THE IMPLEMENTERS OF INDUSTRIAL POLICY : BUREAUCRATS AND THE KOREAN EXPORT MIRACLE*

Philipp Barteska
The University of Hong Kong

Jay Euijung Lee
Stockholm University

February 25, 2026

Abstract

How much does the effect of industrial policy depend on the capacity of the bureaucrats implementing it? We exploit the rotation schedule of managers of South Korea's export promotion offices in 87 countries between 1965 and 2000 to show that a one standard deviation increase in bureaucrat ability boosts product-level exports by 37%, while the policy increases exports by 38% on average. Together, this implies the export promotion policy has no effect when implemented by a bureaucrat one standard deviation below average. Under higher-ability bureaucrats, South Korean exports respond more strongly to destination import demand and expand more for products prone to information frictions, suggesting more effective transmission of market information. We find that performance in the first appointment predicts whether a bureaucrat sees subsequent appointments, highlighting performance-based screening of bureaucrats as a mechanism that increases the policy's effect.

JEL classification: L52, D73, O1, F13, M50

Keywords: Industrial Policy, Bureaucracy, Economic Development, Export Promotion

*Corresponding author: barteska@hku.hk. Faculty of Business and Economics, The University of Hong Kong. Pokfulam Road, Hong Kong. This paper benefited from comments by David Atkin, Jie Bai, Tim Besley, Michael Best, Mike Callen, Swati Dhingra, Alessandra Fenizia, Gordon Hanson, Jonas Hjort, Asim Khwaja, Nathan Lane, Rocco Macchiavello, Isabela Manelici, Benjamin Marx, Cristobal Otero, Elias Papaioannou, Torsten Persson, Steve Pischke, Dani Rodrik, Dan Rogger, Daniel Sturm, James Robinson, Jose Vasquez, Eric Verhoogen, Guo Xu, Noam Yuchtman, as well as discussions by Juan Felipe Riaño, Maddalena Ronchi, Yang Xun, and numerous seminar audiences, especially at the Korea Development Institute (KDI). Philipp Barteska would especially like to thank his advisors: Oriana Bandiera, Gharad Bryan, and Robin Burgess. This study was made possible with the excellent research assistance of Minjeong Kang, Sunae Choi, and Gippeum Lee. We would also like to thank the staff at the KOTRA library, particularly Hyunsun Cho. Financial support from STICERD and the Hayek Programme is gratefully acknowledged.

1 Introduction

Successful cases of national industrialization after WWII—most notably in East Asia—involved dramatic shifts in industrial composition, from production for domestic markets using traditional technologies to export-oriented production using modern technologies. Proponents of state intervention often point to the prevalence of industrial policies during these export-led growth episodes.

While the merits of industrial policy remain debated academically, governments around the world continue to devote substantial resources to such policies (Juhász et al., 2024). In practice, however, only a small number of industrial policy episodes are widely regarded as successful, while many others have fallen short. Much of the research on industrial policy effectiveness has focused on policy design, such as whether governments target the right sectors at the right scale and employ appropriate instruments. Yet formally similar policies often produce sharply divergent outcomes, suggesting that design alone cannot account for the difference.

This paper highlights a key, underexplored factor: implementation. Because industrial policies require a plethora of implementation decisions, their success may depend critically on the individuals charged with carrying them out. We study how the effects of an industrial policy depend on the idiosyncratic ability of the implementing bureaucrats, in the context of South Korea's export promotion policy during its growth miracle. While recent studies have analyzed the effects of different industrial policies, especially in East Asia (Juhász et al., 2023),[1] little is known about how sensitive these policies are to the bureaucracy implementing them.

South Korea's overseas export promotion provides an ideal setting to identify how much industrial policy depends on bureaucratic effectiveness. First, the policy was implemented decentrally by offices in 87 destination countries. Second, the bureaucrats who managed each country office rotated every three years, generating plausibly exogenous variation in implementing capacity within location, i.e., holding fixed time-invariant differences in economic conditions that directly affect the outcome. The rotation schedule also enables tests that alleviate concerns over endogenous bureaucrat appointments. Third, throughout the sample period, the key outcome targeted by each office remained exports to the respective country, allowing for a one-to-one mapping between the implementing bureaucrat and the policy outcome.

Our analysis draws on newly assembled data that allow us to link bureaucrats to both office activities and export outcomes. We compile and cross-check appointment histories for all bureaucrats who managed export promotion offices from 1965 to 2000 using multiple archival sources. We also construct data on intermediate office outputs, such as reports on country-specific market conditions,

[1] Recent papers have assessed policies' average—and sometimes distributional—effects, while considering how these effects are shaped by the policy instruments or underlying economic conditions (Aghion et al., 2015, Juhász, 2018, Barwick et al., 2025, Lane, 2025, Liu, 2019, Mitrunen, 2024, Choi and Levchenko, 2021, Choi and Shim, 2026, Juhász et al., 2024, Kim et al., 2022).

information on potential foreign buyers, and participation in trade fairs, allowing us to open the black box of export promotion and study how bureaucrats differ in their operational choices.

Our main result demonstrates that the effect of the export promotion policy in a destination country strongly depends on the bureaucrat who leads the respective country office. We use a movers design in a two-way fixed effects framework (Abowd et al., 1999) to decompose exports, net of product-year trends, into a destination country component and a bureaucrat component, representing bureaucrat ability to affect exports. We find that increasing bureaucrat ability by one standard deviation increases product-level exports by 37%.[2] This reflects both an expansion in the set of products exported and higher export volumes among products with positive exports. Moreover, the variation in exports due to bureaucrats amounts to 1/7 of that due to countries.[3] These numbers arise from a variance decomposition exercise which corrects for finite-sample bias common in two-way fixed effects frameworks (Kline et al., 2020). A placebo reshuffling of bureaucrats confirms that this bias correction allows us to differentiate bureaucrats' effects on exports from noise. An alternative bias correction method we develop finds near-identical results.[4]

The unbiasedness of the bureaucrat fixed effects is supported by numerous diagnostic tests informed by the rigid appointment schedule of bureaucrats. The identifying assumption is that bureaucrat ability is orthogonal to underlying trends in exports. For our main diagnostic check, we perform event studies of exports around switches of office managers, based on the following predictions. First, we should observe differential pre-trends by bureaucrat ability if good bureaucrats were appointed because of underlying trends. Second, the three-year rotation implies that if good bureaucrats were strategically appointed to countries with growing exports, we should find larger gains when a manager joins an office than the drop when they leave, as there is much less discretion over the timing of departures. Empirically, however, we find parallel pre-trends and symmetric effects of gaining and losing a bureaucrat. In addition to the placebo test described above, we confirm that bureaucrat effects are not driven by noise as they remain stable across different subsamples and retain strong predictive power out of sample.

To benchmark the effect on exports due to bureaucrats, we separately estimate the average effect of the policy, using the staggered rollout of the offices in a dynamic difference-in-differences setup.[5] We find that the policy had a substantial effect on average, increasing exports by 38% ten years

[2]For comparison, from 1965 to 2000, total South Korean exports grew 3000% relative to the U.S. (IMF: Direction of Trade Statistics, 2023).

[3]Recent research on managers in the public sector and firms, using a similar two-way fixed effects framework, has found that managers explain 1/3 - 2/3 as much variation as the organizations they manage (Fenizia, 2022; Metcalfe et al., 2023; Otero and Muñoz, 2025).

[4]This alternative bias correction builds on Best et al. (2023).

[5]The main specification estimates the effect of opening an overseas office relative to a never-treated control group. This choice of control group avoids biases arising from the combination of dynamic treatment effects and staggered rollout in two-way fixed effects regressions. We also test robustness using a not-yet-treated control group (Callaway and Sant'Anna, 2021) and assess sensitivity to parallel trends violations (Rambachan and Roth, 2023).

after an office opening. Thus, the policy would have no effect if implemented by bureaucrats one standard deviation below average.

The differences in bureaucrat ability are explained by one key mechanism: more effective transmission of timely market information. High-ability bureaucrats expand exports more strongly where information frictions are severe—such as for differentiated products and products with volatile demand—and where the value of information is high, including products facing rising import demand in the destination or heightened policy support during the heavy and chemical industry (HCI) drive. Consistent with this interpretation, bureaucrat ability correlates with office activities that facilitate buyer–seller matching, such as listing more import inquiries and participating in more trade fairs.

Appointment patterns suggest that the export promotion organization had the capacity to screen bureaucrats by ability, contributing to the large average effect of the policy. First, bureaucrats are much less likely to be reappointed if exports underperform during their first appointment as manager of a country office. This suggests that the organization learns about a bureaucrat's type during the first appointment. Second, first appointments as country office manager are to less important trade partners, while later appointments are to more important destinations.[6] Hence, the most important offices are rarely run by underperforming bureaucrats.[7] These descriptive patterns align well with archival records on the organization's operations that endorse performance-based pay levels, job assignments, and promotions.[8] A counterfactual scenario with no screening after first appointments reveals that exports would be 6.5% lower than in the baseline.

This paper contributes to a growing literature that studies the effects of industrial policies on firm performance, sectoral outcomes, and aggregate growth (Juhász et al., 2023; Juhász, 2018; Liu, 2019; Lane, 2025; Choi and Levchenko, 2021; Choi and Shim, 2026; Kim et al., 2022). Recent contributions further analyze and compare the performance of different policy instruments, such as production subsidies, investment subsidies, entry subsidies, and consolidation policies (Barwick et al., 2025). In contrast, this paper focuses on the individuals implementing the policy as a major source of variation in industrial policy outcomes. We therefore provide direct evidence on the often-hypothesized but under-researched link between state capacity—at the level of individual bureaucrats—and the effectiveness of industrial policy.

Second, we contribute to the literature on bureaucratic capacity and economic development

[6]Note that as long as the ranking of country importance is time-invariant, this appointment process satisfies the identifying assumption for bureaucrat fixed effects. We classify countries as less important if they have lower fixed effects or later office opening years.

[7]These patterns are consistent with the organization using offices in less important countries to experiment with inexperienced bureaucrats—a strategy more broadly applicable to teachers and managers in other settings, where observable characteristics insufficiently predict the substantial variation in performance.

[8]For example, in a 1982 publication outlining its strategy for the 1980s, the organization highlighted that performance-based rewards and punishments need to be strengthened.

(Besley et al., 2022) by combining aspects of the following two approaches. The first approach studies bureaucrats with a clearly defined output for which they are immediately responsible. These are often front-line bureaucrats rather than bureaucrats who make policy.[9] The second approach studies CEO-like bureaucrats with geographic responsibility who can then plausibly be linked to broad measures of economic activity in their region, but such bureaucrats often manage multiple objectives.[10] The bureaucrats we study uniquely embody advantages of both approaches; they have a focused, well-defined output measure that is also of direct macroeconomic importance—aggregate bilateral exports. Our finding that the effect of export promotion is eliminated when managed by a bureaucrat one standard deviation below the mean demonstrates that bureaucratic staffing decisions are as important as policy design itself for states seeking to promote economic development.

Third, we contribute to research on the role of managers in determining economic outcomes using two-way fixed effects methods, as in recent work on public-sector managers (Fenizia, 2022; Otero and Muñoz, 2025) and private firms (Bertrand and Schoar, 2003; Metcalfe et al., 2023). We differ in two ways. First, we study managers implementing a major state economic policy. Second, we explore the broader organizational learning process—how organizations recognize and manage uncertainty over managerial ability over time.

Fourth, we contribute to the literature on trade and economic development, including work on the impact of exporting on firm growth (Atkin et al., 2017; Alfaro-Ureña et al., 2022) and market-size constraints in development (Goldberg and Reed, 2023), by providing evidence that individual bureaucrats affect firms' access to foreign markets. We also build on research on export promotion, which has focused on average treatment effects (Munch and Schaur, 2018; Hayakawa et al., 2014; Bagir, 2020; Volpe Martincus and Carballo, 2008) and on which sectors and firm types benefit most (Volpe Martincus and Carballo, 2010, 2012). This includes Kim and Kim (2024), who show that public information on potential foreign buyers, provided by the same export promotion agency that we study, significantly boosted exports of small and medium-sized firms by alleviating search frictions. We contribute to this literature by newly unpacking the role of individuals in shaping the effect of export promotion, in a particularly relevant setting: South Korea, the poster child of export-led growth.

The rest of the paper proceeds as follows. Section 2 describes the institutional background. Section 3 introduces the data. Section 4 studies how much individual bureaucrats affect exports

[9]Empirical research often focuses on front-line bureaucrats with well-defined tasks, as modern causal inference tools are particularly well-suited to studying them. Examples include agricultural extension workers (Dal Bó et al., 2021), revenue collectors (Khan et al., 2016), health care providers (Ashraf et al., 2014), teachers (Leaver et al., 2021), procurement officers (Best et al., 2023), and judges (Dahis et al., 2023).

[10]Examples include Chinese provincial governors linked to GDP growth (Jia, 2024), British governors linked to colony-level revenues (Xu, 2018), and Chinese city leaders linked to local industrial comparative advantage (Lin et al., 2024). Gulzar and Pasquale (2017) study Indian bureaucrats responsible for district-level development at large, but focus on one task that is clearly measurable and interpretable: the implementation of the NREGA workfare program.

and the success of the policy. It also investigates how better bureaucrats increase exports and finds that the export promotion organization had the capacity to screen bureaucrats by ability. Section 5 discusses the policy's average effect estimated using office openings. Section 6 concludes.

2 Institutional Background

During our period of study, South Korea's real GDP per capita increased from $1,275 (1962)—below most of Sub-Saharan Africa—to $24,834 (2001).[11] Panel (a) of Figure 1 highlights that export growth was particularly remarkable. Korean exports per capita in the 1950s were below 2% of the U.S. level. From 1960, however, exports increased rapidly and reached parity with the U.S. by the 1990s. During this period, comparable developing countries saw little to modest growth.

It is often argued that this growth is partly due to industrial policies conducted by a well-functioning state.[12] On the other hand, the Korean state was described as aid-dependent and corrupt until the 1960s (Kim and Vogel, 2011).[13] These contradicting narratives, summarized in Section 2.1, motivate us to study how much a particular industrial policy's effect was due to the state's implementation capacity. We focus on overseas export promotion, identified in a 1976 survey of Korean manufacturers as the policy area where government intervention improved most under President Park Chung-hee (1961-1979) relative to Syngman Rhee (1948-1960) (Jones and Il, 1980).

2.1 Bureaucratic Capacity and Korean Economic Development

Qualitative political economy attributes East Asia's rapid economic growth to successful industrial policy (Johnson, 1982; Amsden, 1989; Wade, 1990; Evans, 1995; Woo-Cumings, 1999). These accounts either explicitly or implicitly argue that South Korea's high state and bureaucratic capacity was essential to the economic and export growth it experienced from the mid-1960s. Amsden (1989), perhaps the most influential account of South Korean industrial policy, emphasizes that "the power of the state to discipline big business was greater in Korea—and Japan and Taiwan as well—than in other late-industrializing countries" (p. vi). Jones and Il (1980) further highlight the importance of implementation and adaptation in South Korea's industrial policies, "only possible to governments possessing a well-trained bureaucracy" (p. xxxi, foreword by Edward S. Mason).

Conversely, other historical research describes South Korea's *lack* of bureaucratic capacity in the 1950s and 1960s: "Under Syngman Rhee the bureaucracy was generally both ineffective and disorganised, characterised by widespread corruption and patronage. Not only were policy

[11] Measured in 2017 USD. This corresponds to an increase from 1/16 of U.S. real GDP per capita in 1961 to 1/2 in 2001. Data from Penn World Tables.

[12] Wade (1990) and Cheng et al. (1998) as cited by Besley et al. (2022); Amsden (1989); Juhász et al. (2023). See also the popular book by Studwell (2013). Economists remain broadly skeptical that industrial policies can be welfare-improving, including in the case of South Korea's recent history.

[13] Korea's level of state capacity may be illustrated by its ministerial instability. Under President Rhee (1948-1960), agriculture ministers served an average of just 9 months and commerce ministers just 13 months (Haggard et al., 1991).

instruments used for political purposes, but the staffing of the bureaucracy itself was an important form of patronage [Suh (sic), 1967]." (Cheng et al., 1998; Bark, 1967). The limited state capacity is further underlined by Kim and Baik (2011): "South Korea lacked the expertise necessary for modern government and frequently relied on American advisors to strengthen state capabilities." Since over 90% of the government budget in 1961 was funded by U.S. aid, U.S. advisors were "overseeing and shaping South Korea's major social and economic policies for all practical purposes."

These vastly different perspectives may, in some part, be due to hindsight bias; authors around 1990 may have attributed high bureaucratic capacity to 1960s South Korea partly because, at the time of writing, they were observing a newly rich and democratic country with a well-functioning bureaucracy. Overall, assessing the importance of South Korean bureaucratic capacity on subsequent economic growth is complicated by the difficulty of establishing the link causally. We provide a step forward by causally identifying the importance of individual bureaucrats in the effect of one important policy pursued by the South Korean government.

2.2 Export Promotion by KOTRA: Tasks and Outputs Produced

We study the overseas offices of South 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[d] sales promotion and research, a campaign of public relations and advertising, [and] information service to exporters and importers" (Udell, 1965). Panel (b) of Figure 1 displays the rollout of KOTRA's overseas offices. By 1970, offices had opened in 32 countries, rising to 75 by 1981 before the pace slowed. Panel (c) maps the countries that received offices between 1962 and 2001, illustrating the broad geographic coverage.

Each overseas office targeted their efforts at boosting exports to its host country. Concretely, they contributed to the three main functions of KOTRA that were maintained consistently throughout the period of study. First, KOTRA's "Market Research" division investigated factors related to (1) South Korea's capability to supply a product for exports and (2) the import demand in a particular destination 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 Expansion" division by helping domestic producers and retailers find new trade partners in new and existing markets. This was achieved through direct matchmaking and by collecting and publicizing import inquiries from foreign firms, as well as export inquiries from domestic companies. These inquiries were published in KOTRA's daily newspaper. Business transactions were then mediated between the inquirers and respondents.

Third, the overseas offices aided the "Trade Fair" division with the organization of a South Korean pavilion at international trade fairs. KOTRA's reports attribute substantial export volumes

to these: During a 5-day trade fair in London in 1978, the reported sales by the participating firms amounted to 3% of Korea's exports to the UK that year. The overseas offices managed logistics, such as shipping exhibited products, and coordinated the recruitment, selection, and briefing of exporters. Companies were chosen based on their ability to supply new styles matching market demand. Before each fair, KOTRA offices promoted exporters and their products to potential buyers through ads, letters, calls, and outreach to trade associations.

Compared to other bureaucracies, KOTRA's overseas offices enjoyed considerable discretion in promoting exports. For this reason, we focus on KOTRA's ultimate outcome of interest: exports. It would be 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 occur through fairs, phone calls, or some other channel. Instead, achieving this goal relied on the bureaucrats' knowledge, often tacit and local, and required substantial improvisation. In this regard, KOTRA office managers resembled CEO-like bureaucrats overseeing entire geographic regions, similar to the Indian Administrative Service (Bertrand et al., 2020) or the British colonial administration (Lugard, 1926; Xu, 2018). However, unlike these bureaucrats responsible for a range of policies and outcomes, KOTRA bureaucrats implemented a single policy with a clear target, largely summarized as exports during their appointments. The primary performance measure, as assessed by the head office, was whether export targets were met[14]—supported by Figure 6, which shows that bureaucrats saw fewer future appointments if exports underperformed during their first appointments.

2.3 Assignment of Bureaucrats to KOTRA's Overseas Offices

From 1965 to 2000, KOTRA operated 138 offices in 87 countries. The analysis focuses on the main country offices as data on the outcome—exports—is available at the country level.[15]

The assignment of bureaucrats to offices is determined by the head office's HR team. Our analysis, along with interviews with current and former KOTRA officers, highlights several factors influencing appointments. First, career progression matters. A bureaucrat's first appointment as office head is mostly to unimportant countries, whereas third appointments are almost always to important countries.[16] A second important factor is language skills; a Spanish speaker is more likely to be sent to a Hispanophone country. Third, a bureaucrat who was previously posted to an undesirable location, such as a small, low-income country far from Korea, might be compensated

[14]Archival records and interviews with former bureaucrats offer insight into broad characteristics of the incentive structure. Salary consisted of base pay, determined by seniority and position, as well as performance pay. A 2002 KOTRA publication indicates that an annual performance bonus could reach up to 5 times the monthly base pay. Additionally, a separate living allowance reflected local cost-of-living differences. Performance was primarily assessed at the office level, reflecting the fulfillment of office- and year-specific task allotments and export targets.

[15]Most countries have exactly one office. In most other cases, there is one central office with a couple of subordinate ones. The most notable cases with multiple offices are the largest trade partners (U.S., Japan) and geographically large countries (U.S., Canada, Australia). After our sample period, China joins both of these groups.

[16]Before becoming office head, bureaucrats typically serve in multiple subordinate roles at overseas offices.

by getting posted to a desirable location next. Lastly, connections with KOTRA executives may matter for assignments to desirable locations. It should be noted that factors such as a country's importance and income relative to other countries, as well as its language, are largely time-invariant.

Organizational rules provide substantial rigor regarding the timing of appointments. The regularity of appointments is underscored by the fact that both the modal and median duration is 36 months. Panel (a) of Appendix Figure A.1 plots the distribution of appointment durations. Between appointments, bureaucrats return to South Korea, typically working at KOTRA's headquarters in Seoul and sometimes at regional offices. The timing of their reappointment is also largely pre-determined: The median duration for the gap between appointments is 29 months, and the modal gap is 30 months. Panel (b) of Appendix Figure A.1 plots the distribution of gaps between appointments.

The rotation schedule provides a direct reason for changes in the bureaucrat managing a country office. Its rigidity limits strategic appointments, as only a limited pool of bureaucrats is available when a country is due for a new manager. Moreover, there is little discretion over when a bureaucrat leaves an office. Section 4.1 discusses how these patterns allow us to test our identifying assumption that bureaucrat appointments are orthogonal to underlying export trends in the country of appointment.

2.4 KOTRA and Korea's Largest-Scale Industrial Policy

One reason to study export promotion is the narrative of South Korea's development as export-driven and the prominent role of exports in its industrial policy. Export promotion was closely linked to Korea's largest industrial policy, the Heavy and Chemical Industries (HCI) drive, which began in early 1973 and ended in October 1979.

Appendix Figure A.2 displays how the targeting of KOTRA's activity changed over time. Before the HCI drive, only 15-25% of product-specific reports discussed HCI products.[17] During the HCI drive, this share increased rapidly, reaching nearly half of all reports in the late 1970s. After the HCI drive, the share of reports targeting these sectors remained roughly constant. These patterns show that export promotion formed a part of other well-studied industrial policies.

At the same time, it is worth noting that KOTRA's overseas offices each focused on exports *to a particular destination country*. The cross-destination variation we exploit is therefore plausibly orthogonal to other contemporaneous industrial policies, such as the HCI drive, that targeted specific sectors or regions within South Korea. To account for these policies and for differential sectoral growth more broadly, our regressions include product-year fixed effects.

[17] How we collected and prepared the data on KOTRA's reports is briefly described in Section 3.3 and detailed in Appendix C.3. To classify HCI products, we follow Lane (2025) and include those listed in the enforcement decrees and national sectoral acts underlying the HCI drive. HCI's six broadly defined target sectors included steel, nonferrous metals, shipbuilding, machinery, electronics, and petrochemicals.

3 Data

Our main analyses use data on bureaucrat appointments to explain Korean exports. This is complemented with data on the three main dimensions of KOTRA's overseas activities: market reports, business inquiries, and trade fair participation. Appendix C provides full details on the digitization and coding procedures for each of these data sources.

3.1 Bureaucrat Appointments

We construct a panel of KOTRA bureaucrat appointments to overseas offices from 1962 to 2000 by combining appointment announcements from newspaper archives with a variety of KOTRA internal publications. We identify 398 unique bureaucrats appointed to head country offices, resulting in 705 unique bureaucrat–country pairs. Bureaucrats are identified by their full names, which we harmonize across sources that render names in Chinese characters, Korean characters, and romanizations.

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.

3.3 Intermediate Office Outputs

We complement the export data with measures of concrete bureaucrat activity digitized from KOTRA documents, covering three main functions of KOTRA's overseas operations.

Market Reports. KOTRA's daily publication *Overseas Market News* contain reports on foreign market conditions, tariffs and other trade regulations, as well as guides on how to export. We extract the Table of Contents from 9,602 issues spanning 1965 to 2001. Using dictionary-based country matching and manual coding of high-frequency tokens, we classify each entry by destination country and product category. We link 80,000 entries to at least one specific country. 45,000 of these reports are also linked to a two-digit product.

Business Inquiries. *Overseas Market News* also contains tables listing business inquiries. They are mainly requests from foreign firms seeking to import particular products from Korean firms. We digitize 169,280 inquiries from 1974 to 1997 using both OCR- and LLM-based extraction procedures on scanned pages of the publication. Each inquiry is linked to a destination country and a set of two-digit products.

Trade Fairs. KOTRA organizes Korean pavilions at international trade fairs. We digitize firm attendance records from KOTRA's Annual Reports on International Trade Fairs, covering 1024 fairs between 1969 and 1996. The data record more than 13,000 instances of Korean firm participation, including sales outcomes for over 5,700 firm–fair pairs.

4 Do Bureaucrats Drive Korean Exports?

In this section, we find that the effect of a South Korean overseas export promotion office differs substantially by the bureaucrat managing it: Increasing bureaucrat ability by one standard deviation increases exports to the respective destination country by 37%. This effect is comparable to that of opening an office for the first time, estimated in Section 5. Together, this implies that the policy would have had no effect if all bureaucrats' abilities had been one standard deviation lower.

4.1 Identifying Bureaucrat Fixed Effects

We adapt the two-way fixed effects framework (Abowd et al., 1999) to study how much bureaucrats matter in explaining South Korean exports. We proceed in two steps. We first obtain unbiased estimates of bureaucrat fixed effects and check that identifying assumptions are met. Next, in Section 4.2, we use the estimated fixed effects to quantify the variation in exports explained by bureaucrat abilities, while correcting for variation in sampling error to isolate the true, signal variance.

$$y_{cpt} = \lambda_{pt} + \gamma_c + \theta_{b(c,t)} + \epsilon_{cpt} \tag{1}$$

We model the relationship between the inverse hyperbolic sine (IHS)[18] of South Korean exports—henceforth "exports"—to country c of product p in year t, and the bureaucrat assigned to that country-year, denoted by $b(c,t)$. Exports are modeled as the sum of a product-year component (λ_{pt}), a bureaucrat component ($\theta_{b(c,t)}$),[19] a country component (γ_c), and an error term (ϵ_{cpt}). As in other parts of the paper, exports are defined at the product level, in line with KOTRA's export targets which are typically product-specific. This specification also avoids a few dominant export products driving the results for a country-year.

Equation (1) identifies the bureaucrat and country fixed effects only within the largest connected set of bureaucrats and countries.[20] Additionally, it requires that bureaucrat mobility is as good as random, conditional on product-year and country fixed effects. In other words, bureaucrat assignments need to be orthogonal to ϵ_{cpt}, which we refer to as *underlying trends in exports*. At the same time, this orthogonality condition allows for assignments on the basis of the permanent

[18] For intensive-margin changes, the IHS behaves similarly to the natural logarithm but, unlike the logarithm, is defined at zero. Because export promotion also affects the extensive margin, we use the IHS as our benchmark while exploring robustness to alternative transformations. We show in Section 4.4 that bureaucrat fixed effects based on the IHS are predictive of export changes along both margins. Furthermore, Appendix D.2 shows that the fixed effect estimates themselves are robust to alternative transformations: (1) just using the extensive margin, and (2) explicitly weighting the value of the extensive margin against the intensive margin, a solution to the unit-dependent nature of IHS proposed by Chen and Roth (2023).

[19] To account for the fact that it takes time for a new manager to influence exports, we code each country-year as being managed 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.

[20] Appendix E illustrates the concept of a connected set and a leave-one-out connected set in our setting.

component of country effects γ_c or the permanent component of bureaucrat ability $\theta_{b(c,t)}$. That is, sorting of better bureaucrats to destinations with greater time-invariant South Korean exports, e.g., to larger or richer countries, does not violate this identifying assumption.

Sample Descriptives: Bureaucrats and Countries Form One Large Connected Set

Table 1 describes the structure of the sample. The full sample contains 705 appointments of 398 bureaucrats to head 87 country offices. The largest connected set restriction excludes only one appointment and one country.[21] The largest connected set thus contains 704 appointments of 397 bureaucrats, of whom 184 had appointments to multiple offices.

The bureaucrat–country graph is sufficiently interconnected that 75 countries and 93% of appointments form a leave-one-out connected set, i.e., the set remains connected if any single appointment is removed. The reason is that most country offices remain open for decades and are headed by many different bureaucrats. Column 3 indicates that in the leave-one-out connected set, 72 offices are managed by four or more distinct bureaucrats, 61 offices by six or more, and 49 offices even by eight or more distinct bureaucrats.

Our preferred estimation uses only the appointments in the largest leave-one-out connected set. This set allows us to apply Kline et al. (2020)'s correction method for limited mobility bias in our variance decomposition, helping to distinguish bureaucrat effects from noise, as confirmed by our placebo exercise of reshuffling bureaucrats.[22]

Identifying Assumption: Appointments are Orthogonal to Export Trends

In this section, we discuss how factors shaping bureaucrat appointments relate to the identifying assumption that appointments are orthogonal to export trends.

The central factor behind the movement of bureaucrats is their three-year rotation schedule. As highlighted in Panel (a) of Appendix Figure A.1, a country typically receives a new bureaucrat every three years. This has two important implications for investigating the identifying assumption.

First, suppose KOTRA appoints good bureaucrats to countries with increasing export trends, a violation of the identifying assumption. In that case, we would observe growing exports prior to the appointments of good bureaucrats. Even if KOTRA appoints good bureaucrats to countries where export trends are *due* to increase, the rigid rotation schedule would impede them from perfectly timing all such appointments to the years when exports begin to rise. Then, the increasing trend would have materialized before some appointments, once again leading to differential pre-trends by bureaucrat ability. Section 4.4 tests this prediction in an event-study setup estimating the effect of bureaucrat switches on exports. We find parallel pre-trends, alleviating concerns that KOTRA

[21] Only Zimbabwe lies outside the largest connected set. The data contains only one appointment to Zimbabwe because its office opened near the end of our sample period.

[22] Kline et al. (2020) also show that restricting attention to the leave-one-out connected set substantially reduces limited mobility bias directly.

strategically appoints good bureaucrats to countries with increasing export trends.

Second, while KOTRA has some discretion in new appointments, the rigid rotation schedule dictates the year when a country *loses* any given bureaucrat. Suppose good bureaucrats are strategically appointed to countries with increasing export trends. The greater discretion at the start of an appointment, compared to the end, would imply that exports rise more when gaining a high-ability bureaucrat than they decrease when losing one of the same ability. Importantly, this prediction only requires that discretion is greater at the start of the appointment, which holds even if there are individual exceptions. Section 4.4 tests this prediction in the event-study setup and finds that the effects of gaining and losing a bureaucrat are very symmetric, in addition to confirming the lack of differential pre-trends. This further mitigates concerns about strategic appointments biasing our estimates. The symmetric result also rules out other related hypotheses that would violate the identifying assumption. For example, the assumption would be violated if the South Korean government synchronized new investments into exporting to particular countries with new bureaucrat appointments. However, the symmetric effects we find contradict this hypothesis, unless the investments are (somewhat implausibly) reversed at the end of the bureaucrat's term, despite limited discretion over when a term ends.

In addition to these testable implications, the rotation schedule also limits the pool of bureaucrats available when a country office needs a new manager, reducing opportunities for strategic appointments. From qualitative interviews with KOTRA employees and our analysis, appointments are subject to several constraints: (1) Bureaucrats are more likely to be appointed to a country where they speak the local language. (2) Bureaucrats prefer assignments to high-income, English-speaking countries. (3) Because these preferences are widely shared, KOTRA's Human Resources manages discontent by rotating bureaucrats between low- and high-desirability locations. (4) Bureaucrat tenure influences eligibility to lead more important country offices, as discussed further in Section 4.6. In most cases, a country's language, relative income, importance as Korea's trade partner, and overall desirability change little over time. Therefore, while appointments may be correlated with country fixed effects,[23] the scope for appointing bureaucrats based on (anticipated) export trends is very limited.

Despite these constraints, one may still wonder why KOTRA does not strategically appoint good bureaucrats to countries with growing import demand. A major reason may be that KOTRA prioritized time-invariant country characteristics over trends when gauging export potential, much like it did for office openings.[24]

[23] Recall that a correlation between bureaucrat ability and the time-invariant country effects would not violate the identifying assumption.

[24] The rollout of export promotion offices mainly followed pre-determined gravity variables, as reported in Panel (d) of Figure 7 and discussed in Section 5.3.

4.2 The Share of Variation in Exports Explained by Bureaucrats

In this section, we decompose the variance in South Korean exports to estimate the share attributable to differences between bureaucrats. To focus on the bureaucrat and country components, we first residualize exports for product-year fixed effects:

$$(\text{exports}|pt)_{cpt} \equiv \text{exports}_{cpt} - \hat{\lambda}_{pt} = \theta_{b(c,t)} + \gamma_c + \epsilon_{cpt} \tag{2}$$

where $\hat{\lambda}_{pt}$ is estimated from equation (1).[25] The remaining variation in exports is decomposed into:

$$\mathsf{Var}[(\text{exports}|pt)_{cpt}] = \mathsf{Var}(\theta_{b(c,t)}) + \mathsf{Var}(\gamma_c) + 2\mathsf{Cov}(\theta_{b(c,t)}, \gamma_c) + \mathsf{Var}(\epsilon_{cpt}) \tag{3}$$

As the variation in residualized exports *within* appointment spells—i.e., across years or products for a bureaucrat in a particular appointment—is uninformative for estimating bureaucrat or country fixed effects, we use data collapsed into spell-level averages of residualized exports.[26]

Our primary object of interest is the variation explained by the bureaucrats: $\mathsf{Var}(\theta_{b(c,t)})$. The challenge in estimating $\mathsf{Var}(\theta_{b(c,t)})$ is that a naive plug-in estimator, $\mathsf{Var}(\hat{\theta}_{b(c,t)})$, would be biased upwards due to the sampling error in $\hat{\theta}_{b(c,t)}$, which is aggravated when bureaucrats move countries too infrequently. Such *limited mobility bias* in our setting is illustrated in Appendix Figure A.3.

Our preferred variance decomposition approach follows the bias correction method of Kline et al. (2020). Kline et al. (2020) estimate and correct for the degree of bias in $\mathsf{Var}(\hat{\theta}_{b(c,t)})$ under unrestricted heteroskedasticity, building on Andrews et al. (2008)'s correction method that required homoskedasticity. We use the computational algorithm of Bonhomme et al. (2023) for implementation. Although unreported, the Andrews et al. (2008) correction method delivers quantitatively similar results.[27]

One downside of the analysis based on Kline et al. (2020) is that it does not allow us to make statements about some other moments of the distribution of bureaucrat abilities, e.g., percentiles. An

[25] The calculation of $(\text{exports}|pt)_{cpt}$ follows Chetty et al. (2014b). To remove the effect of pt without biasing the bureaucrat or country effects, $\hat{\lambda}_{pt}$ needs to be estimated using only within-bureaucrat and within-country variation. $\hat{\lambda}_{pt}$ captures macroeconomic shocks, but also long-run changes in South Korea's industrial structure. For instance, $\hat{\lambda}_{\text{cars},1965}$ is very small compared to $\hat{\lambda}_{\text{cars},1995}$. Appendix Table B.1 demonstrates the importance of product-year trends, which explain 35.5% of the total variation in exports.

[26] The bureaucrat and country fixed effects estimated from the collapsed data equal those estimated from the raw, uncollapsed data. For the variance decomposition, we treat the variance of the spell-level averages as total variation, but the variance of the raw (i.e., country×product×year-level) residualized exports is also reported in Table 2 for reference.

[27] The Kline et al. (2020) correction method can only be performed on the leave-one-out connected set which covers 75 countries and 380 bureaucrats, because the bias is an aggregation of observation i-specific bias estimated using a leave-i-out estimator. In contrast, the Andrews et al. (2008) correction method can also be performed on the largest connected set covering 86 countries and 397 bureaucrats. The three approaches—the Kline et al. (2020) method on the leave-one-out connected set, the Andrews et al. (2008) on the leave-one-out connected set, and the Andrews et al. (2008) on the largest connected set—all deliver similar results.

alternative approach involves shrinking each raw fixed effect to minimize its mean-squared-error.[28] We use covariance shrinkage, developed by Best et al. (2023), which extends the standard shrinkage method (Chetty et al., 2014a; Kane and Staiger, 2008) to account for the correlation between the two dimensions of fixed effects—including the negative correlation in the sampling error of bureaucrat fixed effects with that of country fixed effects caused by limited mobility bias. In order to apply covariance shrinkage, we bootstrap equation (2) at the appointment level[29] to estimate the variances and covariances of the sampling errors. Covariance shrinkage has the advantage of yielding shrunk fixed effects for each bureaucrat ($\hat{\theta}^{\text{CovSh}}_{b(c,t)}$), allowing us to compare different parts of the distribution, e.g. the 25th and 50th percentile.

The covariance shrinkage procedure inspires another bias correction method for the variance decomposition. Armed with the bootstrap estimates of the variances and covariances of sampling errors, we can calculate the bias in the plug-in estimators, $\mathsf{Var}(\hat{\theta}_{b(c,t)})$, $\mathsf{Var}(\hat{\gamma}_c)$, and $\mathsf{Cov}(\hat{\theta}_{b(c,t)}, \hat{\gamma}_c)$. We subtract the estimated biases to isolate $\mathsf{Var}(\theta_{b(c,t)})$, $\mathsf{Var}(\gamma_c)$, and $\mathsf{Cov}(\theta_{b(c,t)}, \gamma_c)$. We call this method *bootstrap bias correction* (see Appendix F for details). The bootstrap bias correction uses estimates of the sampling-error variance for all bureaucrats. It therefore overcomes a limitation of Kline et al. (2020), which faces difficulties imputing the sampling-error variance for single-appointment bureaucrats.[30]

4.3 Result: Bureaucrats Drive Korean Exports

Figure 2 plots the cumulative distribution function of raw bureaucrat fixed effects, showing substantial dispersion.[31]

Table 2 reports the variance decomposition results, correcting for bias using the Kline et al. (2020) method (preferred) in columns 1-2 and the bootstrap bias correction method in columns 3-4. The two bias correction methods yield similar results overall. The fact that both correction methods

[28] Although each fixed effect estimator is unbiased under the identifying assumption, it is noisy. Shrinking each estimate toward the mean introduces bias for greater precision.

[29] We bootstrap at the appointment level because, as discussed before, it is the variation across appointments that identifies the bureaucrat and country fixed effects.

[30] Using a leave-i-out estimator to estimate the i-specific variance of the sampling error is not applicable to single-appointment bureaucrats. The Bonhomme et al. (2023) algorithm we use for implementing Kline et al. (2020) handles the issue with a homogeneity assumption: The error variance for a single-appointment bureaucrat is set to equal the average among multi-appointment bureaucrats ever appointed to the same country. This assumption might potentially underestimate the error variance of single-appointment bureaucrats. In fact, Panel (a) of Appendix Figure A.3 shows that the bootstrap estimate of the variance of the sampling error is the largest for single-appointment bureaucrats and steadily decreases in the number of appointments.

[31] While raw fixed effects are unbiased under the identifying assumptions, sampling error inflates their dispersion relative to the true dispersion in underlying abilities. Appendix Figure A.4 compares the distributions of the raw fixed effects and the covariance-shrunk fixed effects. Shrinkage compresses the distribution; for example, covariance shrinkage reduces the difference between the 25th and 50th percentile by 30% and between the 50th and 75th percentile by 18%. Nonetheless, substantial dispersion remains, indicating that much of the dispersion in Figure 2 reflects true differences between bureaucrats.

produce extremely similar estimates of the true dispersion in bureaucrat ability offers confidence that we are appropriately accounting for noise also for single-appointment bureaucrats.

Bureaucrats explain a substantial amount of variation in Korean exports. Columns 2 and 4 show that they explain around 14% of the variation at the spell level. Furthermore, column 1 shows that one standard deviation of bureaucrat ability is estimated to be 0.318,[32] implying an increase in the dollar value of exports of approximately 37%.[33] Moreover, the magnitude is comparable to the policy's average effect—the effect of opening an office—of 0.321 (38%), estimated in Section 5. Hence, opening an office would not impact exports if the office is headed by a bureaucrat whose ability was one standard deviation below the mean.[34]

To put the magnitude into perspective, we compare the effect of a one standard deviation increase in bureaucrat ability to that of geographical distance. Assuming an elasticity of trade to distance of -1 (Anderson, 2011; Head and Mayer, 2014), it amounts to roughly the effect of reducing trade distance from London-Seoul (8,900km/5,500 miles) to Mumbai-Seoul (5,600km/3,500 miles).

Both columns 1 and 3 show that bureaucrats explain around $1/7$ as much variation as countries do. $1/7$ is relatively small compared to effect sizes in recent studies of public-sector managers using similar methodology. For example, managers in Fenizia (2022) explain 1/3 as much variation in the processing of social insurance claims as the offices they manage, and those in Otero and Muñoz (2025) explain 3/4 as much variation in mortality as the public hospitals they manage.[35] The comparatively small contribution of KOTRA bureaucrats is not surprising, given that the country fixed effects represent much more than just the overseas offices. The country fixed effects encapsulate time-invariant gravity variables, such as distance and market size, which are strong predictors of bilateral trade volume. It should be noted, however, that public hospitals, courts, and retail stores also represent more than just factors within the organization, e.g., they also represent the demographics or purchasing power of their location.

We find a negative covariance between bureaucrat and country fixed effects, suggesting that better bureaucrats work in smaller countries. In the bootstrap bias correction (column 3), the covariance is still negative but is smaller and statistically insignificant. The negative covariance is almost exclusively driven by single-appointment bureaucrats.[36] For multi-appointment bureaucrats,

[32] $0.318 = \sqrt{0.101}$

[33] The IHS transformation approximates the log transformation for sufficiently large values, so we approximate the percentage change as $(e^{0.318} - 1) \times 100 \approx 37$. Footnote 18 discusses why we prefer the IHS transformation and how we assess robustness to alternative transformations.

[34] Under the simplification that the average effect of office opening reflects the effect of an office headed by the average bureaucrat.

[35] Other papers studying bureaucrats in non-management roles similarly find that individuals matter more than in our setting. Best et al. (2023) find that individual procurement agents explain similar shares of the variation in procurement prices as the agencies for which they work. Dahis et al. (2023) find that judges matter 2/3 as much as courts in determining the number of cases disposed. Outside the public sector, Metcalfe et al. (2023) find that retail store managers explain 58% as much variation as store fixed effects in determining sales.

[36] The total covariance between bureaucrat and country effects can be decomposed into: a) the covariance within the

the covariance is negative but very close to zero. Moreover, we find that the average bureaucrat effect and the average country effect for multi-appointment bureaucrats are both larger than those for single-appointment ones, an observation we discuss at length in Section 4.6. Overall, bureaucrat and country fixed effects jointly explain around 90% of the spell-level variation in exports (after removing product-year fixed effects).

4.4 Diagnostics

In this section, we perform several diagnostic checks to assess the validity and reliability of our fixed-effect estimates. We begin by presenting event-study evidence that supports the identifying assumption, described in Section 4.1. Analogous event studies show that the estimated effects operate along both the extensive and intensive margins. Under the identifying assumption, the fixed effect estimates are unbiased but may still be noisy. We conclude this section by showing that our results are not driven by noise. Additional diagnostic checks are provided in Appendix D, corroborating the additive separability assumption of bureaucrat and country effects, as well as the robustness to different transformations of export values.

Are Bureaucrat Appointments Orthogonal to Export Trends?

We estimate event-study regressions to alleviate concerns that the appointment of KOTRA bureaucrats may not be orthogonal to underlying export trends, a violation of the identifying assumption that would result in biased estimates.

As explained in Section 4.1, if KOTRA strategically appointed high-ability bureaucrats to countries with increasing export trends, we would expect to observe differential pre-trends by ability. Additionally, we would expect a larger (spurious) increase in exports when a new bureaucrat arrives than corresponding decrease when a bureaucrat of the same ability departs.

$$y_{ept} = \eta_{ep} + \lambda_{pt} + \sum_{k \neq -2} \left(\alpha_k + \beta_k \cdot \hat{\theta}_e^{\text{new}} + \delta_k \cdot \hat{\theta}_e^{\text{old}} \right) \mathbb{1}\{t = T + k\} + \epsilon_{ept} \tag{4}$$

We estimate equation (4), which models how exports change around an event (e), a switch from one bureaucrat to another. It models exports as a time-varying function of the fixed effect of the new bureaucrat ($\hat{\theta}_e^{\text{new}}$) and the old bureaucrat ($\hat{\theta}_e^{\text{old}}$). An event e is uniquely defined by the country (c) and the year of the event (T), defined as the first full year after the new bureaucrat is appointed to country c. The key parameters of interest are β_k and δ_k. β_k measures the change in y associated with a 1 standard deviation increase in the ability of the newly appointed bureaucrat ($\hat{\theta}_e^{\text{new}}$) in year $T + k$, relative to the corresponding change in year $T - 2$.[37] δ_k is the counterpart for the outgoing

group of single-appointment bureaucrats, b) the covariance within the group of multi-appointment bureaucrats, and c) the covariance between the average bureaucrat effect and the average country effect across the two groups.

[37] To make β_k and δ_k more interpretable, we standardize $\hat{\theta}_e^{\text{new}}$ and $\hat{\theta}_e^{\text{old}}$ first.

bureaucrat. We take $T-2$ as the base year since it is the last full year when the old bureaucrat was in charge. We control for trends using product-year fixed effects (λ_{pt}) and for pre-event levels of exports using event-product fixed effects (η_{ep}).

Panel (a) of Figure 3 plots the event-study estimates, $\hat{\beta}_k$ and $\hat{\delta}_k$, obtained from equation (4) for the benchmark outcome variable: the inverse hyperbolic sine of exports. It shows that exports change sharply, over 1-2 years, in the direction of the ability of the incoming bureaucrat and symmetrically against the direction of the outgoing bureaucrat's ability.[38] In addition, the pre-trends in this specification are not statistically distinct from 0 and economically very small, providing further support that appointments are not strategically timed to coincide with increased export potential.

It may be surprising that there is a strong drop in exports upon the appointment of an ineffective bureaucrat. However, South Korean exports were growing at over 35% annually[39] with tremendous churn in exporter-importer matches. As the regression equations include product-year fixed effects, losing a good bureaucrat means exports drop *only relative to this trend*. Hence, our results imply reduced growth—fewer new matches—rather than negative growth.

Export Responses Occur at both Extensive and Intensive Margins.

We next aim to better understand the changes in exports due to the bureaucrat appointment. The results show that the effects operate through both the extensive and intensive margins. Moreover, bureaucrat effects are not confined to products with small export values.

Panels (b)-(f) of Figure 3 plot the event-study estimates, $\hat{\beta}_k$ and $\hat{\delta}_k$, obtained from equation (4) for alternative transformations of export values. All panels broadly confirm the qualitative message from Panel (a): Exports change sharply, over 1-2 years, in the direction of the ability of the incoming bureaucrat and symmetrically against the direction of the outgoing bureaucrat's ability, while pre-trends are statistically and economically absent.

Panel (b) indicates that a 1 standard deviation increase in bureaucrat ability causes a 3 percentage point increase in the number of products that Korea exports to the respective country.

Panels (c)-(f) isolate the intensive margin by setting $y_{cpt} = log(\max\{\text{export value}_{cpt}, m\})$ as the outcome. This transformation ensures a balanced panel by avoiding $log(0)$ while restricting attention to changes in exports only above m, where m = \$1,000, \$10,000, \$100,000, and

[38] Fenizia (2022) notes the possibility that idiosyncratic shocks to exports in ϵ_{ept} could also appear in bureaucrat fixed effects due to finite-sample estimation error, creating a spurious correlation in the post-period even in the absence of a causal relationship. Fenizia addresses the issue by estimating fixed effects for each value of k using data that excludes k. In Panel (b) of Figure 4, we report results from a much more conservative approach: We use bureaucrat fixed effects estimated using data that excludes the country altogether, e.g., for an event study of a bureaucrat switch in the UK office in 1981, we use fixed effects estimated using data excluding the UK. The results continue to support the identifying assumption. The increase in exports when gaining a bureaucrat is not larger than the decrease when losing one.

[39] From 1962 to 1981, South Korean exports rose from 55 million to 21 billion U.S. dollars, implying an annual growth rate of 36.7%.

$1,000,000, respectively.[40] Although export effects of bureaucrat ability persist for all values of m, the magnitudes decline as m increases, from around 55% to about 15%. These patterns are consistent with export promotion alleviating information frictions, which are typically more severe when existing exports to a given product–destination market are small (Volpe Martincus and Carballo, 2010).[41]

Overall, Figure 3 shows that bureaucrats affect both the initiation of exports for new products and the growth of exports for existing products, across all levels of export volumes.

Are the Fixed Effect Estimates Driven by Noise?

Next, we address concerns that the bureaucrat fixed effect estimates capture not only the true underlying bureaucrat ability but also the noise or idiosyncratic variations in the sample data, an issue aggravated when the number of observations per bureaucrat is small. The variance decomposition results discussed in Section 4.2 already account for this noise, but we perform a series of checks to reinforce the reliability of the estimates.

First, we perform a placebo check. We apply the variance decomposition on synthetic data where bureaucrats are randomly reshuffled across the 705 appointments in the full sample. Importantly, we break the true moves of bureaucrats, which identify the bureaucrat (and country) effects, while preserving all other features of the original data: which country-years form an appointment, the number of appointments of each bureaucrat, and the connectivity of the graph.[42] As the true moves are broken, the fixed effects of bureaucrats in the placebo should capture only noise and have no true explanatory power. Columns 5-6 of Table 2 report the results from the placebo variance decomposition.[43] Remarkably, both the variation in bureaucrat fixed effects, as well as the covariance between bureaucrat and country fixed effects, go to zero, indicating that the Kline et al. (2020) bias-correction method successfully removes the contribution of noise to the variance and covariance values. Further, the total variation explained by bureaucrat and country—Var(bureaucrat+country)—is substantially smaller in the placebo. These findings show that the results in columns 1-2 are not spurious, because if they were, we would expect results in columns 5-6 to resemble them.

Second, we test the out-of-sample predictiveness of the bureaucrat fixed effects. An implication of overfitting is that the fixed effect estimates would have low predictive power outside the sample.

[40] As throughout the paper, exports are measured in thousands of contemporaneous U.S. Dollars.

[41] Part of the difference in effect sizes across panels (c)–(f) may reflect mechanical attenuation from increasing the censoring threshold m. To reduce this mechanical attenuation, we exclude country–product combinations that show no variation in an event horizon, i.e., those without exports exceeding m during the event horizon.

[42] We identify the set of unique bureaucrat-country pairs in the original data, randomly reshuffle the matches, and merge them into the original data. For example, if bureaucrat 1 was appointed to country A from 1981 to 1983, a random bureaucrat X replaces bureaucrat 1 in country A for these years.

[43] We confirm that the total number of bureaucrats and countries, as well as the number of bureaucrats by the number of appointment spells, is maintained similarly to the baseline.

We directly test and confirm that our estimated bureaucrat effects are, indeed, predictive out of sample.[44] The most natural and conservative way in our setting to obtain fixed effects that are testable out of sample is to only use *other countries* to estimate the fixed effects. For example, to estimate the fixed effects of bureaucrats appointed to the UK, we obtain their fixed effects when excluding the UK from the sample. The estimation of such out-of-sample fixed effects is only possible due to the interconnectedness of our data. When leaving out one country, we always retain one very large connected set.

Panel (a) of Figure 4 displays a binned scatterplot of residual exports on in-sample (i.e., baseline) and out-of-sample bureaucrat fixed effects. By construction, the slope for the in-sample fixed effects equals 1. More interestingly, the out-of-sample fixed effects do strongly predict residualized exports, with a slope coefficient of 0.51. Furthermore, we replicate the event-study regression depicted in Panel (a) of Figure 3 using the out-of-sample fixed effects. Panel (b) of Figure 4 shows that new and old bureaucrat ability still statistically significantly predict export changes in the expected direction even when ability is estimated only using other countries. The fact that estimated bureaucrat effects are predictive out-of-sample makes it implausible that they are primarily capturing spurious correlations between bureaucrat appointments and underlying export trends.

Last, we conduct a comparative relevance check against ambassadors' effects on exports, reported in columns 7-8 of Table 2. Instead of KOTRA bureaucrats, we consider Korean ambassadors—who also follow rotational assignments—to each country with a KOTRA overseas office.[45] The connectivity of the ambassador-country graph closely mirrors that of the KOTRA bureaucrat-country graph. While ambassadors may influence Korean exports through various channels, such as negotiating trade agreements, resolving bilateral disputes, and networking with local business leaders, export promotion is not their prime focus. Instead, they have broad diplomatic responsibilities, making export promotion just one, if at all, of many duties. In contrast, KOTRA overseas office managers are dedicated solely to increasing exports. Therefore, if our decomposition exercise captures true effects rather than noise, we should expect to find that KOTRA bureaucrats have a larger impact on exports than ambassadors. Column 8 shows that while ambassadors do explain some variation in exports, they account for a much smaller share of the total variation (4%) than KOTRA managers (14%, column 2).

Overall, these results support the interpretation that bureaucrat fixed effects reliably identify the causal impact on exports of an individual bureaucrat managing a KOTRA office.

[44] Additionally, Appendix Table B.1 reports limited additional explanatory power from allowing bureaucrat effects to differ between appointments; the gains in the adjusted R^2 is minimal when we replace bureaucrat fixed effects with bureaucrat×country fixed effects in estimating equation (1) (columns 3 and 4). The stability of the bureaucrat fixed effects across appointments indicates that we are capturing true signals rather than mere noise.

[45] The only country that drops out is Slovenia, which has a KOTRA office but not a Korean embassy.

4.5 Mechanisms: What Do Good Bureaucrats Do Differently?

This section examines why some bureaucrats generate substantially larger export gains than others. We find that high-ability bureaucrats expand exports more where information frictions are severe, increase the responsiveness of exports to market conditions, and reallocate office activities toward broader matching between Korean firms and export opportunities. Together, the evidence points to a common mechanism: High-ability bureaucrats enhance the transmission of information between foreign markets and Korean exporters. This interpretation aligns with Kim and Kim (2024), who show that the public dissemination of foreign buyer information by KOTRA significantly increases exports to targeted destination–industry pairs.[46]

Bureaucrat Ability Matters More When Information Is Valuable

Figure 5 shows that the export effects of bureaucrat ability are systematically stronger when information frictions are more severe and when new export opportunities are emerging.

$$y_{ecpt} = \eta_{ep} + \lambda_{pt} + \sum_{k \neq -2} \Bigg[\left(\alpha_k^1 + \beta_k^1 \, \hat{\theta}_e^{\text{new}} + \delta_k^1 \, \hat{\theta}_e^{\text{old}} \right) \mathbb{1}\{\text{ProductGroup}_{ep} = 1\} + \left(\alpha_k^2 + \beta_k^2 \, \hat{\theta}_e^{\text{new}} + \delta_k^2 \, \hat{\theta}_e^{\text{old}} \right) \mathbb{1}\{\text{ProductGroup}_{ep} = 2\} \Bigg] \mathbb{1}\{t = T + k\} + \epsilon_{ecpt} \tag{5}$$

Panels (a)-(c) report estimates based on equation (5), an event-study regression that estimates how exports evolve for two groups of products when an old bureaucrat (with fixed effect $\hat{\theta}_e^{old}$) is replaced by a new bureaucrat (with fixed effect $\hat{\theta}_e^{new}$). Other than distinguishing between product groups, the specification follows equation (4).[47]

Panel (a) shows that the effect of bureaucrat ability is larger for differentiated products than for non-differentiated goods. This pattern points to a more important role for bureaucrats when informational barriers are higher. Panel (b) shows that the effects are also stronger for products facing more volatile destination-market demand. In such settings, effective bureaucrats may help firms navigate uncertainty by identifying opportunities and conveying timely information about shifting market conditions.

Panels (c) and (d) show that high-ability bureaucrats are particularly effective at enabling exporters to scale up when conditions are favorable in a sector. Panel (c) shows that the effect of bureaucrat ability is higher for heavy and chemical industry (HCI) products during the HCI drive

[46] Kim and Kim (2024) use import inquiries from foreign firms as a measure of buyer information, which we also use in this section. The analysis here differs mainly in its focus on heterogeneity across individual bureaucrats. The analysis also differs by examining a broader set of intermediate office outputs beyond inquiries, and by using the full universe of inquiry data rather than a random subset.

[47] In Panels (a) and (b), each product belongs to either product group 1 or 2, which implies $\hat{\beta}_k^1$ can be interpreted as the effect of ability for product group 1. On the other hand, in Panel (c) product group 2 is defined as *all products*, so $\hat{\beta}_k^1$ can be interpreted as the *differential* effect of ability for product group 1 *relative to all other products*.

period, but not before or after.

Panel (d) reports estimates based on equation (6) which links bureaucrat ability to market conditions by interacting bureaucrat switches with measures of import demand and export supply. The variable demand_{cpt} captures destination-specific import demand for product p—measured by country c's imports of p from all countries but Korea—while supply_{cpt} measures Korea's revealed comparative advantage in product p—measured by Korean exports of p to all destinations but c. The event-study coefficients β_k^D and δ_k^D capture how the export effect of incoming and outgoing bureaucrat ability varies with demand conditions, while β_k^S and δ_k^S capture the corresponding interactions with supply conditions.

$$
\begin{aligned}
y_{ecpt} = \eta_{ep} + \lambda_{pt} + \left(\psi^D + \beta^D \hat{\theta}_e^{\text{new}} + \ \delta^D \hat{\theta}_e^{\text{old}}\right)\text{demand}_{cpt} + \left(\psi^S + \beta^S \hat{\theta}_e^{\text{new}} + \delta^S \hat{\theta}_e^{\text{old}}\right)\text{supply}_{cpt} \qquad (6) \\
+ \sum_{k \neq -2} \Bigg[\alpha_k + \beta_k \hat{\theta}_e^{\text{new}} + \delta_k \hat{\theta}_e^{\text{old}} + \left(\psi_k^D + \beta_k^D \hat{\theta}_e^{\text{new}} + \delta_k^D \hat{\theta}_e^{\text{old}}\right)\text{demand}_{cpt} \\
+ \left(\psi_k^S + \beta_k^S \hat{\theta}_e^{\text{new}} + \delta_k^S \hat{\theta}_e^{\text{old}}\right)\text{supply}_{cpt} \Bigg] \mathbb{1}\{t = T + k\} + \epsilon_{ecpt}
\end{aligned}
$$

Panel (d) shows that high ability bureaucrats cause export responses to be amplified for products experiencing differentially improving demand in a destination, as well as for products which see an increase in Korean comparative advantage (supply).

Overall, Figure 5 suggests that high-ability bureaucrats increase exports primarily by strengthening the link between export opportunities—market conditions and policy—and export outcomes, especially where information frictions are most binding.

Bureaucrat Ability and the Allocation of Office Activities

Table 3 provides complementary evidence on how high-ability bureaucrats organize the activities of the overseas offices they manage. The table reports correlations between bureaucrat fixed effects and measures of office output, including business inquiries, market reports, and trade fair activity.[48] All specifications include country and year fixed effects, ensuring that the estimates are not driven by differences in market size or changes over time in how KOTRA operated.

Columns 1-2 show that offices attract a larger volume of inquiries when managed by a high-fixed-effect bureaucrat. The inquiries also span a broader set of products. At the same time, inquiries under their management tend to be less narrowly specified, with shorter product descriptions. The pattern indicates a shift away from highly detailed, specific requests that may be time-intensive—thereby crowding out more productive activities—or may discourage follow-up by Korean firms.

A similar reallocation appears in the reports which offices produce. Column 3 shows that offices produce more general market reports but fewer product-specific reports when managed by a high-ability bureaucrat, indicating a prioritization of broader market intelligence over narrowly defined

[48]The detailed description of how we collected the data on these office activities are in Appendix C.

product-level information.[49] However, there is a notable exception to this pattern: Conditional on the total number of product-specific reports, high–ability bureaucrats produced more reports related to HCI products during the HCI drive, but not before or after (columns 4-6). The evidence suggests that high-ability bureaucrats target their granular reporting to those sectors where national policy makes such information especially valuable.

Differences also emerge in trade fair activity. Column 7-8 shows that offices managed by higher-ability bureaucrats organize more trade fairs, while not at a larger scale. Speaking to the importance of fairs as a direct venue for generating exports, column 9 shows that higher–ability bureaucrats generate higher total sales at trade fairs.

Taken together, the evidence in Table 3 shows that high-ability bureaucrats reallocate effort toward activities that gather broader-scope information and directly connect buyers and sellers. These activity patterns underpin the stronger export responsiveness shown in Figure 5 by improving the transmission of information.

4.6 Policy Implications: Screening Bureaucrats and the Policy's Effect

Bureaucrats are Screened Based on Performance in 1st Appointments

Our findings thus far show that individual bureaucrats' ability significantly impacts the effectiveness of the export promotion policy. This result highlights the importance of selecting high-ability bureaucrats. What could be a viable way for the state to select high-ability bureaucrats when information about ability is incomplete—potentially even to the bureaucrats themselves?

We show that the descriptive appointment patterns of bureaucrats are consistent with KOTRA recognizing and managing the uncertainty over bureaucrat ability.

First, residualized exports during a bureaucrat's first appointment are predictive of whether they are reappointed in the future. Figure 6 presents a visual illustration. It plots the probability density function of (a) residualized exports by appointment, and (b) covariance-shrunk bureaucrat fixed effects, splitting the sample by whether a bureaucrat has more than one appointment as manager of a country office over their career. The distributions exhibit a substantial mass of lower-ability bureaucrats whose careers only include one appointment.

Importantly, the distributional gap between single- and multi-appointment bureaucrats likely reflects differences in true ability rather than measurement error. It is not driven by greater noise in the fixed effects of single-appointment bureaucrats, as Panel (b) uses covariance-shrunk estimates to account for noise. It is also not driven by first appointments being inherently different, as Panel (a) shows that among multi-appointment bureaucrats, the distribution of residualized exports in first appointments closely resembles that in subsequent appointments. First appointments are therefore

[49] Example of a general report title: "Argentina Abolishes the Import Quota System." Example of a product-specific report title: "TV/Video Recorder/Stereo Market Profitability in West Germany is Worsening."

neither inherently noisier nor associated with systematically worse performance.

This visual observation is buttressed by a regression of a bureaucrat's total number of appointments on residualized exports during the first appointment. We include dummies for the year of first appointment to rule out various confounding factors, such as the mechanical truncation of careers for bureaucrats whose first appointments occur near the end of the sample and systematic cohort differences in bureaucrat characteristics. We find a positive and statistically significant correlation, with a regression coefficient of 0.240 (standard error: 0.112). The correlation survives alternative specifications. For instance, the coefficient increases to 0.430 (standard error: 0.109) when we replace first-appointment residualized exports with an indicator for exceeding the 25th percentile.

Second, first appointments tend to be to less important countries, while third appointments tend to be to the most significant markets. Figure 6 illustrates the shares accounted by the first, second, and third appointments of bureaucrats by country importance. Importance as Korea's trade partner is proxied by how early the KOTRA office opened, as well as the value of the country fixed effect in equation (1). Panel (a) shows that first appointments make up the largest share of appointments to offices that opened late (after 1975), whereas third appointments constitute the lion's share of appointments to offices that opened early (before 1968). Second appointments form an intermediate case. Panel (b) paints a similar picture. In fact, no bureaucrat is appointed to the countries with exceptionally high fixed effects (the U.S., followed by Japan) in their first appointment. Offices in these countries are primarily managed by bureaucrats in their third appointments.

Together, these two patterns align with a model where the organization learns about bureaucrat ability as office manager by initially appointing them to less important countries. It then uses information on performance during the first appointments to screen out low performers from getting reappointed, particularly to important countries.[50] An appointment as a manager of a small country office is plausibly much more informative about a bureaucrat's ability than a subordinate role in a more important country. Therefore, assigning bureaucrats of uncertain ability to manage small countries may be justified, even if placing them in subordinate roles in key markets could have a larger aggregate impact on exports. Pastorino (2024) finds a similar dynamic in a private firm, where worker career paths are shaped by the firm's consideration of how different job assignments vary in the degree to which they reveal a worker's ability.

Crude Counterfactuals—Gains from Screening Bureaucrats

Table 4 provides three simple counterfactuals to assess how much exports could change from screening out ineffective bureaucrats based on their first-appointment performance. We focus on the differential selectivity of single- and multi-appointment bureaucrats in the data, which may be

[50]Note that there could be many other explanations, as the appointments we observe are an equilibrium outcome. For example, we cannot detect voluntary quits or observe various individual-specific or organization-wide Human Resources considerations.

interpreted as the prevailing "screening capacity" of KOTRA.

To investigate the effect of the current degree of screening, Counterfactual (1) simulates export changes in the absence of screening. We eliminate the difference in mean ability between single- and multi-appointment bureaucrats by setting both to the average ability across all bureaucrats. The dollar value of exports increases by 11% for appointments led by single-appointment bureaucrats, while it decreases by 13% for those led by multi-appointment bureaucrats. Since multi-appointment bureaucrats have more appointments, the overall effect is negative, at -6.5%. The sizable effect indicates that the current screening process plays a significant role in shaping export performance.

To assess the potential gains from more intensive screening, counterfactual (2) simulates the change in exports if single-appointment bureaucrats were as effective as multi-appointment ones. We set the mean ability of single-appointment bureaucrats equal to that of multi-appointment bureaucrats. This scenario may be relevant if KOTRA could operate new offices in smaller countries as a screening stage before assigning bureaucrats to the countries in our data. Export value increases by 25% for appointments headed by single-appointment bureaucrats. As these appointments constitute around 30% of all appointments, total export value rises by 6.8%. The similarity in effect sizes between counterfactuals (1) and (2) suggests that KOTRA may only be capturing half of the potential gains from screening bureaucrats according to its current screening capacity.

Counterfactual (3) replaces bureaucrats in the bottom quartile of ability with the median bureaucrat. We obtain these percentiles from the distribution of covariance-shrunk fixed effects, which addresses limited mobility bias—bureaucrats' fixed effects are estimated with varying degrees of sampling error. The value of exports increases by 8.3%, larger than the effects in counterfactuals (1) or (2), highlighting substantial gains from improving selection at the lower end of the ability distribution, where performance is most deficient.

These counterfactuals highlight that strengthening implementation capacity through improved bureaucratic selection could yield significant gains in policy effectiveness.

5 What is the Policy's Average Effect?

In this section, we exploit the staggered rollout of overseas export promotion offices to identify their causal effect on Korean exports to destination countries.[51] The average effect of these offices is policy-relevant in its own right. More importantly for the paper's main research question, it provides a natural benchmark against which to compare the individual bureaucrat effects estimated in Section 4. We estimate a 38% increase in export value 9–11 years after an office opens, comparable to the effect of a 1 standard deviation increase in bureaucrat ability. Assuming that the average effect of an office corresponds to the effect of an office headed by the average bureaucrat, an office with a

[51] For countries with multiple offices, we focus on the opening of the first office.

bureaucrat one standard deviation below average would have no effect on exports.

Our setting is unique in enabling this direct comparison between the average effect of a policy and the heterogeneity in that effect arising from individual bureaucrats' implementation.[52] The comparison is possible because (1) we observe a large number of office openings, and (2) South Korean bilateral exports are well defined even in years without an office.

Panel (b) of Figure 1 depicts the staggered rollout of offices by plotting the cumulative number of offices by year since KOTRA's establishment in 1962. Over the next two decades, offices opened in 75 countries—nearly four new countries each year. After the intensive initial rollout, KOTRA's expansion slowed markedly.[53] Thus, our empirical analysis focuses on the initial office openings (1962-1981). Panel (c) of Figure 1 maps the countries with office openings between 1962 and 2001.

5.1 Identification: Effect of Offices on Exports

To estimate the effect of an export promotion office, the ideal experiment would randomly allocate fully-developed offices to some countries but not others. As this is not feasible, we instead exploit the staggered rollout of offices to countries. We allow for dynamic effects, as offices may require time to become fully operational.

$$y_{cpt} = \lambda_{pt} + \gamma_{cp} + X'_{cpt}\xi + \sum_{k \neq -1} \pi_k D^k_{ct} + \epsilon_{cpt} \tag{7}$$

As a first step, we estimate equation (7). λ_{pt} indicates product-year fixed effects and γ_{cp} indicates country fixed effects that may differ at the product level.[54] D^k_{ct} are dummies equal to 1 if year t is k years after the first office opened in country c. π_k corresponds to the effect of an office that has been open for k periods. X_{cpt} is a vector of time-varying controls. The main specification uses the IHS of the value of South Korean exports as the outcome variable and does not include any controls.

Two central assumptions must be met for $\hat{\pi}_k$ to represent unbiased estimates of the causal effect of an office k years after opening. The first is a parallel trends assumption. We assume that counterfactual export trends, in the absence of an office, do not systematically differ between countries treated in year g and the control group in years $g + k$ (where $k > 0$).[55] The assumption would be violated if office openings were timed and placed based on time-varying export trends, for instance in countries where exports were projected to grow. There is little indication of differential

[52] By contrast, settings in papers using similar two-way fixed effects designs do not permit construction of such a benchmark (Fenizia, 2022; Best et al., 2023; Otero and Muñoz, 2025; Metcalfe et al., 2023).

[53] Only 13 offices opened over the next twenty-year period (1982–2001).

[54] We prefer country-product fixed effects to simpler country fixed effects because the former absorb additional variation unrelated to the effect of office openings. In practice, the estimates are largely unaffected by this choice.

[55] Persistent level differences between the treatment and control group do not violate this assumption. In line with the recent difference-in-differences literature, we carefully select the sample such that either the never-treated or the not-yet-treated form the control group (Callaway and Sant'Anna, 2021).

pre-trends (discussed in Section 5.2), and non-Korean exports to the country do not exhibit an upward trend after office openings (discussed in Section 5.3), suggesting that offices are not placed in countries with faster import demand growth. Moreover, we show that the rollout order of European offices is almost fully explained by pre-determined (1962) import market size, suggesting limited scope for timing office openings either strategically or coincidentally.

The second central assumption is the Stable Unit Treatment Value Assumption (SUTVA), meaning an office affects exports only to the country where it is located and not exports to other countries. If this assumption is violated, the estimated effect may reflect reallocation of export flows across destination countries rather than isolating the absolute effect of offices on exports to a given country. We assess this concern by comparing results using never-treated and not-yet-treated control groups, which may differ in their susceptibility to spillovers. In particular, the not-yet-treated group is more likely to be geographically or economically close to treated countries and thus potentially more exposed to spillover effects than the never-treated group. The similarity of results across these specifications provides reassurance regarding potential SUTVA violations.[56]

A third identifying assumption is the absence of anticipation effects. This assumption would be violated if exports were affected even before office openings. Negative anticipation effects—an Ashenfelter (1978) dip—would lead to overestimation. We find no evidence of anticipation.[57]

Addressing Concerns about Staggered Difference-in-Differences

A recent literature on staggered two-way fixed effects regressions clarifies a number of circumstances under which the identification of causal effects fails. For example, the specification would produce biased estimates if there are dynamic treatment effects and already-treated units are included in the control group. Therefore, our main control group is the set of never-treated countries. For estimation, we construct the treatment group as a panel of ever-treated countries that is balanced within the event horizon of the office opening, and the control group as a panel of never-treated countries that is balanced across all the years included in any event horizon.

A balanced panel of treated countries implies that we cannot estimate the effect of some of the earliest openings, depending on the length of the pre-period in the specification. In specifications with five pre-periods, for instance, countries whose first office opened between 1962 and 1966 are excluded from the treatment group, because our export data starts in 1962.[58] Panels (a) and (b) of Figure 7 show that the pre-trends remain near-zero and effect sizes remain comparable whether

[56]Further, SUTVA violations are most concerning if they generate upward bias in the estimated effect of offices. Alfaro-Ureña et al. (2023) provide a rationale for viewing exports to different countries as complements, implying that SUTVA violations, if present, would be more likely to bias our estimates downward rather than upward.

[57]Positive anticipation may occur if firms increase exports today in expectation that KOTRA will open an office in the future. Such positive anticipation would lead us to underestimate the office effect. In all specifications using the never-treated control group, we find no evidence of positive anticipation.

[58]Territories with office openings in 1962: U.S., Thailand, Taiwan; 1964: Japan, Singapore, Indonesia, South Vietnam; 1965: Philippines, Peru, Kenya, Hong Kong, Iran, UK; 1966: Italy, Netherlands, Panama, ‰Nigeria.

we include countries with openings between 1967 and 1981—allowing for five pre-periods—or between 1964 and 1981—allowing for just two pre-periods but offering greater power by including 14 additional office openings.

The main disadvantage of estimating the treatment effect relative to the never-treated is that these may be the countries least comparable to the treated countries. To increase comparability, we exclude the smallest economies—specifically, all never-treated economies with a population below one million in 1962. Our finding of parallel trends in the pre-period, across various sample criteria and outcome variables,[59] adds support that the never-treated group is comparable.

5.2 Results: Effect of Office Opening on Exports

Panels (a) and (b) of Figure 7 plots $\hat{\pi}_k$, the estimated effects of the first overseas export promotion office in a destination country around the year of the opening. The pre-period coefficients are economically small and statistically indistinguishable from zero, providing support for the parallel trends assumption.

The figure further shows that the opening of an export promotion office is associated with an increase in Korean exports to that destination. The estimates increase over time, suggesting that it takes time for an office to reach its full effect. The point estimates stabilize 9-11 years after the office opening. The average point estimate in those years is 0.321, indicating a 38%[60] higher export value relative to the control group in the higher-powered regression (Panel (b)).

The opening of an office also increases KOTRA activity specific to that country. Appendix Figure A.6 plots the results from estimating equation (7) as before, but replacing the outcome variable with three country-specific activities: the number of market reports, the number of product-specific reports, and the number of business inquiries from potential importers.[61] Like exports, they are each transformed by the IHS to account for zeros in the data. The pre-treatment coefficients are statistically indistinguishable from zero for all panels, again supporting the parallel trends assumption. Panel (c), with just one pre-period, allows for a wider treatment window (openings between 1975 and 1981) and shows an increase in the number of inquiries. Panel (d), covering openings between 1978 and 1981, also shows an overall increase, although the coefficients are noisier and less statistically significant individually.

5.3 Robustness and Validity Checks

In this section, we discuss the parallel trends assumption beyond checking for the absence of differential pre-trends.

[59] Besides the main outcome variable of Korean exports, we also consider: non-Korean exports (as a placebo), an indicator for positive Korean export value, the number of KOTRA market reports, and the number of import inquiries.

[60] We approximate the percentage change using $(e^{0.321} - 1) \times 100 \approx 38$.

[61] The detailed description of how we collected the data on these office activities are in Appendix C.

Other robustness checks are performed in Appendix H. There, we examine the robustness of the results to (1) using a not-yet-treated control group following Callaway and Sant'Anna (2021), and (2) considering only the extensive margin of exports—to show that the increase in exports is not an artifact of the IHS transformation of export values.

No Increase in Import Demand upon Office Opening

A country's export promotion office may be strategically opened at a time when import demand in that country is expected to increase, which would violate the parallel trends assumption. We address this concern in two ways.

First, instead of South Korean exports to a country, we use non-Korean exports to that country as the outcome variable in equation (7). The coefficients from this placebo test are reported in Panel (c) of Figure 7.[62] The point estimates are close to zero both before and after an office opening, with wide confidence intervals, suggesting that office openings do not systematically coincide with increases in import demand.

Second, we control for non-Korean exports, proxying for overall import demand in the destination, while maintaining Korean exports as the outcome. Panels (a) and (b) of Appendix Figure A.5 show that the estimates from this specification remain largely unchanged compared to the baseline. Given South Korea's rapid economic growth, it may be that the relationship between Korean exports and destination import demand changes over time. Panels (c) and (d) show that estimates remain similar even when we control for non-Korean exports interacted with year dummies.

Rollout Follows Pre-Determined Gravity Variables

We further show that the year in which a country's first office opened mainly depended on pre-determined, time-invariant factors. If their effects on exports are also time-invariant, these factors are absorbed in the country fixed effects (γ_c). Even if the their effects vary over time, the pre-determined order of the rollout makes it unlikely that openings coincided with counterfactual changes in exports.

To predict office openings, we draw on insights from the gravity equation, where distance and market size are the strongest predictors of bilateral trade flows. Apart from the U.S., the first KOTRA offices opened in Taiwan, Thailand, Japan, Singapore, Indonesia, and South Vietnam—among the geographically closest non-communist territories.[63]

Within Europe, the distance from South Korea does not vary much between countries,[64] so the main predictor would be market size.[65] As there was no KOTRA office in Europe until 1965, we use 1962 non-Korean exports as a pre-determined measure of a destination's market size. Panel

[62]The coefficients are robust to alternative event-horizon lengths, which alter the sample by including the earliest office openings (Panels (e)–(f) of Appendix Figure A.5).

[63]North Korea, China, the USSR, and North Vietnam were ideological opponents of South Korea.

[64]The distance between Athens and Seoul is 96% of that between London and Seoul.

[65]Other determinants of trade, such as language or cultural distance, also vary little within Europe.

(d) of Figure 7 plots each country's rank by office-opening year against its rank by 1962 market size, for the 17 European countries that received offices during the main rollout period (1962-1981). Notably, the UK was the biggest market (rank 1) and was the first to receive an office (rank 1), while Portugal was the smallest market (rank 17) and was the last to receive an office (rank 17). Across all 17 countries, the correlation is very strong, at 0.87. The rigidity in the rollout schedule counters concerns that KOTRA timed openings based on export trends or projections, which would have violated the parallel trends assumption.

To sum up, the absence of changes in non-Korean exports following office openings, together with evidence that the rollout timing was driven by stable factors, alleviates concerns that openings were strategically timed based on market-specific shocks or anticipation of export growth. These findings support a causal interpretation of the estimated office effects.

6 Conclusion

This paper shows that the effectiveness of an industrial policy can depend critically on the bureaucrats who carry it out. Set in South Korea, the archetypal case of export-led growth and accompanying state intervention, our work offers a unique quantification of the link between the policies of the East Asian miracle and bureaucratic capacity, advancing our understanding of how human and organizational inputs shape macroeconomic development trajectories.

Our findings contribute to the resurgent debate on industrial policy by revealing why successful policies resist straightforward replication, even when baseline economic conditions appear similar. This context-dependence stems from discretionary implementation processes where quality, motivation, and tacit knowledge of bureaucrats prove decisive. This insight extends to other policy domains modeled after past successes—public health systems, education reforms, infrastructure projects, and a broad set of firm-directed policies—particularly when effective implementation requires tailoring interventions to specific beneficiaries or locations.

By showing that bureaucratic capacity matters for exports, we expand the study of state capacity to an economic outcome of substantial interest that is not directly controlled by the state. Our findings further imply that the state can affect the domestic economy through a policy the state implements *outside* its borders, revealing the extended reach of state capacity in global markets.

More broadly, our results speak to why bureaucracies succeed or fail in achieving stated policy goals. While the literature often frames bureaucratic performance through the lens of agency problems, we highlight idiosyncratic bureaucratic capability as a critical factor. We also demonstrate how states can improve bureaucrat selection under conditions of incomplete information—and potentially two-sided uncertainty about bureaucrat capabilities—through experimentation with moderate stakes.

The contemporary resurgence of industrial policy across both developed and developing

economies (Juhász et al., 2024) makes our findings particularly timely. As nations respond to climate imperatives, technological competition, supply chain vulnerabilities, and geopolitical realignments, they are deploying increasingly ambitious industrial policies. However, our research suggests that the effectiveness of new industrial policies may depend not only on policy design but also on parallel investments in the bureaucratic infrastructure needed to deliver them.

Bibliography

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

Aghion, P., J. Cai, M. Dewatripont, L. Du, A. Harrison, and P. Legros (2015). Industrial Policy and Competition. *American Economic Journal: Macroeconomics 7*(6), 1–32.

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*. 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 171*(3), 673–697.

Ashenfelter, O. (1978). Estimating the Effect of Training Programs on Earnings. *The Review of Economics and Statistics 60*(1), 47–57.

Ashraf, N., O. Bandiera, and B. K. Jack (2014). No margin, no mission? a field experiment on incentives for public service delivery. *Journal of Public Economics 120*, 1–17.

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.

Bagir, Y. K. (2020). Impact of the presence of embassies on trade: Evidence from turkey. *World Trade Review 19*, 51–60.

Bark, D. S. (1967). *Public Personnel Administration in Korea*. Ph.D. thesis, University of Minnesota. unpublished.

Barwick, P. J., M. Kalouptsidi, and N. B. Zahur (2025, 03). "Industrial Policy Implementation: Empirical Evidence from China's Shipbuilding Industry". *The Review of Economic Studies 92*(6), 3611–3648.

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). 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.

Chen, J. and J. Roth (2023). Logs with Zeros? Some Problems and Solutions. *The Quarterly Journal of Economics 139*(2), 891–936.

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). 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). 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.

Choi, J. and Y. Shim (2026). Industrialization and the Big Push: Theory and Evidence from South Korea. *Review of Economics and Statistics*.

Dahis, R., L. Schiavon, and T. Scot (2023). Selecting Top Bureaucrats: Admission Exams and Performance in Brazil. *The Review of Economics and Statistics 107*(2), 408–425.

Dal Bó, E., F. Finan, N. Y. Li, and L. Schechter (2021). Information Technology and Government Decentralization: Experimental Evidence From Paraguay. *Econometrica 89*(2), 677–701.

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 (2023). Demand-Side Constraints in Development: The Role of Market Size, Trade, and (In)Equality. *Econometrica 91*(6), 1915–1950.

Gulzar, S. and B. J. Pasquale (2017). Politicians, Bureaucrats, and Development: Evidence from India. *American Political Science Review 111*(1), 162–183.

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.

Hayakawa, K., H. Lee, and D. Park (2014). Do Export Promotion Agencies Increase Exports? *The Developing Economies 52*(3), 241–261.

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.

Jia, R. (2024). Pollution for promotion. *The Journal of Law, Economics, and Organization*, 8756–6222.

Johnson, C. (1982). *MITI and the Japanese Miracle: The Growth of Industrial Policy, 1925-1975*. Stanford University Press.

Jones, L. P. and S. Il (1980). *Government, Business, and Entrepreneurship in Economic Development: The Korean Case*. 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 (2024). Measuring 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 16*, 213–242.

Kane, T. J. and D. O. Staiger (2008). Estimating Teacher Impacts on Student Achievement: An Experimental Evaluation. Working Paper 14607, National Bureau of Economic Research.

Khan, A. Q., A. I. Khwaja, and B. A. Olken (2016). Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors. *Quarterly Journal of Economics 131*(1), 219–271.

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

Kim, C. and D. Kim (2024). Trade Search Frictions: Public Information, Exports, and Concentration. Working Paper.

Kim, M., M. Lee, and Y. Shin (2022). The Plant-Level View of an Industrial Policy: The Korean Heavy Industry Drive of 1973. Working Paper.

Kim, T. and C. J. Baik (2011). Taming and Tamed by the United States. In B.-K. Kim and E. F. Vogel (Eds.), *The Park Chung Hee Era : The Transformation of South Korea*, pp. 58–84.

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

Lane, N. (2025). Manufacturing Revolutions: Industrial Policy and Industrialization in South Korea. *The Quarterly Journal of Economics*.

Leaver, C., O. Ozier, P. Serneels, and A. Zeitlin (2021). Recruitment, Effort, and Retention Effects of Performance Contracts for Civil Servants: Experimental Evidence from Rwandan Primary Schools. *American Economic Review 111*(7), 2213–46.

Lin, J. Y., S. Liu, M. Xu, and H. Zhang (2024). Learn in order to practice: The effect of political rotation on local comparative advantage in china. *SSRN Electronic Journal*.

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*.

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

Mitrunen, M. (2024). War Reparations, Structural Change, and Intergenerational Mobility. *The Quarterly Journal of Economics 140*(1), 521–584.

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 (2025). Managers and Public Hospital Performance. *American Economic Review*.

Pastorino, E. (2024). Careers in Firms: The Role of Learning about Ability and Human Capital Acquisition. *Journal of Political Economy 132*(6), 1994–2073.

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

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics 48*(1), 7–35.

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

Studwell, J. (2013). *How Asia Works: Success and Failure in the World's Most Dynamic Region*. Grove Press and Grove/Atlantic, Inc.

Udell, J. I. (1965). Report on KOTRA for Korea Trade Promotion Corporation. Technical report, United States Operations Mission/Korea, Agency for International Development/Washington.

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.

Figure 1: Korean Exports and Export Promotion Offices

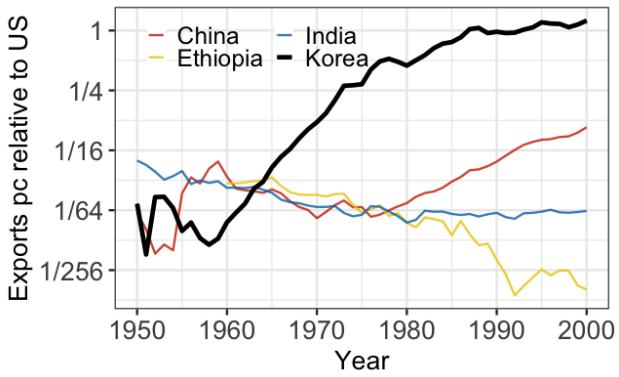

(a) Korean Exports relative to the U.S. by Year

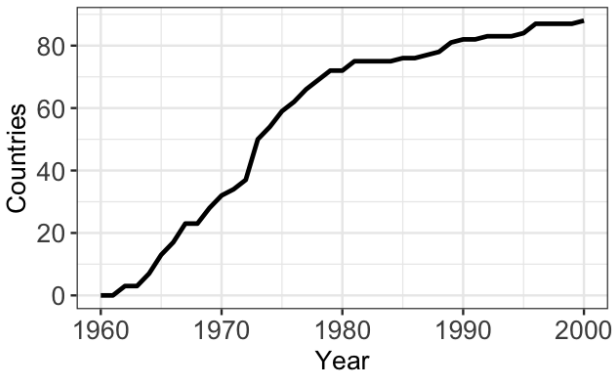

(b) Countries with Export Promotion Offices by Year

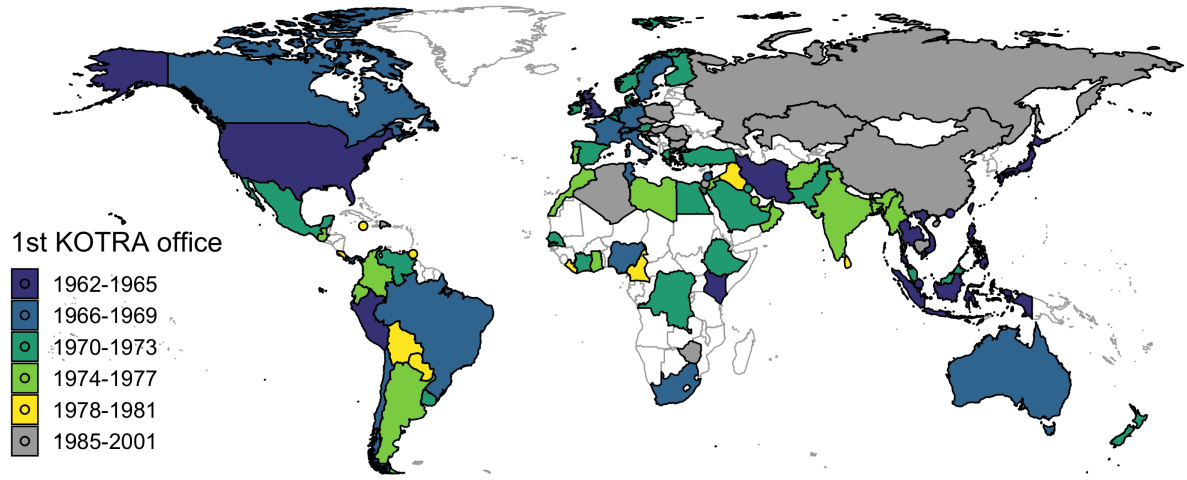

(c) Map of Export Promotion Offices

Notes: Panel (a) displays exports per capita relative to the U.S. from 1952 to 2000, for South Korea and a selected group of other countries. Data on exports and population are obtained from International Monetary Fund (2023): Direction of Trade Statistics. Panel (b) presents the cumulative number of countries with KOTRA offices by year. Panel (c) shows the geographic distribution of Korean export promotion offices; colors indicate the opening year of the first office in each country. Most offices remain open from the opening year for the entire period of study. White indicates territories which have not received an office by 2001. Back to Sections 2, 2.2, and 5.

Figure 2: Distribution of Raw Bureaucrat Fixed Effects

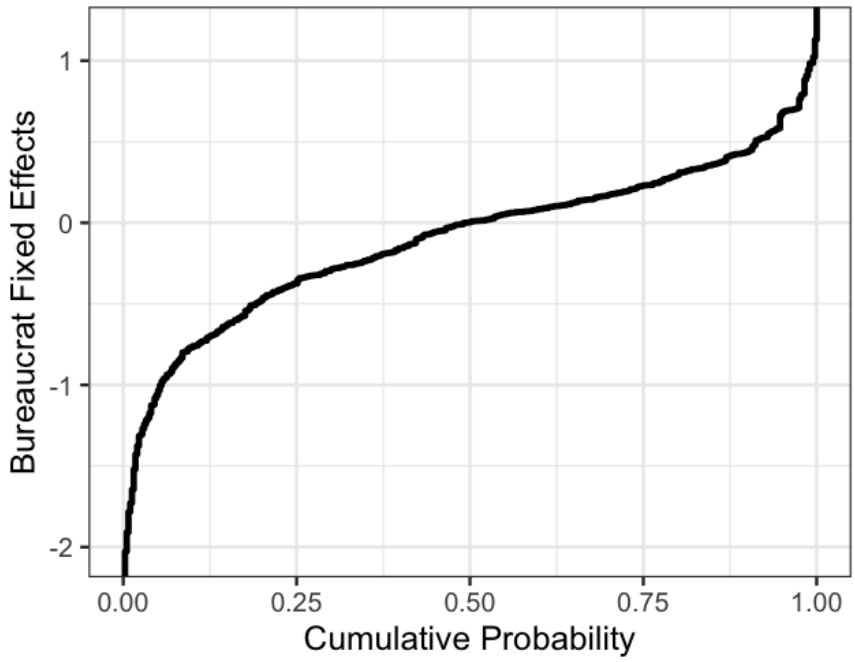

Notes: The figure shows the cumulative distribution function of bureaucrat fixed effects, estimated from equation (2). Back to Section 4.3.

Figure 3: Export Response to Bureaucrat Switches by Extensive and Intensive Margins

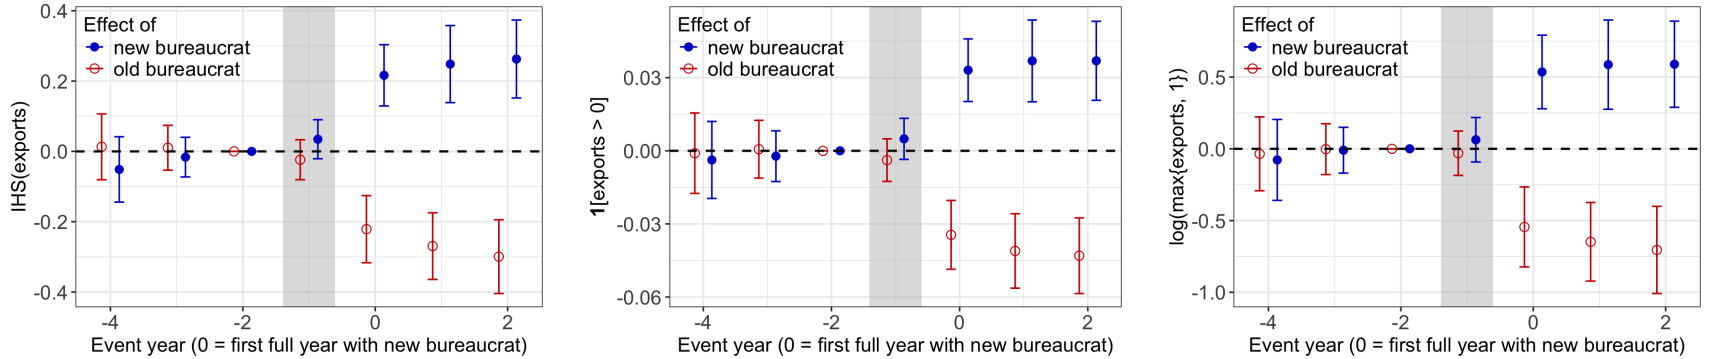

(a) Benchmark result, IHS(exports), combining extensive and intensive margin

(b) Extensive margin

(c) Intensive margin, only using variation $> \$1,000$

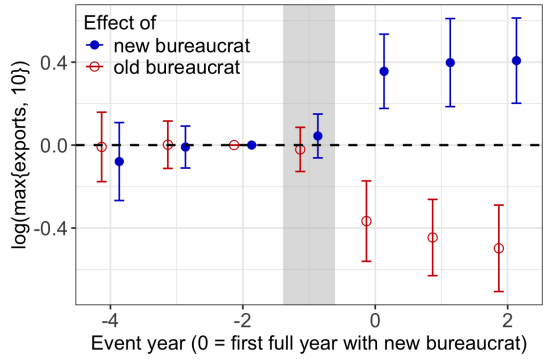

(d) Intensive margin, only using variation $> \$10,000$

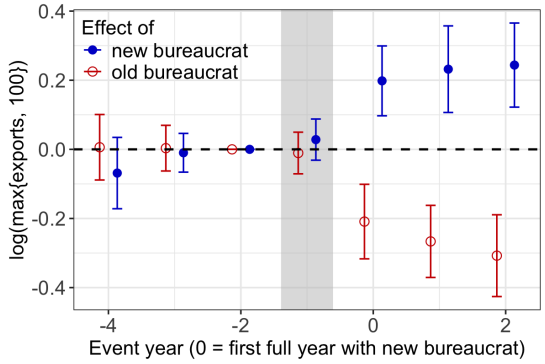

(e) Intensive margin, only using variation $> \$100,000$

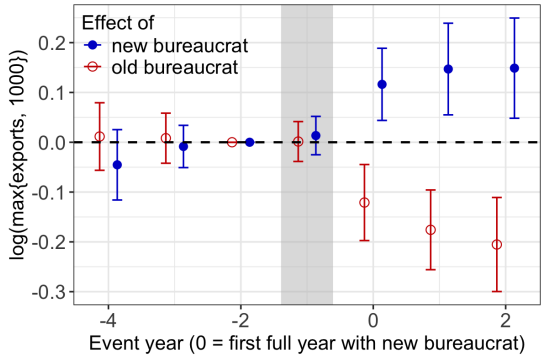

(f) Intensive margin, only using variation $> \$1,000,000$

Notes: The figure shows the effect of the incoming and outgoing bureaucrats' fixed effects on exports, around the time that the bureaucrat managing a country office changes. Specifically, we plot $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (4). In Panel (a), the outcome variable is the inverse hyperbolic sine (IHS) of exports to the country of the bureaucrat switch. Panel (b) investigates effects on the extensive margin: whether a product is exported to the destination. Panels (c)-(f) investigate the intensive margin by using $log(\max\{\text{exports}, m\})$, where m is a censoring threshold, as the outcome. Panels (c)-(f) use $m =$ \$1,000, \$10,000, \$100,000, and \$1,000,000, respectively (throughout the paper, exports are measured in thousands of contemporaneous U.S. dollars). For each value of m, the estimated effects only reflect changes in export values above m. We exclude country-product combinations without variation—i.e., without exports above m—over the event horizon to reduce mechanical attenuation from increasing m. Panels (b)-(f) suggests that both the extensive and intensive margins contribute to the overall effect in Panel (a). The bureaucrat fixed effects are estimated separately from equation (2). The switch occurs in year -1 so year 0 is the first full year that the new bureaucrat manages the country office. The omitted year is -2, the last full year with the old bureaucrat. Since equation (4) includes generated regressors, we report bootstrapped standard errors based on 500 bootstrap iterations, which are found to be more conservative than standard errors clustered by country (unreported). Back to Section 4.4.

Figure 4: Bureaucrat Fixed Effects Predict Exports Out of Sample

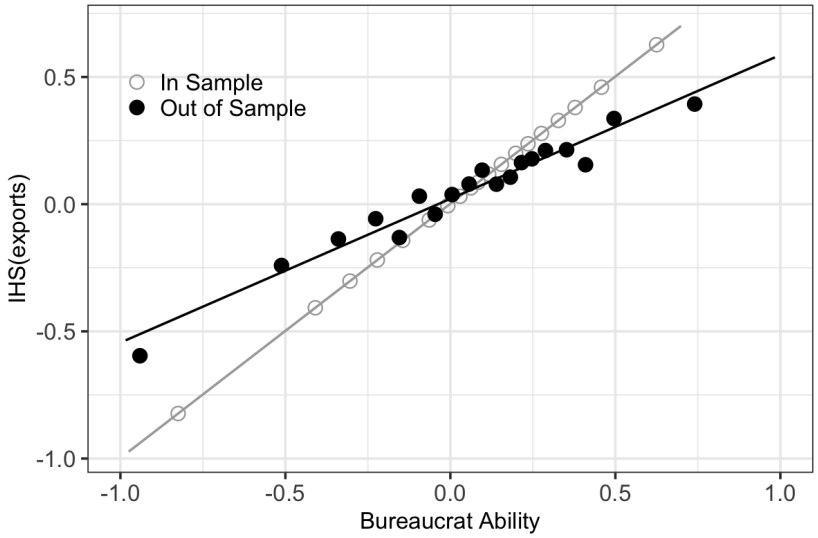

(a) Exports are greater under bureaucrats with high out-of-sample fixed effects.

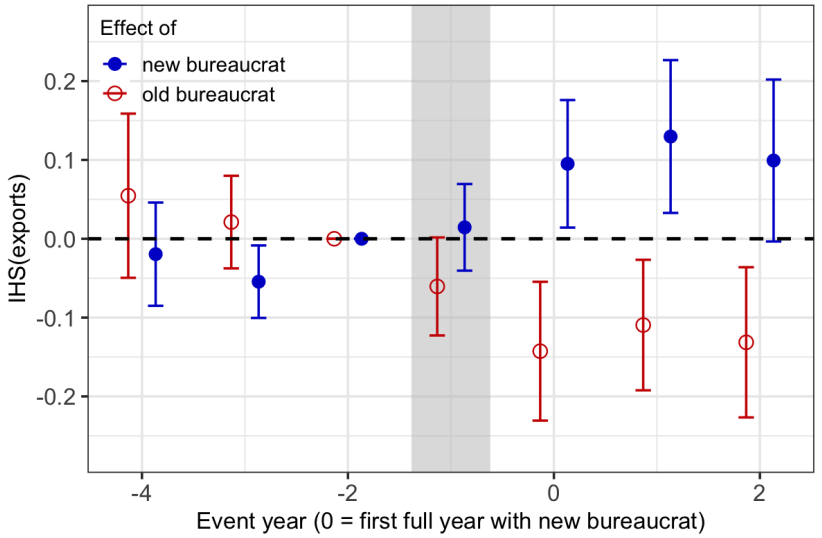

(b) Exports respond to out-of-sample fixed effects of new (incoming) and old (outgoing) bureaucrats.

Notes: Panel (a) displays a binned scatterplot of residualized exports on in-sample (i.e., baseline) and out-of-sample bureaucrat fixed effects. On the vertical axis, we plot exports, i.e., the inverse hyperbolic sine of export value, residualized for product-year fixed effects ($\hat{\lambda}_{pt}$) and country fixed effects ($\hat{\gamma}_c$). $\hat{\lambda}_{pt}$, $\hat{\gamma}_c$, and the in-sample bureaucrat fixed effects are all estimated from equation (2) using data from all country-years. Hence, by construction, each in-sample dot lies on the 45-degree line: In-sample fixed effects translate one-to-one into higher exports. Out-of-sample fixed effects, on the other hand, are estimated only using other countries. For instance, to predict exports to the UK in 1982, we obtain the fixed effect of the bureaucrat in charge of the UK office in 1982 using data from all country-years except the UK. The slope coefficient of the regression of residualized exports on these out-of-sample, i.e., *other country*, fixed effects is 0.51. Panel (b) uses these out-of-sample bureaucrat fixed effects to replicate Panel (a) of Figure 3; it shows the effect of the incoming and outgoing bureaucrats' out-of-sample fixed effects on exports to the country experiencing the switch. Because out-of-sample fixed effects are not available for single-appointment bureaucrats, we estimate two complementary versions of equation (4) to maximize power. One specification uses out-of-sample fixed effects for outgoing bureaucrats and in-sample fixed effects for incoming bureaucrats; the other reverses this assignment. Panel (b) then plots the out-of-sample coefficients, the objects of interest, from each specification. The corresponding in-sample coefficients, though unreported, are nearly symmetric to the out-of-sample ones. Standard errors are clustered by country. Back to Section 4.4.

Figure 5: Export Responses are Mediated by Product Characteristics, Market Conditions, and National Policy

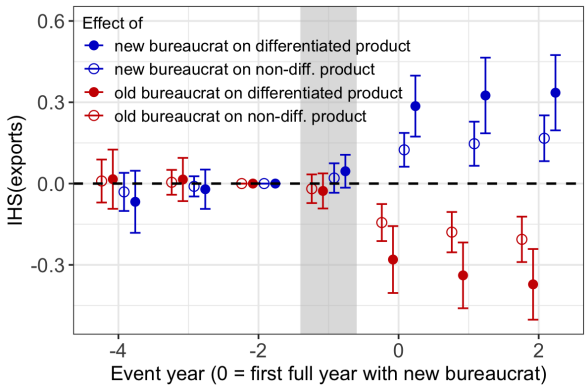

(a) Stronger effects in differentiated products

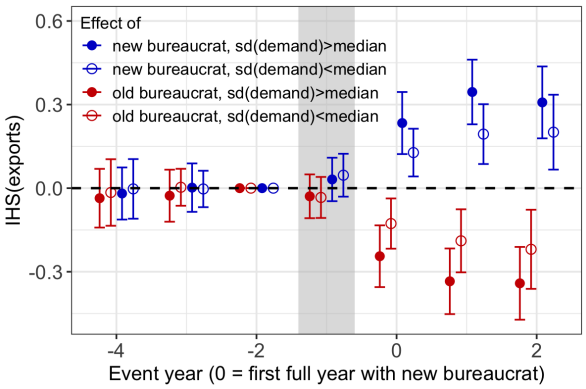

(b) Stronger effects in products with greater demand uncertainty

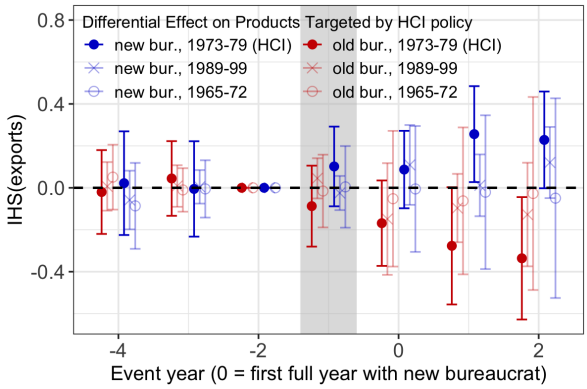

(c) Stronger effects in HCI products only during the national HCI drive

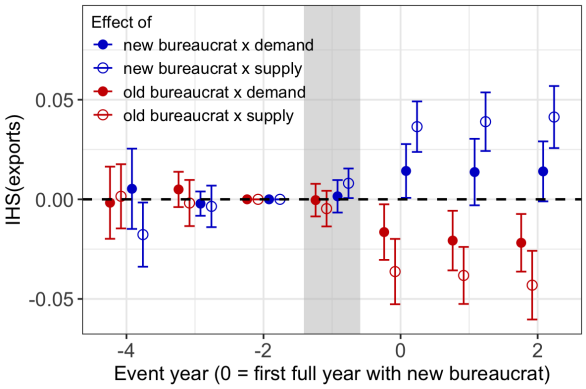

(d) Stronger effects in products with improving market conditions (increasing demand or supply)

Notes: The figure shows how the effect of bureaucrat ability on exports varies by product characteristics and market conditions. Panels (a) and (b) plot $\hat{\beta}_k^g$ and $\hat{\delta}_k^g$ for product groups $g \in \{1, 2\}$ from estimating equation (5). Panel (a) plots the effect of bureaucrat ability separately for differentiated and non-differentiated products, defined following Rauch (1999); the effects are stronger for differentiated products, indicating that higher-ability bureaucrats disproportionately raise exports of differentiated goods. Panel (b) plots the effect of bureaucrat ability separately for products with high and low demand uncertainty—defined as above- or below-median variance in import demand prior to the bureaucrat switch. The effects are stronger for products with higher demand uncertainty. Panel (c) plots $\hat{\beta}_k^1$ and $\hat{\delta}_k^1$ estimated from equation (5), where product group 1 consists of heavy and chemical industry (HCI) products and product group 2 includes all products. Estimates are reported separately for events occurring before (1965–1972), during (1973–1979), and after (1989–1999) the HCI drive. The differential effect of bureaucrat ability on HCI products is present only during the HCI policy period. Panel (d) plots $\hat{\beta}_k^D$, $\hat{\delta}_k^D$, $\hat{\beta}_k^S$, and $\hat{\delta}_k^S$ from estimating equation (6). These are the effect of bureaucrat ability interacted with destination-country import demand (measured by imports of the product from countries other than Korea) and Korean export supply (measured by Korea's exports of the product to other destinations). The effects are stronger when market conditions are favorable, indicating that higher-ability bureaucrats translate positive demand and supply conditions into exports more effectively. For all panels, the bureaucrat fixed effects are estimated separately from equation (2). Since the specifications include generated regressors, we report bootstrapped standard errors based on 500 bootstrap iterations. Back to Section 4.5.

Figure 6: Bureaucrats are Screened Based on Performance in 1st Appointments

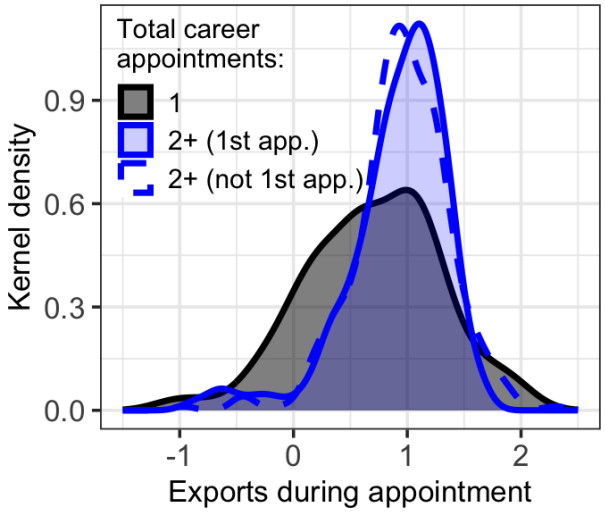

(a) Residualized exports by appointment

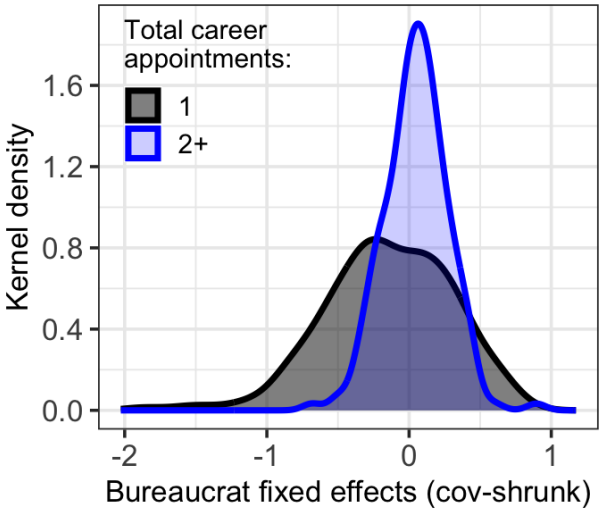

(b) Bureaucrat fixed effects

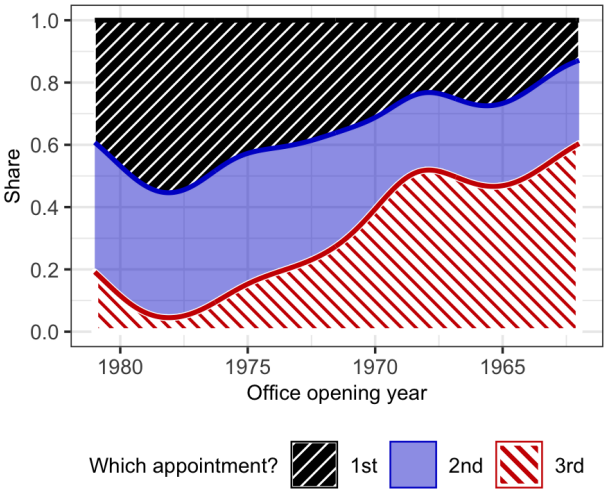

(c) Shares of 1st, 2nd, 3rd appointments by office opening year (in reverse order)

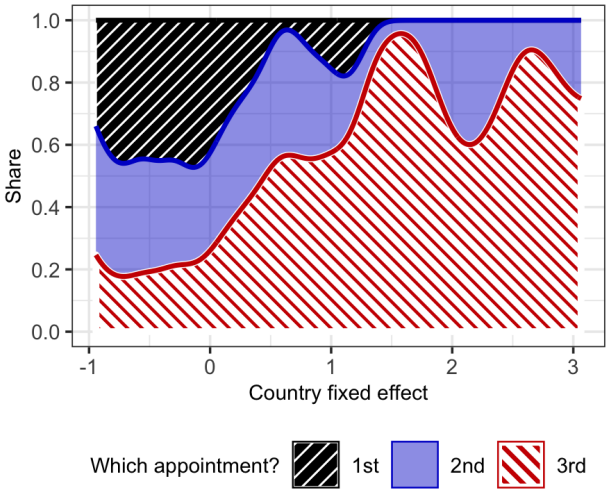

(d) Shares of 1st, 2nd, 3rd appointments by country fixed effect

Notes: Panels (a) and (b) show the kernel density of (a) residualized exports during a bureaucrat's appointment and (b) bureaucrat fixed effects (covariance-shrunk), separately for bureaucrats with one appointment (black) and for those with two or more appointments (blue). The distributions for single-appointment bureaucrats have much fatter left tails. Panel (a) also shows that the distribution of residualized exports does not change much between the first appointment and later appointments for multi-appointment bureaucrats. Panels (c) and (d) show the (smoothed) shares of appointments by whether it is the 1st, 2nd, or 3rd appointment of a bureaucrat, and by two different measures of country importance. Panel (c) depicts shares by office opening year; Panel (d) by country fixed effect value. Countries increase in importance from left to right, reflected in earlier openings and larger fixed effects. We consider offices opened during the main rollout period (1962-1981). To avoid distortion driven by limited office availability, we only include bureaucrats whose first appointments started after the rollout period (1981 or later). More important countries tend to be filled by bureaucrats in their 2nd, and especially 3rd, appointments. Back to Sections 2.2 and 4.6.

Figure 7: Opening an Export Promotion Office Increases Korean Exports,
but Does Not Coincide with Increases in Import Demand

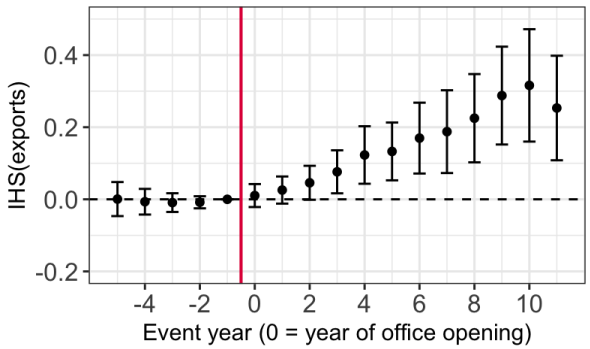

(a) Korean exports. Sample includes openings from 1967 to 1981

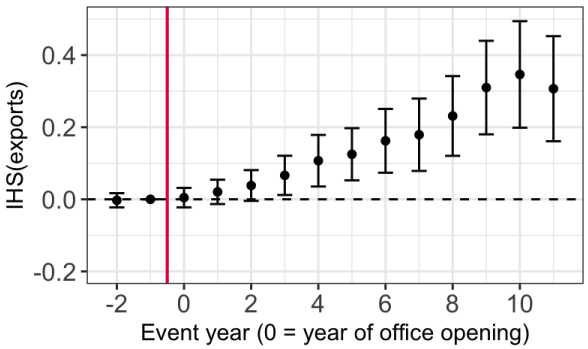

(b) Korean exports. Sample includes openings from 1964 to 1981

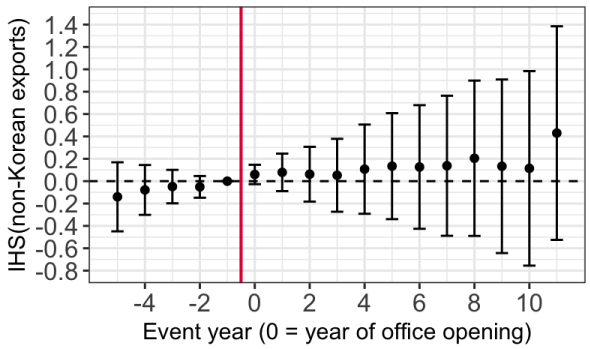

(c) Non-Korean exports. Sample includes openings from 1967 to 1981

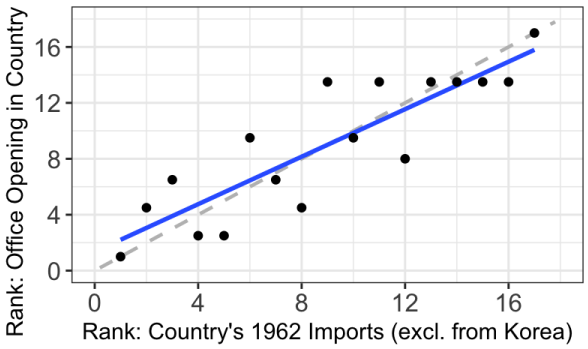

(d) Order of office openings in Europe follows pre-determined market size

Notes: Panels (a)-(c) plot event-study coefficients ($\hat{\pi}_k$) of the effect of opening export promotion offices on the inverse hyperbolic sine (IHS) of Korean/non-Korean exports, obtained from estimating equation (7) with a never-treated control group. Standard errors are clustered at the country level. Panel (d) is a scatterplot of the rank of office opening year and the rank of pre-determined market size (1962 total imports from all countries but Korea), for each European country that received a KOTRA office during the main rollout (1962-1981). The solid blue line plots the linear fit, whereas the dashed gray line is the 45-degree line. The rank correlation between 1962 imports and office opening is 0.87. When multiple countries have the same opening year, we assign the average rank to them. Back to Sections 5.2 and 5.3.

Table 1: Appointment Descriptives

	Full Sample	Connected Set	Leave-One-Out Connected Set
	(1)	(2)	(3)
# Bureaucrats	398	397	380
# Countries/Offices	87	86	75
# Appointments	705	704	676
# Bureaucrats ≥ 2 Offices	184	184	180
# Offices ≥ 2 Bureaucrats	83	83	75
# Offices ≥ 4 Bureaucrats	75	75	72
# Offices ≥ 6 Bureaucrats	62	62	61
# Offices ≥ 8 Bureaucrats	49	49	49

Notes: The table reports the sample characteristics for KOTRA's overseas offices and their managing bureaucrats. Column 1 reports these for the full sample of KOTRA's country offices. Column 2 restricts the sample to the countries and bureaucrats that form the largest connected set, while column 3 includes only those in the largest leave-one-out connected set, i.e., the set of countries and bureaucrats that would remain connected if any one appointment is removed. "# Bureaucrats" indicates the number of distinct bureaucrats. "# Countries/Offices" indicates the number of distinct offices, which is the same as the number of distinct countries. "# Appointments" indicates the number of distinct bureaucrat-country pairs. "# Bureaucrats ≥ 2 Offices" indicates the number of bureaucrats that managed at least two offices in the sample period. "# Offices $\geq x$ Bureaucrats" indicates the number of offices managed by more than x bureaucrats over the sample period. Back to Sections 3.1 and 4.1.

Table 2: Variance Decomposition of Exports

		Baseline		Robustness check: Bootstrap bias correction		Placebo check: Bureaucrats randomly reshuffled to countries		Relevance check: Ambassadors	
		Component	% Share	Component	% Share	Component	% Share	Component	% Share
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Var(exports|pt), spell-level		0.732	100	0.732	100	0.737	100	0.781	100
		(0.071)		(0.067)		(0.077)		(0.095)	
Var(bureaucrat)		0.101	13.8	0.100	13.7	0.007	0.9	0.033	4.2
		(0.051)		(0.029)		(0.026)		(0.056)	
Var(country)		0.721	98.6	0.682	93.2	0.590	80.1	0.615	78.8
		(0.102)		(0.097)		(0.073)		(0.139)	
Cov(bureaucrat, country)		-0.088	-12.0	-0.044	-6.0	-0.005	-0.6	-0.015	-1.9
		(0.047)		(0.039)		(0.014)		(0.051)	
Var(bureaucrat+country)		0.646	88.3	0.693	94.7	0.588	79.8	0.617	79.0
		(0.078)		(0.066)		(0.075)		(0.104)	
Var(exports|pt), raw		4.404		4.404		4.362		4.552	
Number of observations		1703465		1703465		1753145.1		1616256	
Number of bureaucrats		380		380		387.9		387	
by no. of spells in sample:	1	200		200		207.8		195	
	2	96		96		98.6		105	
	3	56		56		54.4		69	
	4	24		24		21.4		17	
	5	4		4		5.7		1	
Number of countries		75		75		78.3		74	

Notes: The table shows the results of variance decomposition according to equation (3). Columns 1-2 contain our baseline results. We correct for limited mobility bias following Kline et al. (2020), implemented via the algorithm of Bonhomme et al. (2023). Columns 3-4 use the bootstrap bias correction method, described in the main text, instead. Columns 5-6 use synthetic data where bureaucrats are randomly reshuffled to countries, preserving the original appointments—i.e., country-years that form one appointment—as well as the number of appointment spells for each bureaucrat. Columns 7-8 include fixed effects for Korean ambassadors instead of KOTRA office managers. The Kline et al. (2020) bias correction method is applied to columns 5-8 as well. Because the algorithm relies on numerical approximations of traces of large matrix inverses, the decomposition results contain a small degree of approximation noise. Columns 5–6 introduce additional randomness due to the random reshuffling of bureaucrats. To reduce approximation noise, all reported variances and covariances are averages over 100 iterations of the algorithm, holding the original sample fixed. The reported numbers of observations, bureaucrats, and countries are likewise averaged over 100 iterations. In columns 5-6, they are decimal numbers because each random reshuffling produces a different leave-one-out connected set; their similarity to the ones in columns 1-2 implies that the connectivity of the graph is largely maintained in the random reshuffling. The parentheses contain bootstrap standard errors of the variance components, generated by resampling the data at the appointment level 100 times. For each bootstrap sample, we again average over 100 iterations of the algorithm, such that 10,000 iterations are used in total to construct the standard errors. Back to Sections 4.3 and 4.4.

Table 3: Correlates of Bureaucrat Fixed Effects

	Bureaucrat Fixed Effects (std.)								
	Inquiries by foreign firms		Reports on market conditions				Trade fairs		
	All	If any inquiry	All	Before HCI	During HCI	After HCI	All	If any fair	If any fair (years with reported sales)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Total # of inquiries (std.)	0.164***	0.154**							
	(0.061)	(0.061)							
Total # of unique SITC codes (std.)	0.086*	0.103*							
	(0.050)	(0.052)							
Avg # of products in an inquiry		0.016							
		(0.018)							
Avg char. len. per product in an inquiry		-0.0078**							
		(0.0034)							
Total # of general reports (std.)			0.148**	0.102	-0.159	0.005			
			(0.065)	(0.106)	(0.120)	(0.011)			
Total # of product-specific reports (std.)			-0.308***	-0.102	-0.214***	-0.192			
			(0.060)	(0.071)	(0.060)	(0.152)			
Total # of HCI reports (std.)				-0.005	0.192***	0.061			
				(0.273)	(0.061)	(0.113)			
Number of fairs							0.174***	0.221***	
							(0.043)	(0.044)	
Avg number of firms per fair								0.00064	
								(0.00691)	
Log of total sales in fairs									0.067*
									(0.036)
N	1623	1566	2091	218	400	760	2091	468	203
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Country FE	✓	✓	✓	✓	✓	✓	✓	✓	✓

Notes: This table shows correlations between bureaucrat fixed effects estimated from equation (2) and measures of activity produced by the overseas office managed by each bureaucrat. Observations are at the country-year level. All specifications include country and year fixed effects to control for differences in market size across countries and for changes over time in how KOTRA operated. Variables are standardized where indicated. Standard errors are clustered at the bureaucrat level and reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Back to Section 4.5.

Table 4: Counterfactuals

Counterfactual		**Change in exports** compared to status quo (where screening occurs after 1st appointment)			
		Decomposition		Total Effect	Total Effect
		Appointments with single-appointment bureaucrats	Appointments with multi-appointment bureaucrats	All appointments in leave-one-out connected set	All appointments ... using covariance-shrunk bureaucrat effects
(1)	No screening occurs after 1st appointment	0.106 (11.2%)	-0.139 (-13.0%)	-0.067 (-6.5%)	-0.052 (-5.1%)
(2)	Screening occurs before 1st appointment	0.223 (25.0%)	0	0.066 (6.8%)	0.054 (5.6%)
(3)	Replace bottom quartile with median	-	-	-	0.080 (8.3%)
	Number of bureaucrats	200	180	380	380
	Number of appointments	200	476	676	676

Notes: The table reports changes in the inverse hyperbolic sine (IHS) of exports and the corresponding percentage change in exports (approximated using $(e^{\Delta \text{IHS(exports)}} - 1) \times 100$) under different counterfactual scenarios. We focus on the differential selectivity of single- and multi-appointment bureaucrats in the data, which we label as "screening." Counterfactual (1) sets the average ability of both groups of bureaucrats to equal the average ability across *all* bureaucrats, i.e., raising it for single-appointment bureaucrats and lowering it for multi-appointment ones. Counterfactual (2) sets the average ability of single-appointment bureaucrats to be as high as that of multi-appointment bureaucrats. Counterfactual (3) sets the ability of bureaucrats in the bottom quartile to equal the median ability among all bureaucrats. The total effect is a weighted average of the change in exports for appointments with single-appointment and multi-appointment bureaucrats, weighted by their respective number of appointments. Counterfactuals (1) and (2) are not subject to limited mobility bias as they only compare average ability between two groups, each estimated from a large number of observations. Counterfactual (3), however, is based on individual bureaucrat fixed effects, so we apply covariance shrinkage to adjust for limited mobility bias. The other counterfactuals, too, are repeated using the covariance-shrunk fixed effects, for comparability with counterfactual (3). Back to Section 4.6.

Online Appendix

A Appendix Figures

Figure A.1: Appointment Durations and Time Between Appointments

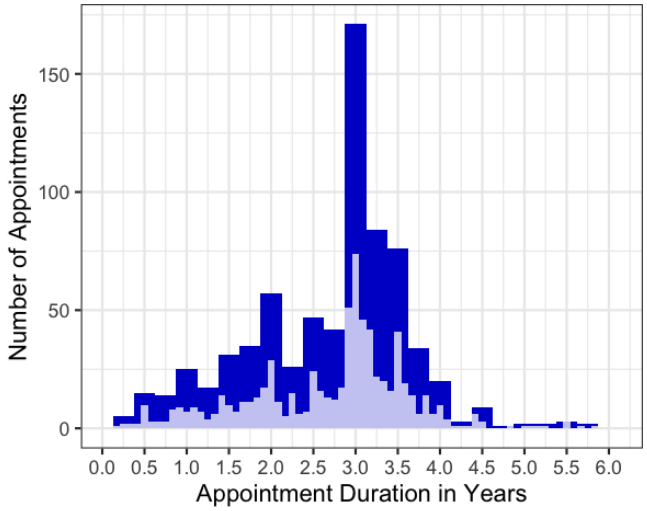

(a) Distribution of Appointment Durations. Median and Modal Duration: 36 Months.

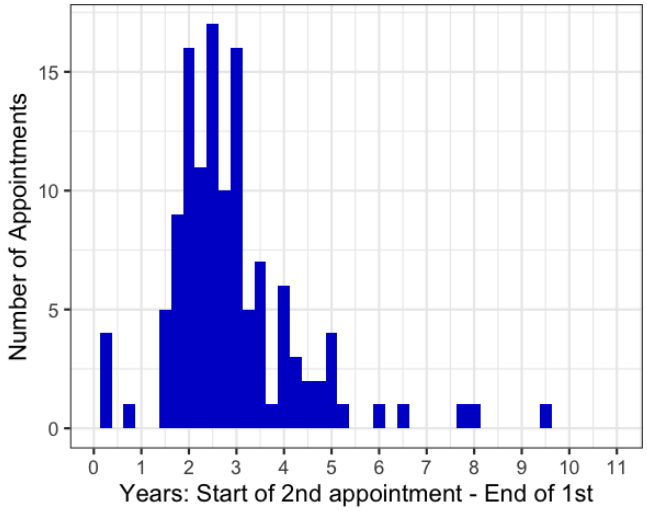

(b) Distribution of Time Between Appointments. Median: 29 Months. Mode: 30 Months.

Notes: Panel (a) shows the distribution of appointment durations. The blue (light blue) bars indicate the number of appointments by quarterly (monthly) duration. As each quarter contains multiple months, the blue bars always weakly exceed the light blue ones. Panel (b) shows the distribution of gaps between appointments, measured in quarters. Back to Sections 2.3 and 4.1.

Figure A.2: Targeting of Export Promotion Activity by Product.
Export Promotion Activity Moves in Parallel with Heavy Chemical Industry Drive

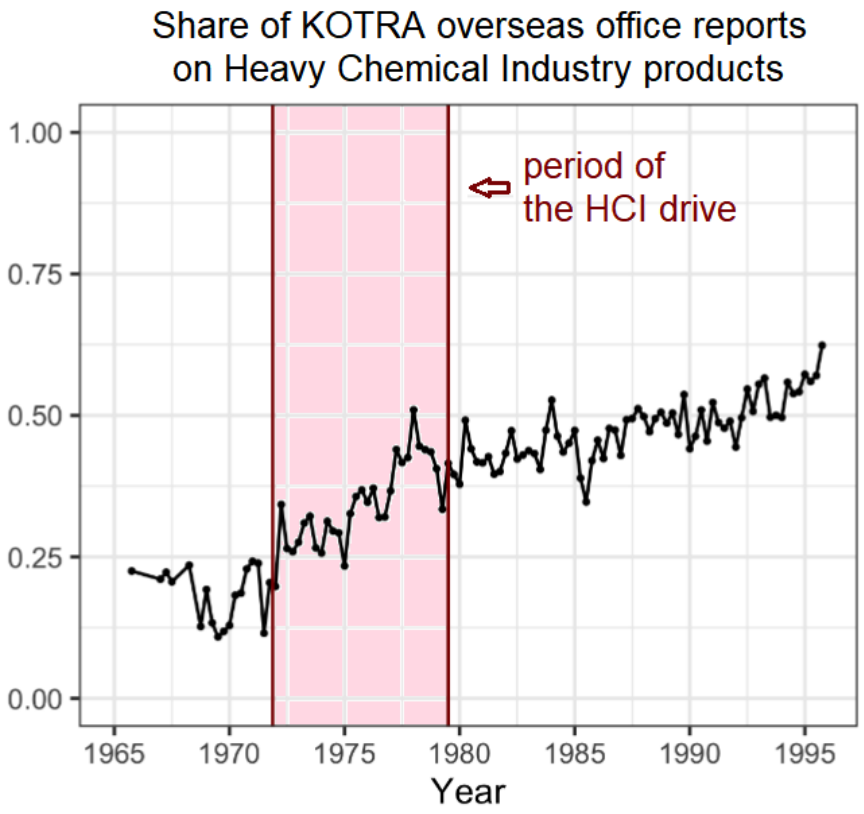

Notes: Targeting of export promotion activity by product. For each quarter, the vertical axis presents the share of overseas office reports that could be linked to the target products of the Heavy Chemical Industry (HCI) drive, relative to the number of reports that could be linked to any product. Back to Section 2.4.

Figure A.3: Demonstration of Limited Mobility Bias: Variances and Covariances of Sampling Errors Converge to Zero As Mobility Increases

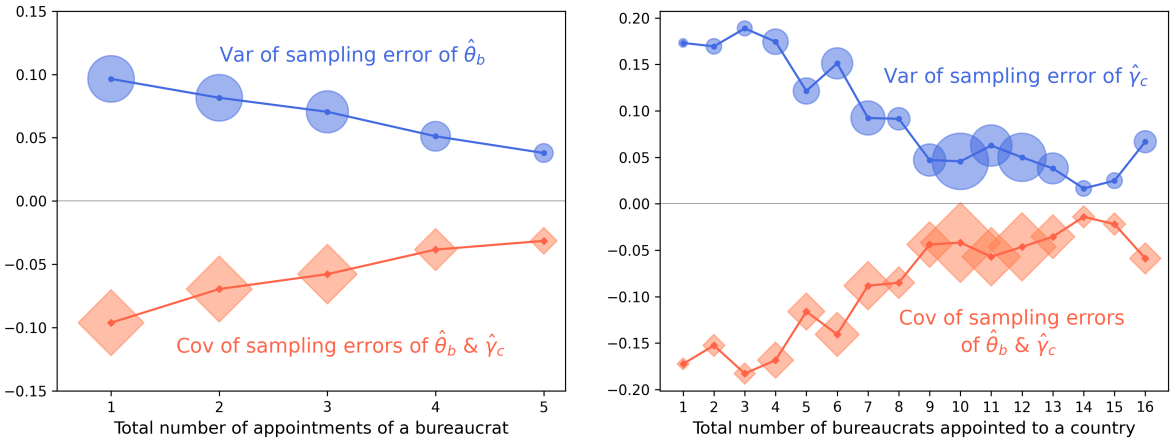

(a) Sampling errors of bureaucrat fixed effects

(b) Sampling errors of country fixed effects

Notes: This figure shows that the variances and covariances of sampling errors of bureaucrat and country fixed effects gradually tend to zero, as the size of the sample most relevant for their identification increases. Although all bureaucrat and country fixed effects are jointly estimated from the rotation of bureaucrats over the entire connected set of bureaucrats and countries, the observations most relevant in estimating $\hat{\theta}_b$ are the countries bureaucrat b is directly connected to (same for $\hat{\gamma}_c$). To estimate the variance and covariance of the sampling errors, we first bootstrap equation (2). Then, we calculate the variance of $\hat{\theta}_b$ across bootstrap samples for each bureaucrat b, the variance of $\hat{\gamma}_c$ across bootstrap samples for each country c, and the covariance of $\hat{\theta}_b$, $\hat{\gamma}_c$ across bootstrap samples for each (b, c) pair (see Appendix F for details). Panel (a) plots the averages of the bootstrap estimates of the variance of $\hat{\theta}_b$ and the covariance by the total number of appointments of b over the sample period. Panel (b) plots the averages of the bootstrap estimates of the variance of $\hat{\gamma}_c$ and the covariance by the total number of bureaucrats appointed to c over the sample period. Sampling error—and therefore the spurious negative correlation between the fixed effects of matched pairs (see Andrews et al. (2008))—is the worst for bureaucrats (countries) who are directly connected to a very small number of countries (bureaucrats). The divergence from zero for countries with 16 bureaucrats is explained by the fact that they are also Korea's most important trade partners, with exceptionally large values of $\hat{\gamma}_c$. Back to Sections 4.2.

Figure A.4: Distribution of Bureacrat Fixed Effects, Raw vs. Covariance-Shrunk

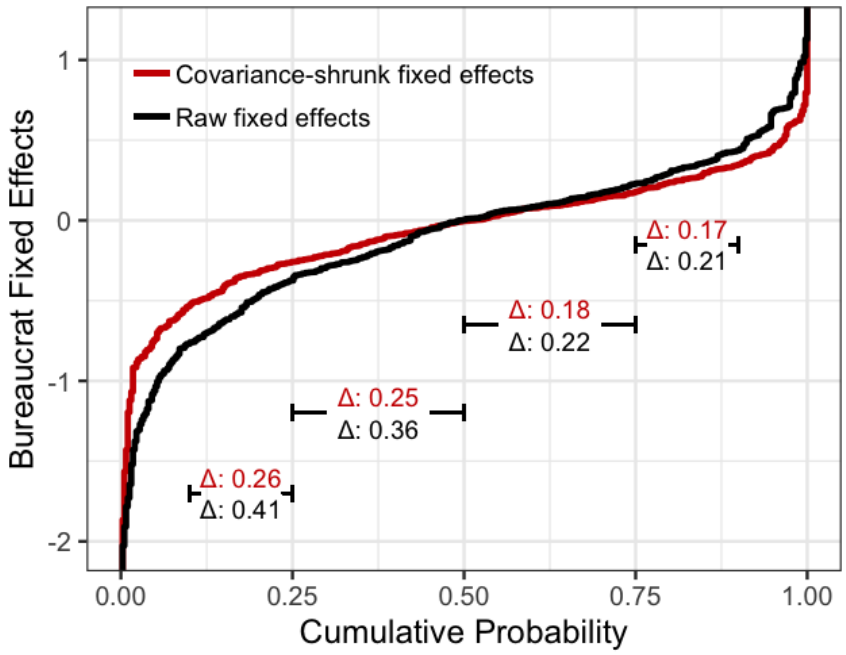

Notes: The figure shows the cumulative distribution function of the raw bureaucrat fixed effects estimated from equation (3), as well as covariance-shrunk bureaucrat fixed effects. Moving from the 25th to the 50th percentile would imply an increase in the inverse hyperbolic sine (IHS) of exports of 0.25 based on the covariance-shrunk fixed effects, compared to 0.36 under the raw fixed effects. Moving from the 50th to the 75th percentile would imply an increase of 0.18 under covariance-shrinkage, compared to 0.22 under the raw fixed effects. Back to Section 4.3.

Figure A.5: Robustness of Office Opening Results to Controls

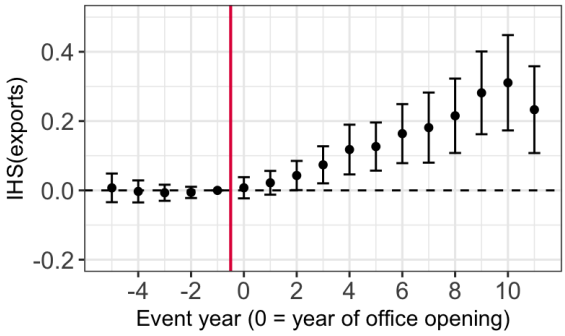

(a) Control for non-Korean exports. Sample includes openings from 1967 to 1981

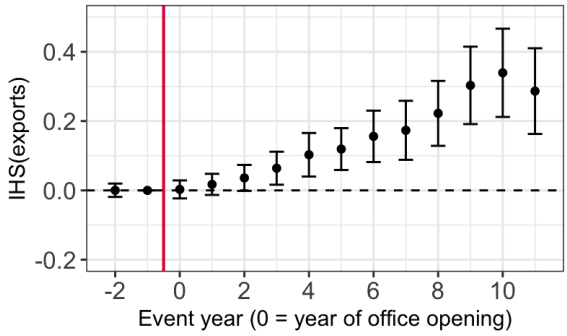

(b) Control for non-Korean exports. Sample includes openings from 1964 to 1981

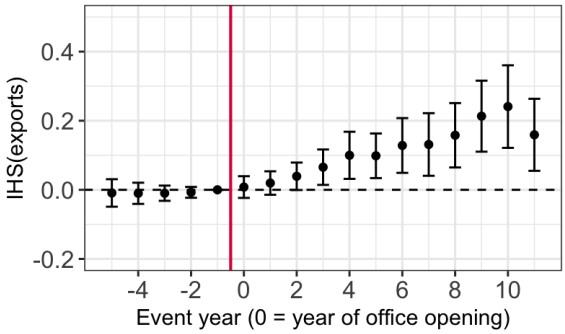

(c) Control for non-Korean exports $\times$ year dummies. Openings from 1967 to 1981

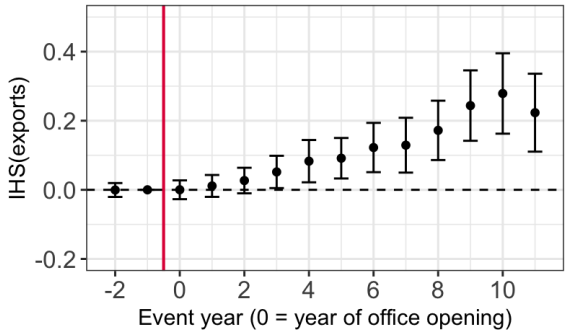

(d) Control for non-Korean exports $\times$ year dummies. Openings from 1964 to 1981

Notes: The figure reports results from estimating equation (7), with a never-treated control group. For Panels (a)-(d), the outcome variable is the inverse hyperbolic sine of Korean exports to the country of the office opening. Panels (a)-(d) perform robustness checks on the results in Panels (a) and (b) of Figure 7 by controlling for non-Korean exports or non-Korean exports interacted with year dummies. Standard errors are clustered by country. Back to Section 5.3.

Figure A.6: Opening an Office Increases KOTRA Activity on That Country

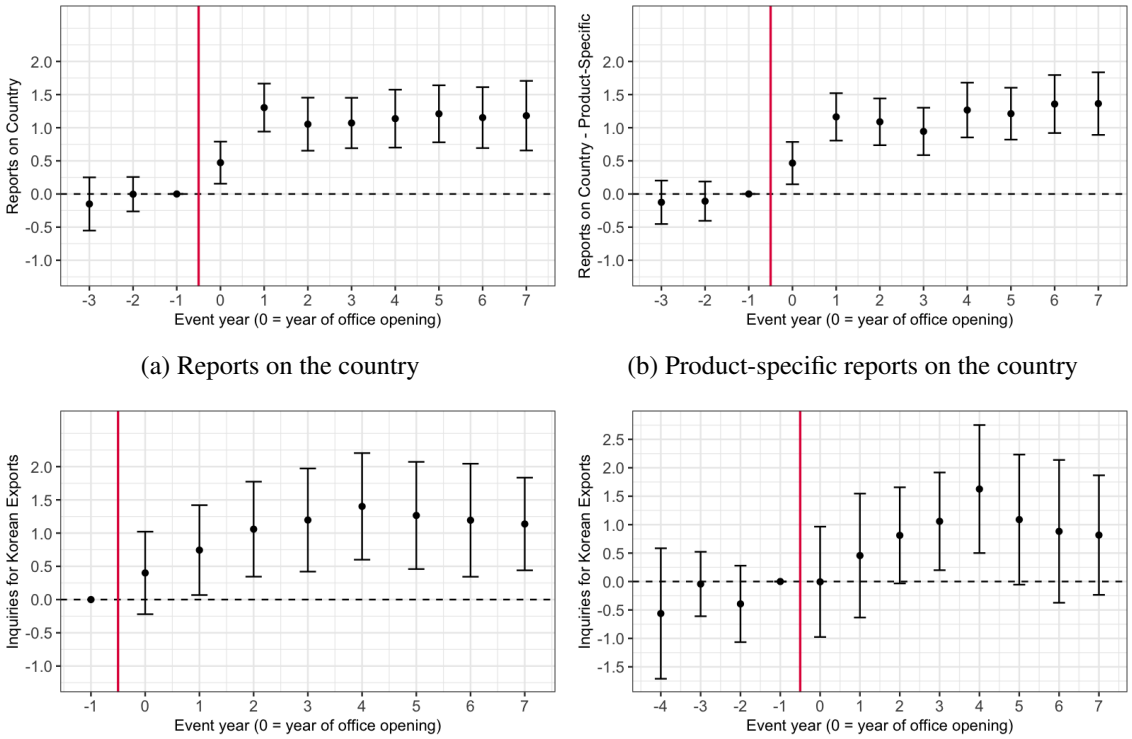

(a) Reports on the country

(b) Product-specific reports on the country

(c) Inquiries from potential importers

(d) Inquiries from potential importers. Openings from 1978

Notes: This figure reports results from estimating equation (7), with a never-treated control group, for different outcome variables measuring KOTRA activity on a specific country. These are (1) the number of reports about a country, (2) the number of product-specific reports, and (3) the number of inquiries from potential importers in that country interested in trade with Korea, each transformed by the inverse hyperbolic sine (IHS). The data on reports cover the years 1965–2001, so Panels (a) and (b) exclude offices that opened before 1968. The data on inquiries cover 1974–1997, so Panel (c) excludes offices that opened before 1974. Including openings from 1975 allows only one pre-period in Panel (c). Panel (d) instead has multiple pre-periods, at the cost of restricting the treatment group to countries with openings between 1978 and 1981. Back to Section 5.2.

B Appendix Tables

Table B.1: Bureaucrat Fixed Effects Contribute to Variation in Exports and Are Stable Across Appointments

	Share of variation in IHS(exports) explained by various FE			
	(1)	(2)	(3)	(4)
Adjusted 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

Notes: This table reports the R^2 from regressing the inverse hyperbolic sine (IHS) of Korean exports on different combinations of fixed effects. We progressively add fixed effects as we move one column to the right. Column 3 corresponds to estimating equation (1). Column 4 adds bureaucrat-country fixed effects, which subsumes bureaucrat fixed effects as well as 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. Back to Sections 4.2 and 4.4, and Appendix D.1.

C Data

C.1 Bureaucrat Appointments

The most comprehensive source of bureaucrat appointments consists of contemporaneous reports on appointments to KOTRA's overseas offices published in major South Korean newspapers. In most years, appointments were announced on two main dates, typically in January and July, with postings most often beginning in April and October. The information was usually reported in three major newspapers (*Dong-A Ilbo*, *Chosun Ilbo*, and *Khughyang Shinmun*). Because these sources overlap, almost no round of announcements went unreported. For nearly all rounds of appointments, the information can be corroborated using at least two of the these sources.

We complement the newspaper announcements with various publications that document the bureaucrat managing each office at a given point in time. We obtained and digitized the names of bureaucrats in (i) monthly KOTRA publications aimed at non-Korean importers (1966-1971), (ii) KOTRA's directory of its overseas network (1977, 1991-1994, 1998-2000), (iii) KOTRA's reports on trade fairs (1969, 1971-1997), and (iv) a full directory of all overseas office managers using the Korean Business Directory (1987-1998) published by the Korean Chamber of Commerce and Industry.

Overall, we identify 138 offices in 87 distinct countries that operated between 1962 and 2001. We identify 475 unique bureaucrats and 974 unique appointments to offices. Of these, 398 bureaucrats were appointed to head country offices, resulting in 705 unique bureaucrat-country pairs. Table 1 provides additional descriptive statistics, detailing the remarkable connectedness of the country-bureaucrat graph, discussed in Section 4.

Bureaucrats are identified by their names, requiring us to avoid two types of errors. First, we may erroneously code two bureaucrats as the same individual, for example if different bureaucrats share a name. A priori, this issue could have been serious, as 45% of bureaucrats in our sample share the last names *Kim*, *Lee*, and *Park*. However, substantial diversity in first names largely mitigates this concern.[66] After extensive consistency checks, it appears highly unlikely that any two bureaucrats in the data share the exact same full name.

Second, and more challenging in practice, we must determine whether slightly different names truly refer to distinct individuals. This task is complicated by the fact that, over time, our sources shift from Chinese characters to Korean characters to render names. In addition,

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

in the few cases where names are romanized, inconsistent spellings are used—for example, both *rhee* and *lee* appear for the same Korean syllable. We address these challenges in four steps: (1) matching rare or unusual names to more common variants to detect likely digitization errors; (2) harmonizing alternative renderings of the same syllables (e.g., *rhee* and *lee*); (3) identifying offices with likely mistakes, e.g., the managing bureaucrat's name flips back and forth; and (4) reconstructing each bureaucrat's career to identify inconsistencies, overlaps, or missing years. Following these steps meticulously allowed us to construct a consistent panel of unique bureaucrats covering all offices and years.

C.2 Inquiries

Our inquiry data come from *Overseas Market News*, KOTRA's newspaper that was published almost daily and made publicly available to Korean firms. The publication contained tables reporting business inquiries from foreign companies seeking to trade or cooperate with Korean firms, along with procurement opportunities announced by foreign agencies. We conducted OCR text recognition on the scanned pages of the *Overseas Market News* available in the KOTRA archives and constructed a novel dataset of 169,280 inquiries from September 13, 1974 to November 26, 1997.

Each inquiry row lists the name, street address, and country of the foreign company or agency and a short description of the product or service sought. The tables are organized by inquiry type. The main types include i) *import inquiries*—foreign firms seeking to import from Korea; ii) *export inquiries*—foreign firms seeking to export to Korea; iii) *cooperation inquiries*—foreign firms seeking joint ventures or other forms of partnership with Korean firms; and iv) *tender notices*—information on public procurement projects for which Korean firms could submit bids. Among these categories, import inquiries by foreign firms are most directly related to Korea's exports and are by far the most numerous, accounting for more than 70% of all inquiries in the dataset. They also follow the most standardized and consistent reporting format. For these reasons, the analysis in Section 4.5 focuses on import inquiries. Appendix Figure C.7 contains an example import inquiry table.

Data Digitization and Extraction Methods

We employ two independent extraction procedures to identify and parse inquiry tables from the scanned pages of *Overseas Market News*, allowing us to cross-validate results and assess extraction accuracy.

The first method relies on conventional optical character recognition (OCR) applied to all scanned pages. The OCR output provides the recognized text along with the coordinate

locations of each text on the page. Using these coordinates and text, we detect sections on a page that match the structure of an inquiry table. The contents of each identified table region are then parsed into structured rows and columns.

The second method employs a large-language-model (LLM)-based table extraction procedure. For every page identified as containing an inquiry table by the OCR method, we submit a table extraction request to the OpenAI API.

Then for each table, we deduce the inquiry type based on key terms in the column titles extracted by the two methods, and compare the inquiry types in the two methods. Where the inquiry type of a table is identified in both methods, the two methods agree 97% of the time. Discrepancies occur when there are errors in the extracted column titles due to unclear text in the scanned page. In these cases, we submit independent OpenAI extraction requests multiple times to reach a majority verdict on the inquiry type. For pages with tables that lack column titles as they are continuations of tables in previous pages, we assign them the inquiry type identified on the preceding pages.

Country Identification and Matching

Each inquiry row contains the country associated with the foreign firm or agency. We normalize these extracted country names and match them to the country names in our main dataset on the value of Korean exports. Matching is done via fuzzy string comparison, allowing for minor OCR errors, transliteration differences, and typographical noise. For unmatched cases, the entries are flagged for manual inspection. The resulting harmonized country names allow us to merge the inquiry data to our main dataset.

Product Classification

The product description column in each inquiry is less standardized than the country or company name and address columns, varying in total character length, the number of products listed, and the level of detail. To quantify this information, we first classify the listed products into SITC Revision 4 codes at the 2-digit level. We submit the full product description text from each inquiry row to the OpenAI API, which returns a mapping between each product and its most likely two-digit SITC code.[67] We performed manual verification on a small subsample to assess accuracy. This output enable aggregation of inquiries by product category as well as measuring the breadth of the different SITC code coverage.

[67] As an example, the product–SITC code mapping for the first inquiry row shown in Appendix Figure C.7 is: {"sphygmomanometers": 87, "stethoscopes, diagnostic set": 87}. SITC code 87 stands for "Professional, scientific and controlling instruments and apparatus."

Data Quality and Limitations

While every effort was made to ensure accuracy, the data are subject to several limitations inherent in historical document digitization. OCR errors in country names, product descriptions, and column titles occasionally occur, especially in pages with lower quality scans or where the original printing quality happened to be less good. These introduce misclassifications of foreign country, SITC product code, or inquiry table type. However, we expect such errors to occur more or less randomly across pages.

Figure C.7: Example Import Inquiry Table

國 名	商 社 名 및 住 所	對 韓 輸 入 希 望 品 目	問議番號
Australia	Healthjoy Distributors 56 Willsmere Road, Kew, Vic. 3101 Tlx : TF 8623308	Sphygmomanometers; Stethoscopes, Diagnostic set	4376-광-043 (MELD900085)
U.S.A.	Ancel Int'l Inc. 5000 Cedar Plaza Parkway Suite 304, St. Louis, MO 63128 Fax : 314-842 3222	Chair & parts	4376-광-044 (CHID900206)
Spain	SA Transgrup Castillejos, 408-412 08024 Barcelona Tel : 51246 Fax : 3- 347-4368	Sewing-machine accessories and replacement parts	4384-중-200 (MADD900149)

Notes: This figure displays a snippet of an import inquiry table from *Overseas Market News* dated March 9, 1990. The four column titles read, from left to right, country name; company name and address; products sought for import from Korea; and inquiry reference ID.

C.3 Market reports

In addition to inquiry tables, we digitize the topical content of *Overseas Market News* by extracting the Table of Contents (ToC) from each issue. The ToC lists the titles and subtitles of all reports and articles appearing in that issue.

Data Extraction Methods

We extract ToC information from the bookmark metadata (PDF outline) embedded in the scanned issue files, using the `pikepdf` library in Python. Each bookmark entry yields a title and, where present, a subtitle, along with the date of the issue. We compile these across all issues into a single corpus of 149,515 ToC entries from 9,602 issues spanning 1965 to 2001. Coverage in the 1960s is incomplete; we have 26 issues for 1965, none for 1966, 197 for 1967, 51 for 1968, and 190 for 1969. From 1970 onward, coverage is nearly complete,

averaging around 290 issues per year—consistent with daily publication on business days. Because the publication was intended for a Korean audience, report titles appear in Korean and Hanja (Chinese characters used in Korean writing).

Country and Product Classification

To classify reports as referring to a certain country or product, we separately tokenize the ToC corpus for Korean and Hanja scripts using the `quanteda` text analysis package in R. We remove standard stopwords for each script, a curated list of high-frequency but uninformative tokens (generic terms for imports, exports, trends, demand, and similar). From the cleaned tokens we construct frequency counts of unigrams, bigrams, and trigrams, retaining all tokens appearing at least ten times in the corpus. This yields 2,427 Hanja tokens and 2,356 Korean tokens (4,783 in total).

Research assistants review each token and, where applicable, code it as referring to a destination country or a product category (two-digit SITC codes). Of the 4,783 tokens, 325 are coded as country references (256 mapped to specific countries and the remainder to broader regions, such as the Middle East, Europe, or the Americas) and 615 are coded as product references. To assist with coding tokens, each token is presented alongside (i) a machine translation[68] and (ii) the six most frequent bigrams and trigrams in which the token appears. For example, the Hanja token 鐵 ("iron") is mapped to SITC code 67 (Iron and steel), while the Korean token 신발 ("shoes") is coded as SITC code 85 (Footwear). The remaining tokens correspond to generic economic vocabulary (e.g., "trend," "increase," "demand") or institutional terms. In total, about 20 percent of all high-frequency tokens receive a country or product classification.

We then use the token-level classification to code the countries and products to which each ToC title and subtitle refer. Products, as well as multi-character country names (e.g., 일본, 中國, 인도네시아) are matched directly via string search. Single-character Hanja abbreviations for countries—a convention common in Korean newspapers of this era (e.g., 美 for the United States, 日 for Japan, 英 for the United Kingdom)—are potentially ambiguous, as the same characters can appear as components of unrelated words. We accept a single-character country match only if the character appears at the start of the title or subtitle, or if it is immediately adjacent to trade-related terms such as 對 ("toward"), 輸出 ("export"), 輸入 ("import"), or 貿易 ("trade"). We also apply substring precedence rules to

[68] To translate tokens into English, we use Papago for Korean tokens. We use both Papago and Baidu for Hanja tokens, translating from Japanese and Chinese respectively, as standard translation software does not recognize Hanja appearing in Korean texts.

avoid double-counting; e.g., 인도 ("India") is not counted when it appears as part of 인도네시아 ("Indonesia"). Each ToC entry is assigned up to four country codes.

This procedure produces country-level and product-level measures of KOTRA's reporting intensity over time.

Data Quality and Limitations

The ToC-based measures are subject to several limitations. First, ToC extraction relies on bookmark metadata embedded in the PDF files; issues without PDFs or without the necessary metadata are not covered. From 1970 to 2000, there are 9,705 Monday-through-Saturday days, of which we extract ToC data for 8,995. Half of the gap of 710 days is accounted for by five months missing entirely from the archive (August, October, and December 1972; November 1976; March 1980) and by public holidays. Coverage in the 1960s is less complete, as noted above.

Second, report titles are necessarily a coarse summary of article content and may omit secondary products or countries discussed in the full text. Third, manual coding of tokens into countries and SITC categories involves judgment, especially for Hanja terms with multiple possible meanings. However, because coding is restricted to high-frequency tokens and performed by trained research assistants following a common protocol, we do not expect systematic misclassification.

C.4 Trade Fairs

We digitize reports on KOTRA-organized pavilions at international trade fairs from KOTRA's Annual Reports on International Trade Fairs.

Data Extraction Methods

Research assistants code each trade fair report, recording the list of Korean firms that attended the KOTRA pavilion and, where available, information on their sales at the fair (value sold in U.S. dollars, expected future sales, and buyer firm names) as well as the products exhibited. Appendix Figure C.8 shows an example of a table containing information on sales for each Korean firm.

The data cover 1,024 fairs across 153 locations in 25 years between 1969 and 1996 (the annual reports for 1970, 1974, and 1991 are not available in the archive). Trade fair activity expanded over the sample period, from 12 fairs with 51 participating firms in 1969 to 53 fairs with 856 firms by 1996. The most frequent fair locations are Tokyo (60 fairs), Paris (58), Chicago (46), New York (41), and Frankfurt (30).

Overall, there are more than 13,000 instances of Korean firms attending KOTRA pavilions. A typical fair has 12.1 participating Korean firms (interquartile range 7-14). Sales values in U.S. dollars are recorded for over 5,700 of these instances. The products a firm exhibited at a firm is recorded for more than 8,000 firm-fair pairs.

Figure C.8: Example Fair Sales Table

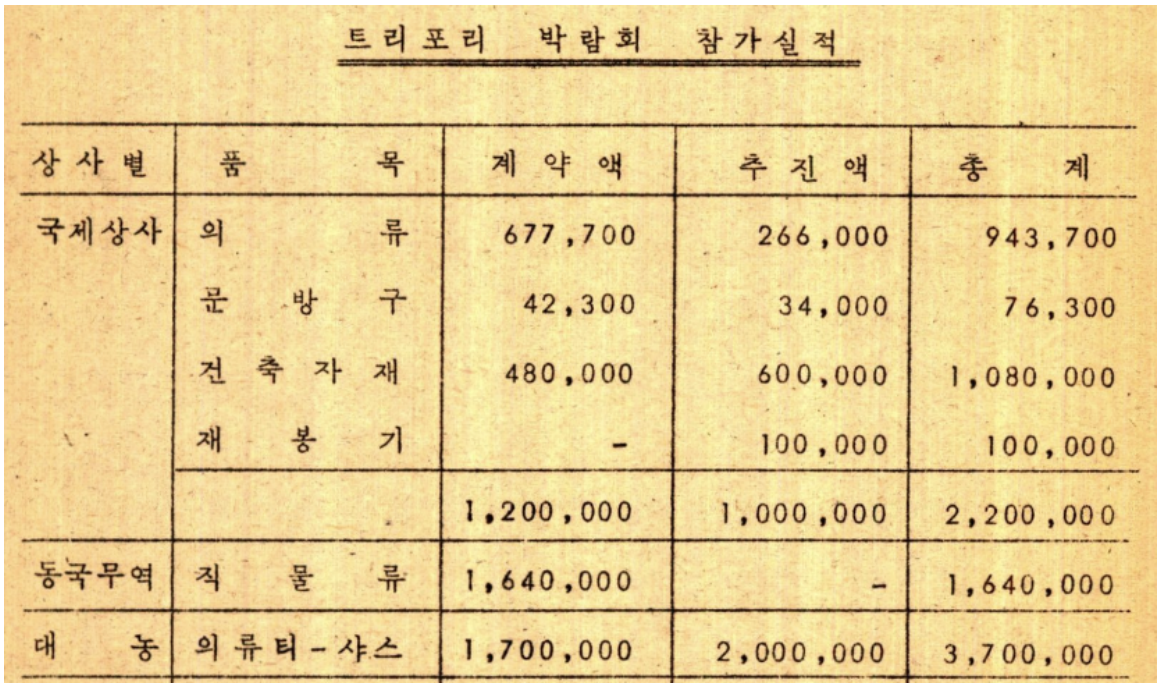

트리포리 박람회 참가실적

상사별	품목	계약액	추진액	총계
국제상사	의류	677,700	266,000	943,700
	문방구	42,300	34,000	76,300
	건축자재	480,000	600,000	1,080,000
	재봉기	-	100,000	100,000
		1,200,000	1,000,000	2,200,000
동국무역	직물류	1,640,000	-	1,640,000
대농	의류티-샤스	1,700,000	2,000,000	3,700,000

Notes: This figure shows a snippet of the summary table on sales performance from the 15th Session of Tripoli International Fair in March, 1977. The columns report, from left to right, firm name, product category, contracted amount, additional amount under negotiation, and total transaction value. All amounts are in contemporaneous U.S. dollars.

D Further Diagnostics of the Variation Explained by Bureaucrats

D.1 Misspecification checks

This section explores the additive separability between bureaucrat and country effects that is implicit in equation (3). It finds the following: (1) Residuals by quartiles of bureaucrat and country effects do not indicate misspecification. (2) Bureaucrat effects are stable between appointments. (3) Upon switches between bureaucrats, expected jumps in exports occur consistently for many different types of transitions between high, middle, and low-ability bureaucrats.

First, given that violations of additive separability would result in residuals with high absolute values for certain kinds of bureaucrat-country pairs, we check for systematic patterns in the residuals. Following the literature (Fenizia, 2022; Otero and Muñoz, 2025), we divide our observations into a 4×4 grid based on the quartile of the estimated bureaucrat fixed effect and the quartile of the estimated country fixed effect. If bureaucrat ability mattered more in small countries, for example, we would expect large positive (negative) residuals for observations in the top (bottom) quartile bureaucrats and the bottom quartile countries. Appendix Figure D.1 shows that mean residuals do not exhibit any clear pattern. Further, mean residuals are small for each combination of bureaucrat and country quartiles—between -0.05 and 0.05 in absolute value. This allays concerns about the existence of match effects.

Second, we explore how much the effect of a bureaucrat differs across their appointments. If bureaucrat effects differed greatly between appointments, this could indicate misspecification because either (1) bureaucrat-country are not additively separable, i.e., there are strong match effects, or (2) the estimated bureaucrat effects mainly pick up noise that is not correlated between appointments.[69] Appendix Table B.1 reports the changes in the R^2 when we regress exports on different combinations of fixed effects.[70] Adding bureaucrat fixed effects increases the adjusted R^2 by 0.018 (column 2 vs. column 1), about 18.6% of the increase in explanatory power from adding country fixed effects.[71] To understand

[69] Both of these points, especially (2), also constitute a reason to test whether bureaucrat fixed effects are predictive out of sample (see Section 4.4).

[70] While informative, the R^2 analysis is subject to limited mobility bias. In the formal variance decomposition exercise of Table 2, the issue is addressed by adopting the Kline et al. (2020) bias correction method.

[71] Similar to the results from the variance decomposition, the explanatory power of individual effects is somewhat smaller than in other recent papers studying the role of public sector managers. The absolute increase in R^2 when adding bureaucrat fixed effects, as well as the comparative increase in R^2 when adding

whether bureaucrat effects differ between appointments, we compare the R^2 when including bureaucrat-country fixed effects (column 4) to when bureaucrat and country fixed effects are separately added (column 3; corresponding to equation (2), our main specification). The increase in explanatory power (R^2) is negligible, suggesting that bureaucrat effects are relatively stable across appointments and that match effects, if present, are quantitatively not relevant in our setting. Further, it supports the argument that bureaucrat fixed effects are not driven by noise, because if they were, allowing for appointment-specific effects would likely increase explanatory power by more than is observed.

Third, we provide a further non-parametric check that our bureaucrat fixed effects obtain meaningful variation across the different types of transitions between high- and low-ability bureaucrats. Appendix Figure D.2 shows time trends in residualized exports around the year when an office experiences a change in the managing bureaucrat. It classifies switches between bureaucrats into terciles of effectiveness of the new and old bureaucrat, closely following Card et al. (2013), Card et al. (2016) and Best et al. (2023). Appendix Figure D.2 shows that the main takeaways from Figure 3 are present for transitions across all terciles of old and new bureaucrats. In the pre-period, exports are highest when the old bureaucrat is in the top tercile and lowest when the old bureaucrat is in the bottom tercile. In the post-period, the effect of the old bureaucrat's tercile becomes less important, and instead, the effect of the new bureaucrat's tercile becomes dominant. In event year one—the second full year of the incoming bureaucrat—exports are lowest when the new bureaucrat is in the bottom tercile and the highest when the new bureaucrat is in the top tercile. Moreover, exports change sharply and in the expected direction, precisely when there is a change in the tercile between the old and new bureaucrats. Exports increase the most upon a switch to the highest tercile and (relatively) decrease the most upon a switch to the lowest tercile.[72]

Overall, our main specification in equation (2) assuming additive separability passes multiple diagnostic checks.

bureaucrat fixed effects vs. when adding organization/location fixed effects, is smaller than in Fenizia (2022) (studying managers of organizations that process insurance claims) or Otero and Muñoz (2025) (studying hospital CEO's).

[72]The figure also shows little sign that export trends are evolving differently for different tercile transitions, suggesting that drift in effectiveness and switches are uncorrelated, in support of the identifying assumption discussed in Section 4.1.

Figure D.1: Residuals By Estimated Bureaucrat and Country Effects.
No Clear Pattern Suggesting Misspecification.

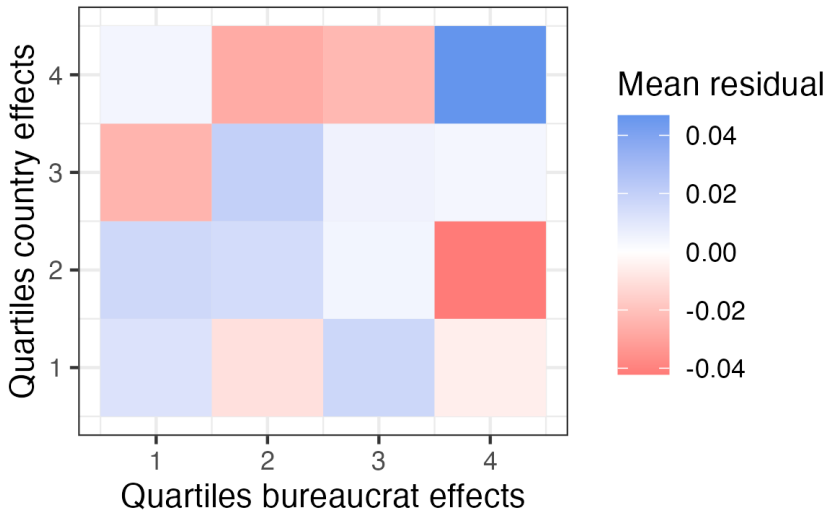

Notes: This figure shows mean residuals from equation (2) with cells defined by quartiles of estimated bureaucrat and country effects.

Figure D.2: Mean Residualized Exports Around Bureaucrat Switches.
Consistent Effects Across Terciles of New and Old Bureaucrats.

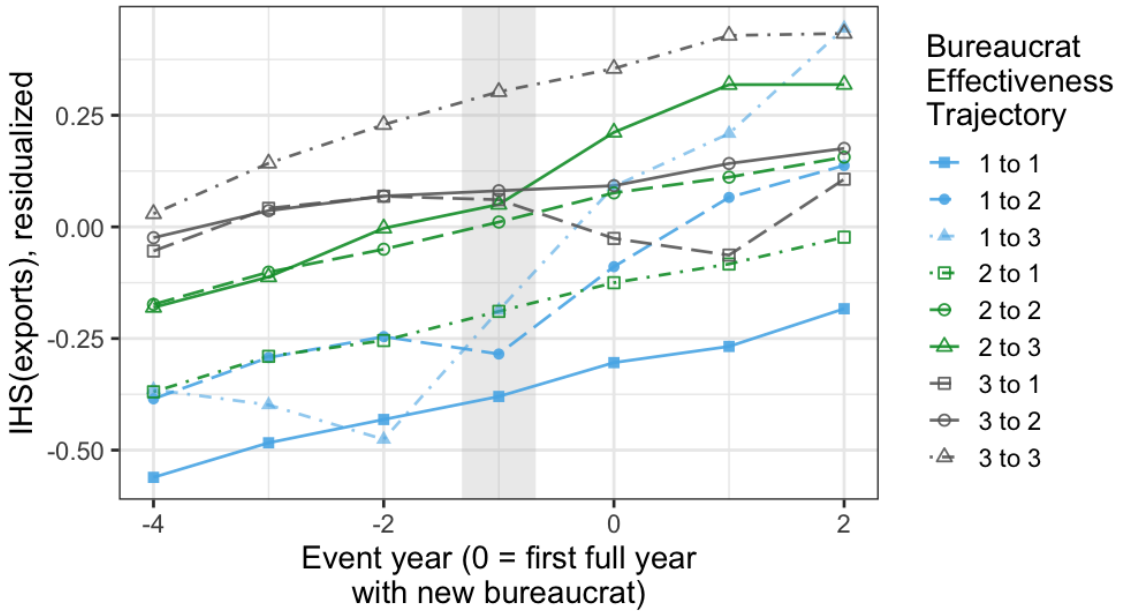

Notes: The figure shows time trends in exports around the time of bureaucrat switches. The switch occurs in year -1 so year 0 is the first full year that the new bureaucrat manages the country office. Year -2 is the last full year with the old bureaucrat. The vertical axis measures the average residualized exports for each event year and trajectory group. Exports are residualized for product-year and country fixed effects, in line with how bureaucrat fixed effects are estimated in equation (2).

D.2 Sensitivity of FE to Alternative Transformations of Exports

In this section, we investigate how much the bureaucrat fixed effect estimates would change under alternative transformations of export values.

The weakness of the IHS transformation is that it is sensitive to the scaling of export values, e.g., whether in dollars, thousands of dollars, millions of dollars, etc. That is because the IHS weights the extensive and the intensive margins differently depending on the scale.

First, to focus solely on the extensive margin, Appendix Figure D.3 Panel (a) displays a scatterplot comparing the benchmark bureaucrat fixed effects estimates for $y_{cpt} = \text{IHS}(\text{exports}_{cpt})$, to alternative estimates based on $y_{cpt} = \mathbb{1}\{\text{exports}_{cpt} > 0\}$. The alternative fixed effects purely account for the share of products with any positive Korean export value to country c in year t. Naturally, the precise estimates change greatly due to this substantial change in the outcome variable. Nevertheless, Panel (a) shows a strongly increasing, largely linear relationship between the benchmark fixed effect estimates and the estimates based on just the extensive margin.

Next, we adopt Chen and Roth (2023)'s proposal to address the weakness of the IHS by explicitly assigning a weight to the extensive margin relative to the intensive margin. Appendix Figure D.3 Panels (b) and (c) display a scatter plot comparing the benchmark estimates to alternative estimates based on the following transformation:

$$y_{cpt} = \begin{cases} log(\text{exports}_{cpt}), & \text{if exports}_{cpt} > 0, \\ -m, & \text{if exports}_{cpt} = 0. \end{cases}$$

The choice of m determines the weight of the extensive margin. Panel (b) displays benchmark estimates vs. alternatives based on extreme values of m: 0.05, 2.00, 10.00. For example, when $m = 0.05$, we set the value of going from 0 to 1 unit of export value (0 to 1,000 nominal U.S. dollars) of a product to be worth the same as increasing the export value of an already-exported product by 5%. Although the dispersion of the bureaucrat fixed effects expands with large values of m, the scatter plot appears linear for all three values of m, indicating an extremely strong positive correlation with the benchmark fixed effects.

Panel (c) shows that the benchmark estimates' dispersion of fixed effects lies between estimates based on $m = 0.50$ and $m = 1.00$, thus providing bounds for the weight implicitly assigned to the extensive margin by the IHS. Furthermore, these two values of m result in very similar fixed effect estimates, speaking to the robustness of our results to different weighting of the extensive margin.

Figure D.3: Bureaucrat Fixed Effects Under Alternative Transformations of Exports

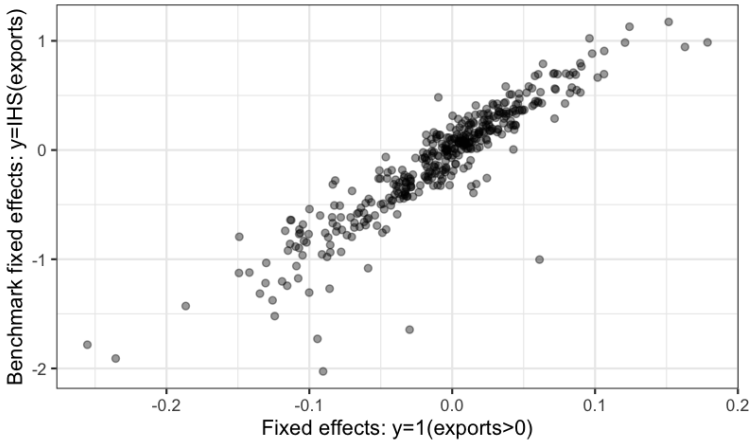

(a) Compared to FE based on the extensive margin

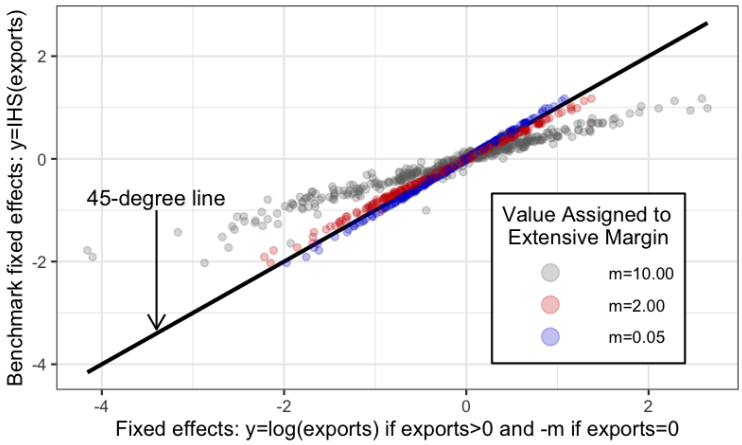

(b) Compared to FE based on explicitly assigning *extreme* values to the extensive margin

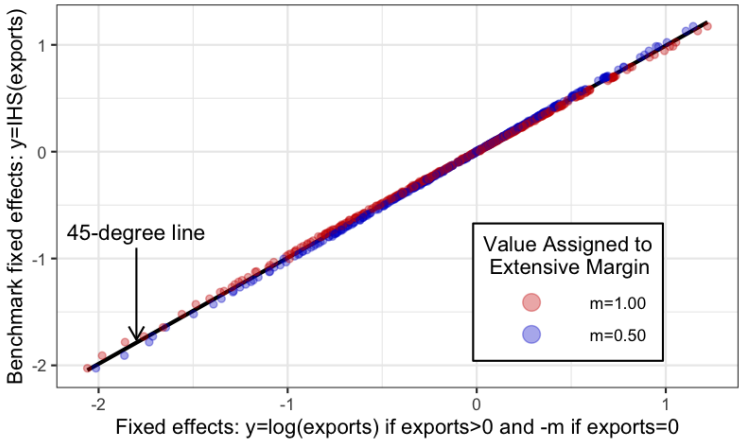

(c) Compared to FE based on explicitly assigning *moderate* values to the extensive margin

Notes: This figure displays the relationship between the benchmark bureaucrat fixed effects (outcome variable y: inverse hyperbolic sine (IHS) of export value) and those estimated under alternative transformations of export value. The bureaucrat fixed effects are estimated from equation (2).

E Connected Set and Leave-One-Out Connected Set

Appendix Figure E.1 below illustrates the concept of the connected set—moves of bureaucrats connecting countries.

Panel (a) shows two separate connected sets. Panel (b) shows how a single bureaucrat b_3 can connect the two sets, resulting in one larger connected set. In a sparsely connected network as in Panel (b), where bureaucrats move between only a few countries, we have the limited mobility bias problem: Idiosyncratic shocks are not averaged away in the estimates of bureaucrat and country fixed effects in equation (2). For instance, if there is a positive shock in Korean exports to country c_2 during b_3's appointment to c_2, the shock will be absorbed into $\hat{\gamma}_{c_2}$ and $\hat{\theta}_{b_3}$, biasing them upwards. The bias propagates through the network. $\hat{\gamma}_{c_1}$ would be biased downwards to compensate for the positively biased $\hat{\theta}_{b_3}$. Similarly, $\hat{\theta}_{b_2}$ would be biased downwards to compensate for the positive bias in $\hat{\gamma}_{c_2}$. The finite-sample estimation errors in the bureaucrat fixed effects imply that the variance of the estimated fixed effects overstates the true variance in bureaucrat ability.

Panel (c) displays a country-bureaucrat graph where Mexico, Peru, and the U.S. constitute a single leave-one-out connected set. This is the sample of countries and bureaucrats that remains connected even when removing any single appointment (bureaucrat-country pair) from the data. By restricting attention to such a leave-one-out connected set, the limited mobility bias is greatly attenuated (Kline et al., 2020). More importantly, the leave-one-out connected set allows us to obtain consistent estimates of the true variance in bureaucrat ability by applying the Kline et al. (2020) correction for limited mobility bias, under weak assumptions on the form of the bias.

The bureaucrat-country graph in our data is densely connected. Even the leave-one-out connected set contains most of the bureaucrats and countries of the full sample (380 bureaucrats out of 398, and 75 countries out of 87). Panels (d)-(f) of Appendix Figure E.1 illustrate why. They show how many connections the UK has with other countries thanks to just three bureaucrats.

Figure E.1: KOTRA Bureaucrats' Rotation Results in a Single Connected Set

(a) Country-bureaucrat graph composed of two connected sets

(d) Connections between the UK and other countries due to the bureaucrat appointed to manage the UK office in 1981

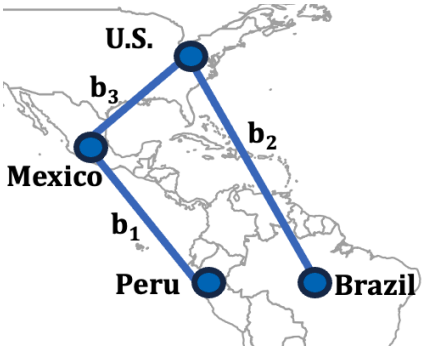

(b) Country-bureaucrat graph composed of a single connected set

(e) Connections with the UK due to 1981 and 1984 appointments

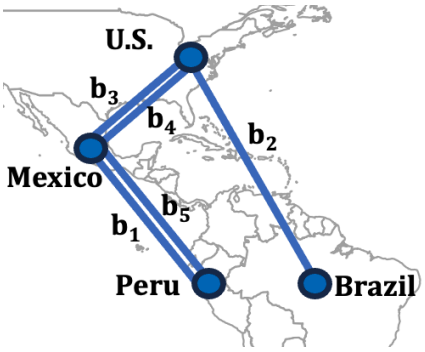

(c) Leave-one-out connected set includes U.S., Mexico, and Peru, but not Brazil

(f) Connections with the UK due to 1981, 1984, 1987 appointments

Notes: This figure illustrates how the country-bureaucrat graph of our data fulfills the requirement to form a single connected set. The left Panels (a)-(c) display hypothetical country-bureaucrat graphs. The nodes indicate the countries, and the edges indicate the bureaucrats connecting two countries. For example, b_1 is observed in both Mexico and Peru. b_2 is observed in both Brazil and the U.S. The right Panels (d)-(f) display the connections between the UK and other countries in the actual data, formed thanks to the three bureaucrats appointed to the UK office in 1981, 1984, and 1987.

F Bootstrap Bias Correction

In this section, we slightly change the notation in two ways for simplicity and clarity.

1. We simplify the notation of $\theta_{b(c,t)}$ with θ_b, for bureaucrat $b = 1, ..., n$. Country fixed effects remain as γ_c, for country $c = 1, ..., C$.

2. We distinguish the dimension over which a variance, covariance, or expectation is computed through subscripts. For example, $\mathsf{Var}_b(\hat{\theta}_b)$ is the variance of $\hat{\theta}_b$ across bureaucrats, $\mathsf{Cov}_{bc}(\theta_b, \gamma_c)$ is covariance of θ_b and γ_c across bureaucrats and countries, and $\mathsf{Var}(\hat{\theta}_b)$—no subscript—is the variance of $\hat{\theta}_b$ for an individual bureaucrat b. Based on this notation, the variance decomposition equation (3) becomes[73]

$$\mathsf{Var}_{bc}[(\text{exports}|pt)_{bc}] = \mathsf{Var}_b(\theta_b) + \mathsf{Var}_c(\gamma_c) + 2\mathsf{Cov}_{bc}(\theta_b, \gamma_c) + \mathsf{Var}_{bc}(\epsilon_{bc})$$

The plug-in estimators of the variances and covariance of bureaucrat and country fixed effects, $\mathsf{Var}_b(\hat{\theta}_b)$, $\mathsf{Var}_c(\hat{\gamma}_c)$, and $\mathsf{Cov}_{bc}(\hat{\theta}_b, \hat{\gamma}_c)$, are biased due to sampling errors in the fixed effects estimates. The bootstrap bias correction computes the degree of bias by estimating the variances and covariances of the sampling errors via bootstrap.

Let us first focus on the variance of bureaucrat fixed effects. Equation (2) will produce an estimate of θ_b: $\hat{\theta}_b = \theta_b + v_b$, where v_b is sampling error. The variance of the estimated bureaucrat fixed effects, $\mathsf{Var}_b(\hat{\theta}_b) = \mathsf{Var}_b(\theta_b) + \mathsf{Var}_b(v_b) = \sigma^2_\theta + \sigma^2_v$, will overstate the true dispersion in bureaucrat quality, σ^2_θ.

By the law of total variance, $\sigma^2_v \equiv \mathsf{Var}_b(v_b) = \mathsf{E}_b(\mathsf{Var}(v_b)) + \mathsf{Var}_b(\mathsf{E}(v_b)) = \mathsf{E}_b(\sigma^2_{v_b}) + \mathsf{Var}_b(0) = \mathsf{E}_b(\sigma^2_{v_b})$. Hence, an estimate of σ^2_v is given by a weighted average of the sampling-error variances of individual bureaucrats:

$$\sigma^2_v = \mathsf{E}_b(\sigma^2_{v_b}) \Rightarrow \hat{\sigma}^2_v = \frac{\sum_{b=1}^n \hat{\sigma}^2_{v_b} w_b}{\sum_{b=1}^n w_b} \tag{F.1}$$

where the weights w_b reflect the weight of bureaucrat b (the number of country-product-year observations of b) in the original sample.

The variance of the sampling error for bureaucrat b, $\hat{\sigma}^2_{v_b}$, is estimated by bootstrapping[74]

[73] Recall that fixed effects estimated on the raw data and those estimated on the data collapsed at the bureaucrat-country level are identical.

[74] We cluster-bootstrap at the appointment (bureaucrat-country match) level, the level of the variation identifying the bureaucrat and country fixed effects.

equation (2) K times [75] to get the variance of $\hat{\theta}_b$ across bootstrap samples $k = 1, ..., K$:

$$\sigma^2_{v_b} = \mathsf{Var}(\hat{\theta}_b) \Rightarrow \hat{\sigma}^2_{v_b} = \widehat{\mathsf{Var}}_k(\hat{\theta}^k_b) = \frac{\sum_{k=1}^K (\hat{\theta}^k_b - \bar{\hat{\theta}}_b)^2 w^k_b}{\sum_{k=1}^K w^k_b - 1} \tag{F.2}$$

where $\hat{\theta}^k_b$ is the estimated fixed effect of b in bootstrap sample k, the weights w^k_b are the weight of b in bootstrap sample k, and $\bar{\hat{\theta}}_b = \frac{\sum_{k=1}^K \hat{\theta}^k_b w^k_b}{\sum_{k=1}^K w^k_b}$.

Then, $\hat{\sigma}^2_\theta = \widehat{\mathsf{Var}}_b(\hat{\theta}_b) - \hat{\sigma}^2_v$. $\hat{\sigma}^2_\theta = 0.100$, reported in column 3 of Table 2.

The variance of the country fixed effects can be estimated similarly. Where $\hat{\gamma}_c = \gamma_c + u_c$ is the estimated fixed effect of country c, with u_c representing the sampling error, we can estimate the dispersion of country fixed effects, σ^2_γ, by:

$$\hat{\sigma}^2_\gamma = \widehat{\mathsf{Var}}_c(\hat{\gamma}_c) - \hat{\sigma}^2_u. \tag{F.3}$$

$\hat{\sigma}^2_\gamma = 0.682$, reported in column 3 of Table 2.

The covariance of the estimated bureaucrat and country fixed effects is also a biased estimator of the true covariance:

$$\begin{aligned}
\mathsf{Cov}_{bc}(\hat{\theta}_b, \hat{\gamma}_c) &= \mathsf{Cov}_{bc}(\theta_b, \gamma_c) + \mathsf{Cov}_{bc}(\theta_b, u_c) + \mathsf{Cov}_{bc}(\gamma_c, v_b) + \mathsf{Cov}_{bc}(v_b, u_c) \\
&= \mathsf{Cov}_{bc}(\theta_b, \gamma_c) + \mathsf{Cov}_{bc}(v_b, u_c) \\
&= \sigma_{\theta\gamma} + \sigma_{vu}
\end{aligned}$$

By the law of total covariance, $\sigma_{vu} \equiv \mathsf{Cov}_{bc}(v_b, u_c) = \mathsf{E}_{bc}(\mathsf{Cov}(v_b, u_c)) + \mathsf{Cov}_{bc}(\mathsf{E}(v_b), \mathsf{E}(u_c)) = \mathsf{E}_{bc}(\sigma_{v_b u_c}) + \mathsf{Cov}_{bc}(0, 0) = \mathsf{E}_{bc}(\sigma_{v_b u_c})$. Then,

$$\sigma_{vu} = \mathsf{E}_{bc}(\sigma_{v_b u_c}) \Rightarrow \hat{\sigma}_{vu} = \frac{\sum_{b=1}^n \sum_{c=1}^C \hat{\sigma}_{v_b u_c} w_{bc}}{\sum_{b=1}^n \sum_{c=1}^C w_{bc}} \tag{F.4}$$

where w_{bc} reflect the weight of the (b, c) pair (the number of product-year observations of each (b, c) pair) in the original sample.

The covariance of the sampling errors for bureaucrat b and country c in an (b, c) pair is again estimated from the bootstrap results:

$$\sigma_{v_b u_c} = \mathsf{Cov}(v_b, u_c) \Rightarrow \hat{\sigma}_{v_b u_c} = \widehat{\mathsf{Cov}}_k(\hat{\theta}^k_b, \hat{\gamma}^k_c) = \frac{\sum_{k=1}^K (\hat{\theta}^k_b - \bar{\bar{\hat{\theta}}}_b)(\hat{\gamma}^k_c - \bar{\bar{\hat{\gamma}}}_c) w^k_{bc}}{\sum_{k=1}^K w^k_{bc} - 1} \tag{F.5}$$

where w^k_{bc} is the weight of (b, c) in bootstrap sample k, and $\bar{\bar{\hat{\theta}}}_b = \frac{\sum_{k=1}^K \sum_{c=1}^C \hat{\theta}^k_b w^k_{bc}}{\sum_{k=1}^K \sum_{c=1}^C w^k_{bc}}$, $\bar{\bar{\hat{\gamma}}}_c = \frac{\sum_{k=1}^K \sum_{b=1}^n \hat{\gamma}^k_c w^k_{bc}}{\sum_{k=1}^K \sum_{b=1}^n w^k_{bc}}$.

Finally, $\hat{\sigma}_{\theta\gamma} = \widehat{\mathsf{Cov}}_{bc}(\hat{\theta}_b, \hat{\gamma}_c) - \hat{\sigma}_{vu}$. $\hat{\sigma}_{\theta\gamma} = -0.044$, reported in column 3 of Table 2.

[75] We choose $K = 500$, which we confirm to be a sufficiently large number so that the sampling-error variance and covariance estimates stabilize.

G Robustness of Appointments by Country Importance

This section examines the pattern shown in Panels (c) and (d) of Figure 6: bureaucrats are assigned to the least important countries in their first appointments, to the most important countries in their third appointments, with second appointments forming an intermediate case.

Appendix Figure G.1 highlights the flows of bureaucrats across terciles of country importance, from first to second and third appointments, or to exiting KOTRA. A similar picture emerges whether we consider terciles of office opening year (in reverse order, consistent with Panels (c) and (d) of Figure 6) or terciles of country effect. First, around half of the bureaucrats exit after their first appointment. The exit rate appears unrelated to the tercile of the first appointment. Second, between appointments one and two, bureaucrats move across all terciles of countries, although there is some persistence. Third, between appointments two and three, almost all bureaucrats leave the first tercile (latest third of openings and the lowest third of country fixed effects), the least important countries. In contrast, bureaucrats mostly stay in the third tercile, the most important countries.[76]

It should be noted that more important countries may also be more desirable for bureaucrats. Thus, this pattern of appointments is also broadly in line with an alternative mechanism where progressively better postings are used as career incentives.

Overall, there is a clear pattern of bureaucrats being moved towards more important countries as their careers progress.

[76] From appointments one to two, there is already somewhat limited flow from the most important third to the least important third of countries, particularly in terms of opening years.

Figure G.1: Bureaucrat Flows By Appointment and Country Importance

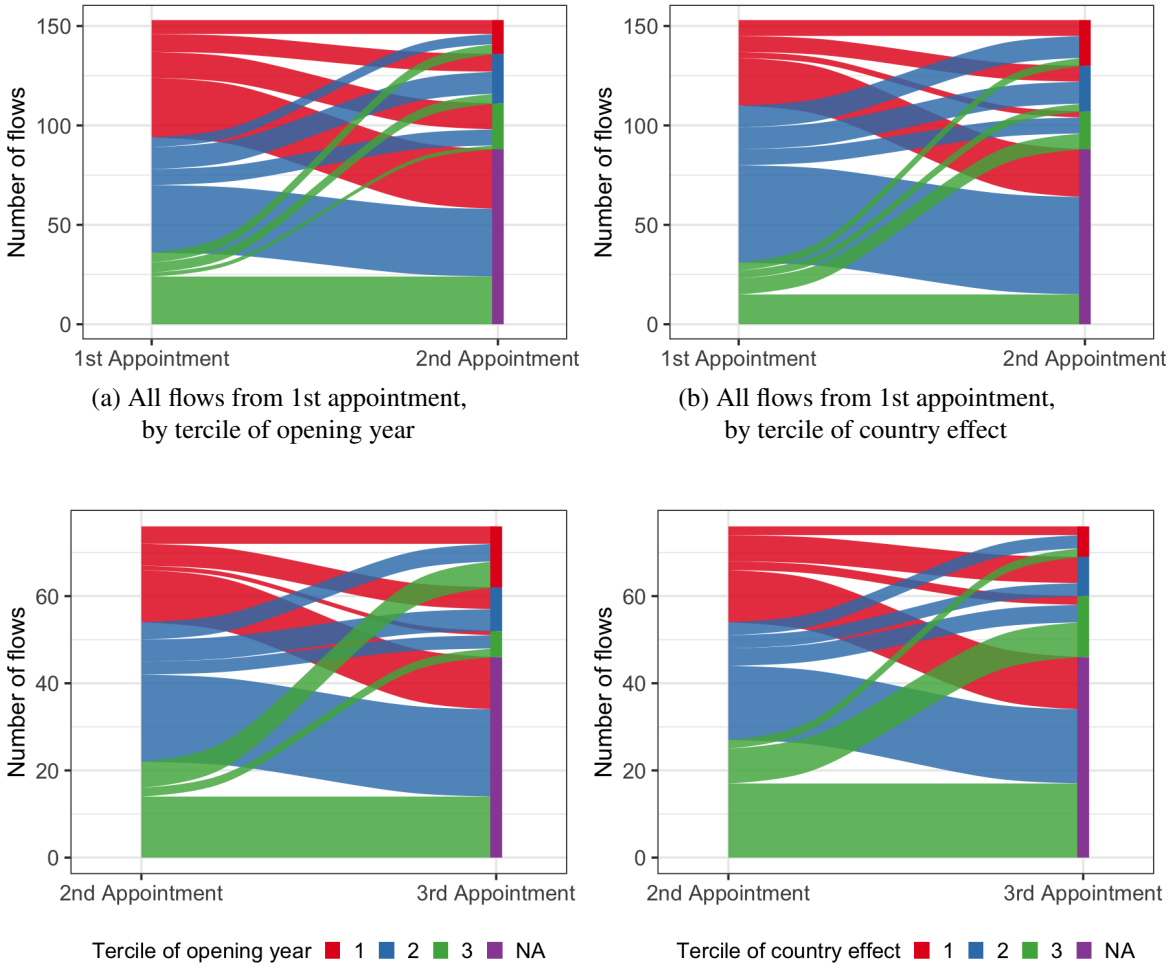

(a) All flows from 1st appointment, by tercile of opening year

(b) All flows from 1st appointment, by tercile of country effect

(c) All flows from 2nd appointment, by tercile of opening year

(d) All flows from 2nd appointment, by tercile of country effect

Notes: This figure shows the flow of bureaucrats from their 1st to 2nd, and from 2nd to 3rd appointments. We split the offices into terciles of office opening year in reverse order (left panels) and terciles of country effects (right panels). The third tercile by either metric contains the most important trade partners of Korea. We consider offices that were opened during the main rollout period (1962-1981). "NA" indicates that the bureaucrat is not subsequently observed as an overseas office manager, including cases where they exit KOTRA, are assigned to domestic offices, or are posted to an overseas office at a lower rank in the hierarchy. To avoid a distortion in the relationship driven by a bureaucrat's appointment being limited to offices that are already open, we only include bureaucrats whose first appointments started after the rollout period, i.e., 1981 or later. Early openings and large country fixed effects proxy for the importance of a country as Korea's trade partner. For both metrics, the third terciles represent the most important group of countries, and the first terciles represent the least important group.

H Robustness of Office Opening Results

H.1 Not-Yet-Treated Control Group

This section performs robustness checks on the results to using an alternative control group: the not-yet-treated countries.

This complements the main results that use never-treated countries as the control group. The absence of differential pre-trends across sample periods and outcome variables alleviates concerns about violations of the parallel trends assumption. The never-treated control group also sidesteps bias in estimating (7) caused by already-treated units being present in the control group when there is treatment timing heterogeneity and dynamic treatment effects.

Panel (a) of Appendix Figure H.1 shows that our estimates are not importantly driven by our choice of countrol group. We plot the results from a not-yet-treated control group, using the Callaway and Sant'Anna (2021) estimator. The estimates are of very similar magnitude and precision to the ones from the main estimation strategy (Panels (a) and (b) of Figure 7). However, the coefficients in periods -4 to -2 are negative and statistically significant, albeit small. We thus investigate the sensitivity of the estimates to violations of the parallel trends assumption (PTA) following the approach proposed by Rambachan and Roth (2023). Under this approach, the causal effect of an office is partially identified under a large class of restrictions that impose that the post-treatment violations of parallel trends are not too different from the pre-trends. For this analysis, we zoom into the causal effect on exports 10 years after an office opens.

Panel (b) adopts the "relative magnitudes" approach of Rambachan and Roth (2023). It shows that the 95% robust confidence set of the year-10 causal effect of an office does not contain zero when allowing for parallel trend violations up to one times the largest pre-treatment violation ($\bar{M} = 1$). Conceptually, $\bar{M} = 1$ means we allow for possible violations of parallel trends driven by confounding economic shocks that are of a similar magnitude to confounding economic shocks in the pre-period. The fact that we can rule out a null effect 10 years later is rather striking, given the flexibility and the long period of the violations to PTA permitted. The approach bounds the violation of parallel trends across *consecutive periods*, implying that identified sets are larger for later post-treatment periods since the treatment and control groups have more time to diverge (e.g., the identified set for the second period will be twice as large as for the first period).

Figure H.1: Robustness of Office Opening Results to
a Not-Yet-Treated Control Group

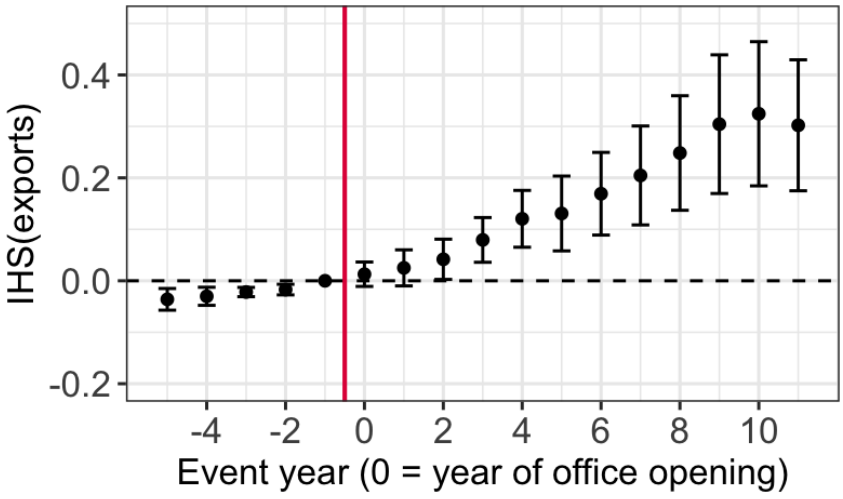

(a) Results with a not-yet-treated control group

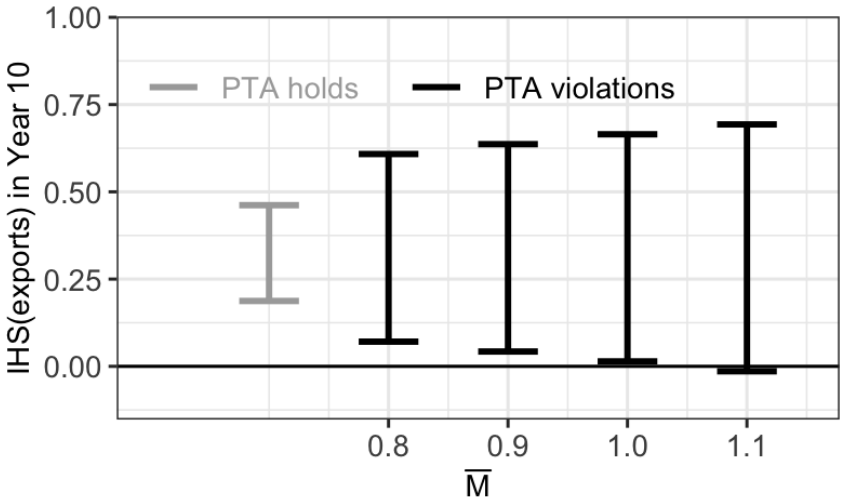

(b) Confidence intervals for the year-10 estimate, allowing for violations of the parallel trends assumption of up to $\bar{M}\times$ the largest pre-treatment violation

Notes: The figure reports results from estimating equation (7), with a not-yet-treated control group. Panel (a) plots the average treatment effect on the treated based on the Sant'Anna and Zhao (2020) estimator, a doubly-robust version of the Callaway and Sant'Anna (2021) estimator for difference-in-difference settings with staggered rollout. The outcome variable is the inverse hyperbolic sine (IHS) of Korean exports to the country of the office opening. Bootstrapped standard errors are obtained clustering at the country level. Panel (b) reports the sensitivity of the year-10 estimate in (a), estimated under the parallel trends assumption (PTA), to violations of PTA (Rambachan and Roth, 2023). The left-most bar (gray) in (b) reproduces the 95% confidence interval of the year-10 estimate in (a). The other bars (black) represent the 95% confidence intervals for different degrees of violation. We bound the maximum post-treatment violation of parallel trends by $\bar{M}$ times the maximum pre-treatment violation of parallel trends.

H.2 Extensive Margin

This section examines whether the effects of office openings on the IHS of exports, shown in Panels (a)-(b) of Figure 7, reflect changes along the product-level extensive margin. We estimate equation (7) with y_{cpt} changed to a dummy indicating whether the export value of product p to country c in year t is strictly positive. The question becomes: Does a KOTRA office in a country increase the likelihood that Korea exports a given product to that country? Alternatively: Does a KOTRA office in a country increase the share of products that Korea exports to that country?

Appendix Figure H.2 reports the estimated effects of offices on the extensive margin. It indicates a 5 percentage point increase in the likelihood of a product being exported to a destination country 10 years after an office opening. Although coefficient values are not directly comparable, the trajectory of the point estimates in Panel (a) is very similar to that in the main results (Figure 7). Again, the pre-treatment coefficients are very close to zero, corroborating that the parallel pre-trends in Figure 7 are not due to the IHS somehow obscuring differential pre-trends. The fact that the coefficients slope upwards *starting with* the office opening supports the causal interpretation of the estimated effects. The coefficient estimates become economically sizable as early as one year after the opening and statistically significant at the 5%-level from two years after the opening. As in Figure 7, the coefficients stabilize after around ten years.

The results remain similar in Panel (b), where we expand the sample to include openings from 1964 at the cost of a shorter pre-period.

Figure H.2: The Effect of Offices on the Extensive Margin of Exports

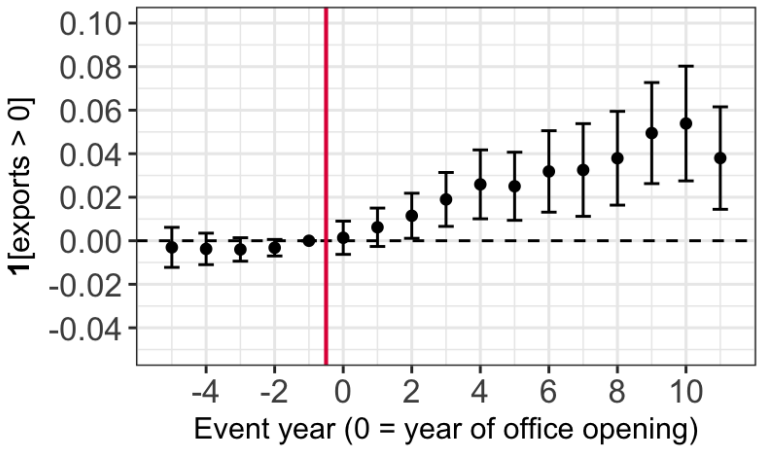

(a) Sample includes openings from 1967

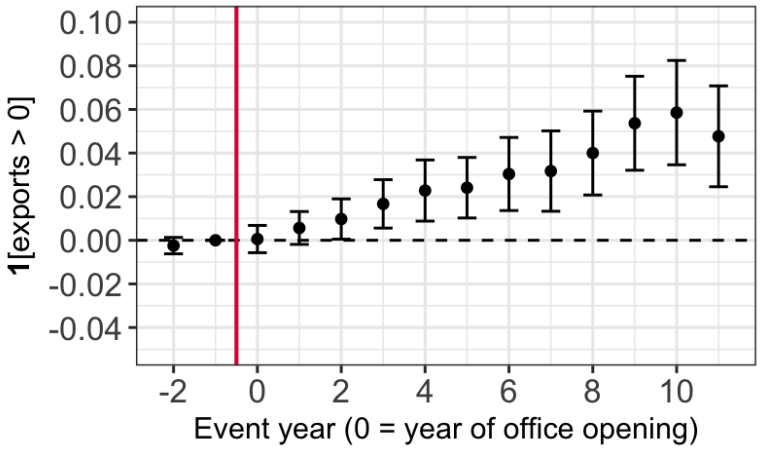

(b) Sample includes openings from 1964

Notes: The figure reports results from estimating equation (7), with a never-treated control group. The outcome variable is a dummy indicating whether South Korea had a positive export value in a product to the country of the office opening—hence corresponding to a linear probability model. In Panel (a), we have the main sample, which includes office openings from 1967. In Panel (b), we increase power by expanding the sample to include office openings from 1964. Standard errors are clustered at the country level.