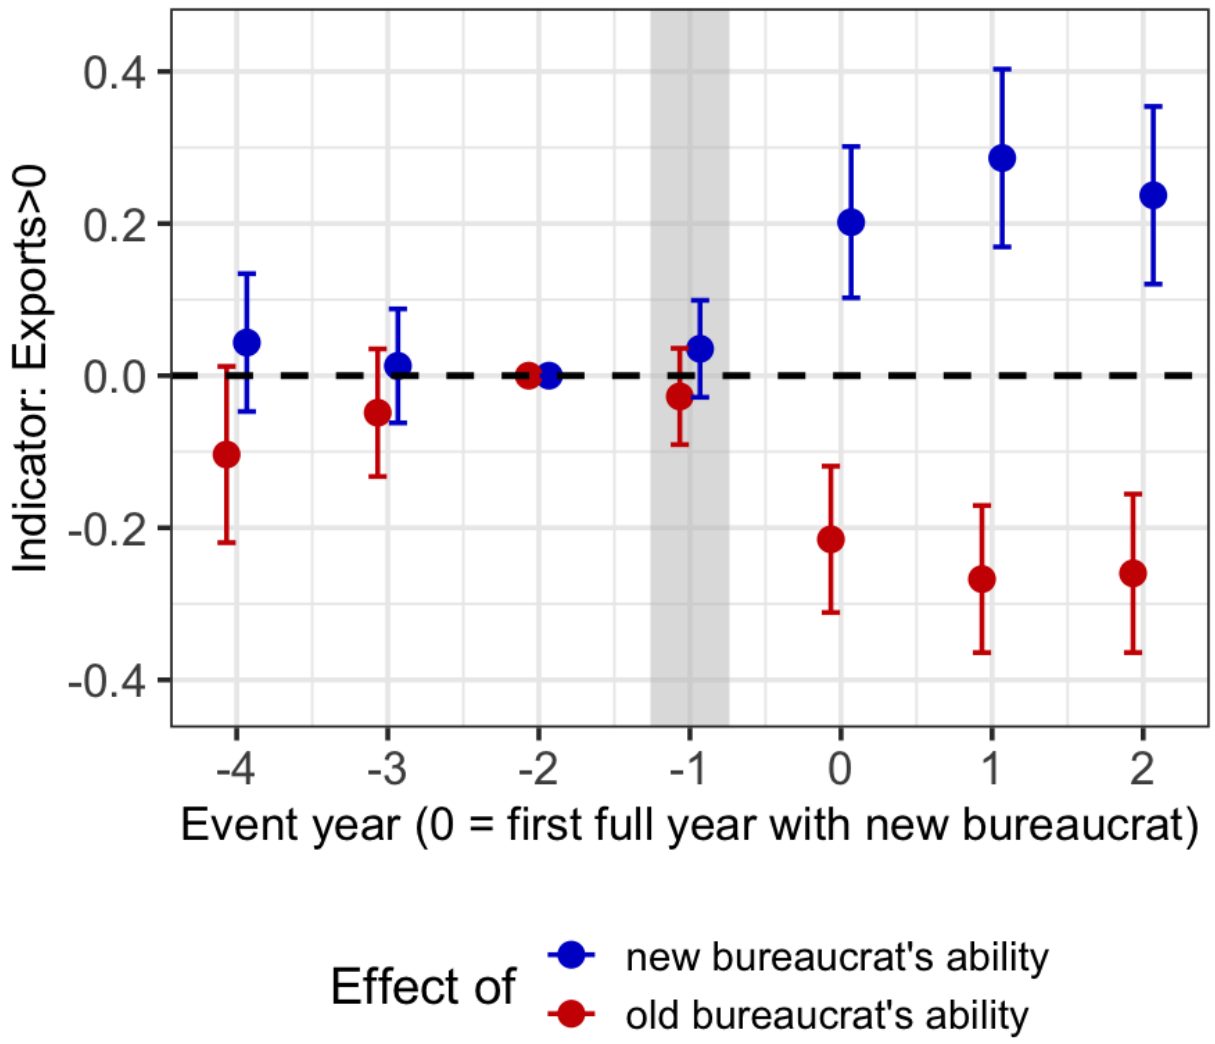

Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects on the likelihood of positive exports in a given product around the time that the director of a country office changes. These estimates are $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (6). Observations are included for a given event-horizon if Korea exports this product to this country in one year during the event horizon but not all years in the event horizon. The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat heads the country office. Year -2 is the last full year that the old bureaucrat headed the office. The y axis measures the effect of bureaucrat effectiveness on exports. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure A.12: Event study – the intensive margin response to switches between bureaucrats

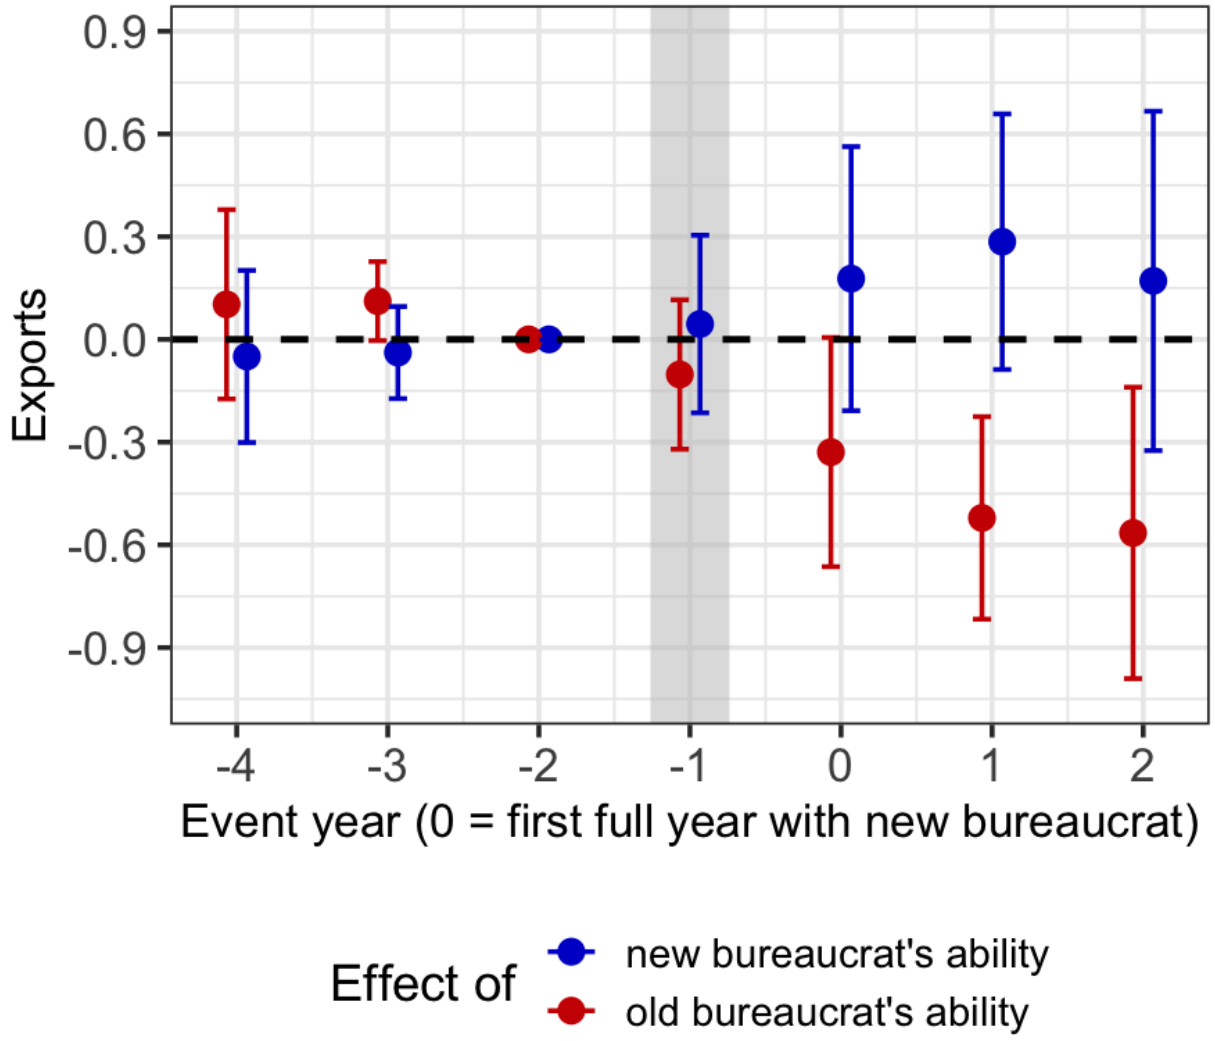

Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects on exports around the time that the director of a country office changes. This only includes the intensive margin effect as observations are included for a given event-horizon if Korea exports this product to this country in all years during the event horizon. These estimates are $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (6). The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat heads the country office. Year -2 is the last full year that the old bureaucrat headed the office. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure A.13: Event study – Exports increase upon greater switches in experience

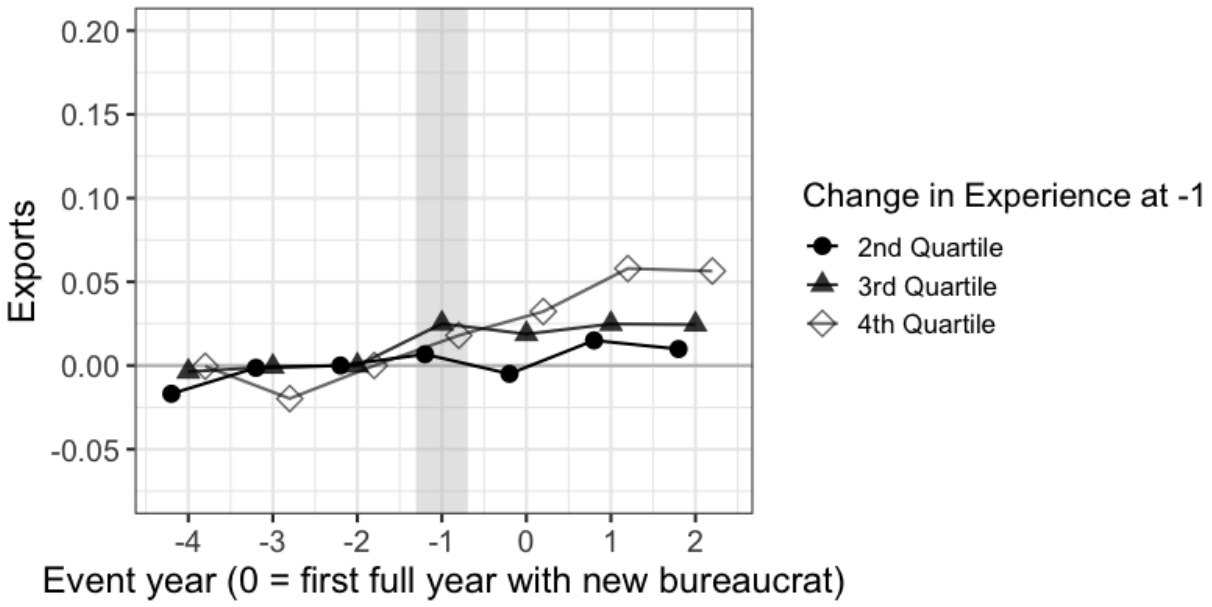

Notes:

Figure A.14: Event study – quasi-random experience$_p$ increases reaction to market conditions. Estimates of main effect become imprecise.

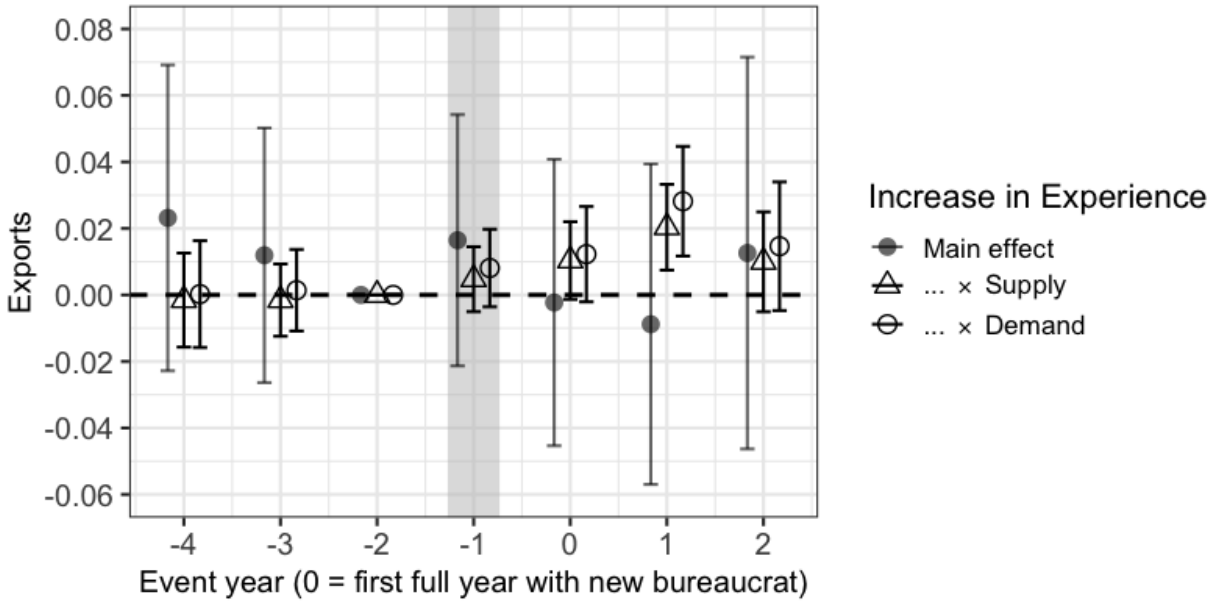

B Appendix Tables

Table B.1: Pre-determined market size determines office opening when distance is similar

	Opening	Non-Korean imports 1962	Predicted	Predicted (Omit own)
UK	1965	1	1965	1966
Italy	1966	4	1967	1967
Netherlands	1966	5	1967	1969
W Germany	1967	2	1966	1966
Switzerland	1967	8	1970	1972
France	1969	3	1966	1966
Sweden	1969	7	1969	1970
Austria	1970	12	1973	1973
Belgium	1972	6	1969	1969
Spain	1972	10	1972	1972
Denmark	1973	9	1972	1972
Norway	1973	11	1973	1973
Finland	1973	13	1973	1973
Greece	1973	15	1973	1973
Turkey	1973	16	1973	1974
Ireland	1973	14	1973	1973
Portugal	1974	17	1974	NA

Notes: The column 1st Opening displays the year in which a country's first office actually opened. The column Non-Korean imports in 1962 ranks the countries by the size of imports from countries other than Korea in 1962. The next column assigns the year of the nth 1st opening to the nth country as ranked by non-Korean imports in 1962. Italy is assigned the 4th opening year (1967). The final column does so while neglecting a country's own opening. Hence, Italy is assigned the 5th opening year (1967) - as this is the 4th when omitting the actual opening in Italy.

Table B.2: The effect of EP on exports depends on the individual bureaucrat. Bureaucrat effects do not differ between appointments.

	Exports			
	(1)	(2)	(3)	(4)
Share of Variation explained by FE				
Adj. R^2	0.345	0.442	0.460	0.464
R^2	0.355	0.451	0.469	0.473
Year-product FE	Yes	Yes	Yes	Yes
Country FE		Yes	Yes	Yes
Bureaucrat FE			Yes	Yes
Bureaucrat-Country FE				Yes
Observations	1,772,452	1,772,452	1,772,452	1,772,452
Bureaucrats	397	397	397	397
Countries	87	87	87	87

Results from estimating equation (2) reported. An observation is a product-country-year. The dependent variable is exports after residualizing by product-year and country fixed effects. A country is included for all the years that it has an office and is linked to a bureaucrat. A product is included for all the years in which Korea exported it to any country. Product refers to 4-digit SITC Rev. 2 codes. S.D. of ihs exports : 2.45, s.d. of ihs exports | tp, c: 1.83. The increase in R^2 due to bureaucrat FE is most meaningfully compared to the increase due to country FE – 0.018 compared to 0.097. These levels are lower than reported in the variance decomposition as the latter bundles all observations within an appointment while this table retains separate observations for each product, thus including variation that cannot be explained by product-invariant explanatory variables such as country FE and bureaucrat FE.

Table B.3: The extensive margin's importance to each event changes little over time. Across decades, the intensive margin becomes relevant to more products.

Year of switch	Events	No. products with extensive margin change		No. products with exports> 0 throughout	
		Mean	Median	Mean	Median
1965-1969	21	96.0	76.0	17.6	8.0
1970-1974	61	119.9	108.0	30.4	16.5
1975-1979	88	138.1	124.5	37.4	27.0
1980-1984	117	169.4	153.5	62.1	47.0
1985-1989	102	163.3	149.0	52.3	24.0
1990-1994	112	144.6	144.0	82.9	55.0
1995-1999	132	154.8	150.0	127.8	89.5

This table gives the mean and median number of products across events. It first does so for products with extensive margin changes during the event horizon, i.e. products with both positive and 0 exports to the respective country. It also reports the number of products with only positive exports throughout the event horizon, i.e. products with positive exports to the respective country in each year of the event horizon.