PERSONNEL IS POLICY (IMPLEMENTATION) : BUREAUCRATS AND THE KOREAN EXPORT MIRACLE*

Philipp Barteska
Harvard Kennedy School

Jay Euijung Lee
Stockholm University

April 14, 2025

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 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 a country's import demand, suggesting a 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: philippbarteska@hks.harvard.edu. Harvard University, John F. Kennedy School of Government, 79 John F. Kennedy Street, Cambridge, MA 02138, United States.

This paper benefited from comments by David Atkin, Jie Bai, Tim Besley, Michael Best, Mike Callen, Swati Dhingra, Alessandra Fenizia, 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 countries dramatically shifting their industry composition from producing for domestic markets using traditional technology to producing for global markets using modern technology. Proponents of state intervention often highlight the prevalence of industrial policy during these export-led growth episodes. However, achieving success with industrial policies is challenging because states often lack the necessary information, such as local knowledge required in allocating subsidies or loans between firms, facilitating FDI, or connecting firms to markets. The state relies on bureaucrats to collect this information, but their ability and motivation to gather, interpret, and act on it can vary significantly.[1] The challenge for the state is compounded if it cannot observe these idiosyncratic capabilities to select the best talent. As a result, the effectiveness of industrial policies may depend critically on the bureaucracy implementing them. Differences in bureaucratic capacity may help explain why attempts to replicate successful policies have fallen short.

In this paper, we study how much the effect of an industrial policy depends on the idiosyncratic ability of the responsible 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, Lane, and Rodrik, 2023),[2] 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, we observe the bureaucrats who managed each country office from 1965 to 2000, from a novel dataset we compiled from archival records. These bureaucrats rotated every three years, providing variation in implementing capacity within location,

[1] Recent research has found substantial variation in the effectiveness of bureaucrats in procurement and the processing of social insurance claims (Best, Hjort, and Szakonyi, 2023; Fenizia, 2022).

[2] Recent papers have described policy instruments and assessed policies' average – and sometimes distributional – effects, while considering how these effects are shaped by the policy instruments or underlying economic conditions (Aghion, Cai, Dewatripont, Du, Harrison, and Legros, 2015; Juhász, 2018; Barwick, Myrto, and Bin, 2025; Lane, 2025; Liu, 2019; Mitrunen, 2024; Choi and Levchenko, 2021; Choi and Shim, 2024; Juhász, Lane, Oehlsen, and Pérez, 2024; Kim, Lee, and Shin, 2022).

i.e., holding fixed time-invariant differences in economic conditions affecting the outcome. The rotation schedule is the direct reason for moves of bureaucrats, not other potentially confounding factors. Indirectly, the rotation schedule allows for tests through which we rule out major remaining endogeneity concerns. Third, throughout the sample period, the key outcome targeted by each office remained exports to the respective country. By studying bureaucrats who targeted exports, a prime economic outcome,[3] this paper stands apart from existing work on bureaucratic effectiveness, which has often focused on narrower domains for measurability of performance, such as tax collection, public goods provision, and aspects of public sector administration (Besley, Burgess, Khan, and Xu, 2022).

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, Kramarz, and Margolis, 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 exports by 37%.[4] Moreover, the variation due to bureaucrats amounts to 1/7 of that due to countries.[5] These numbers arise from a variance decomposition exercise which corrects for finite-sample bias common in two-way fixed effects frameworks (Kline, Saggio, and Sølvsten, 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.[6]

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

[3]Exporting is important for economic growth and development more broadly. For evidence highlighting the effect of exports on development outcomes at the firm level, see Atkin, Khandelwal, and Osman (2017). For evidence at the macro level, see Hausmann, Hwang, and Rodrik (2007) and Atkin, Costinot, and Fukui (2025). For support that market-size factors may be decisive in economic development, see Goldberg and Reed (2023).

[4]For comparison, from 1965 to 2000, South Korean exports grew 30-fold relative to the U.S. (IMF: Direction of Trade Statistics, 2023).

[5]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, Sollaci, and Syverson, 2023; Otero and Muñoz, 2025).

[6]This alternative bias correction builds on Best, Hjort, and Szakonyi (2023).

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 absolute effects when a manager joins an office than when they leave it, as there is much less discretion regarding the year in which a bureaucrat leaves an office. Empirically, however, we find parallel pre-trends and symmetric effects of gaining and losing a bureaucrat, in support of the identifying assumption. 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.[7] We find that the policy had a substantial effect on average, increasing exports by 38% ten years after an office opening. Thus, the policy would have no effect if implemented by bureaucrats one standard deviation below average.

The differences between the abilities of office managers are explained by one key mechanism: Better bureaucrats transmit timely market information more effectively. We show that, upon the appointment of a high-ability bureaucrat, exports of a product rise much more responsively to simultaneous growth in the product's import demand in the destination.[8]

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

[7]The main specification estimates the effect of opening an overseas office relative to a never-treated control group. This choice of control group avoids the potential biases arising in a two-way fixed effects regression caused by the combination of dynamic treatment effects with the treatment's staggered rollout. Recent advances in the difference-in-differences literature allow us to test robustness using a not-yet-treated control group (Callaway and Sant'Anna, 2021) and to test how sensitive the results are to violations of the parallel trends assumption (Rambachan and Roth, 2023).

[8]In addition, the exports of a product rise much more responsively to simultaneous growth in the product's export supply from Korea under a high-ability bureaucrat. Import demand is measured by exports of the same product from countries other than South Korea to the same destination, while export supply is measured by exports of the same product from South Korea to other destination countries.

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.[9] Hence, the most important offices are rarely run by underperforming bureaucrats.[10] These descriptive patterns align well with archival records on the organization's operations that endorse performance-based pay levels, job assignments, and promotions.[11] 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 the literature on the effects of industrial policies (Juhász, Lane, and Rodrik, 2023; Juhász, 2018; Liu, 2019; Lane, 2025; Choi and Levchenko, 2021; Choi and Shim, 2024; Kim, Lee, and Shin, 2022). Unlike previous studies, this paper focuses on understanding the variation in industrial policy outcomes driven by the bureaucrats implementing the policy. In doing so, we provide evidence for the often-hypothesized but under-researched link between state capacity and the effectiveness of industrial policy.

Second, we contribute to the literature on bureaucratic capacity and economic development (Besley, Burgess, Khan, and Xu, 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.[12]

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

[10] 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.

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

[12] Front-line bureaucrats with specific tasks are often studied, as the empirical tools from the "credibility revolution" (Angrist and Pischke, 2010) are typically more appropriate for studying them. As summarized by Besley, Burgess, Khan, and Xu (2022), these include: "agricultural extension workers", "revenue collectors", "health care providers", "teachers", "procurement officers", and "judges" (Akhtari, Moreira, and Trucco, 2022; Aman-Rana, 2025; Ashraf and Bandiera, 2018; Bandiera, Best, Khan, and Prat, 2021; Best, Hjort, and Szakonyi, 2023; Brown and Andrabi, 2021; Dahis, Schiavon, and Scot, 2023; Dal Bó, Finan, Li, and Schechter, 2021; Khan, 2025; Khan, Khwaja, and Olken, 2016, 2019; Leaver, Ozier, Serneels, and Zeitlin, 2021; Mehmood, 2022). Further relevant paper are Bergeron, Tourek, and Weigel 2024; Chetty, Friedman, and Rockoff 2014b; Janke, Propper, and Sadun 2019; Muñoz and Prem 2024; Otero and Muñoz 2025; Rasul and Rogger 2017.

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 of responsibility.[13] The bureaucrats we study embody advantages of both approaches, as they manage offices targeting the same well-defined outcome variable, observed at a geographic level: exports to the country where the office is located. This unique setting allows for a quantification of the role of bureaucrats in industrial policy implementation during a growth miracle.

Third, this paper builds on research into the role of managers in determining economic outcomes. It investigates managers by applying methods from the labor literature on worker and firm heterogeneity first proposed by Abowd, Kramarz, and Margolis (1999).[14] Fenizia (2022) and Otero and Muñoz (2025) follow a similar approach to show that public-sector managers matter in the processing of social insurance claims and in public hospitals. A closely related literature similarly finds large effects of managers on private firms (Bertrand and Schoar, 2003; Metcalfe, Sollaci, and Syverson, 2023).[15] This paper differs by studying managers implementing a major economic policy and highlighting how the organization recognizes and manages uncertainty over manager ability.

Fourth, we contribute to the literature on trade and economic development. One particular area we build on examines the impact of exporting on firm growth (Atkin, Khandelwal, and Osman, 2017; Alfaro-Ureña, Manelici, and Vasquez, 2022) and market-size constraints in development (Goldberg and Reed, 2023). We provide evidence suggesting that individual bureaucrats affect firms' access to foreign markets. A closely related line of research investigates export promotion.

[13] As summarized by Besley, Burgess, Khan, and Xu (2022), these include: "provincial governors and GDP growth (Jia 2017), governors and colony-level revenue generation (Xu 2018)" - published versions: Jia (2024), Xu (2018). Besley, Burgess, Khan, and Xu (2022) further include among these bureaucrats with CEO characteristics those with tasks more distant to economic growth, e.g., office managers in charge of processing social insurance claims (Fenizia, 2022). An interesting case forms Gulzar and Pasquale (2017) which studies Indian bureaucrats responsible for district-level development outcomes writ-large. The authors focus on one of their tasks with clearly measurable and interpretable outcomes: the implementation of the NREGA workfare program.

[14] For further discussions regarding the methodology, see Card, Heining, and Kline (2013); Card, Cardoso, and Kline (2016); Kline, Saggio, and Sølvsten (2020); Bonhomme et al. (2023).

[15] Other studies, using more heterogeneous methodologies, also find that managers matter greatly in the private sector: Mollick (2012); Lazear, Shaw, and Stanton (2015); Bandiera, Prat, Hansen, and Sadun (2020); Frederiksen, Kahn, and Lange (2020); Hoffman and Tadelis (2021); Adhvaryu, Nyshadham, and Tamayo (2022); Adhvaryu, Kala, and Nyshadham (2022); Friebel, Heinz, and Zubanov (2022); Patault and Lenoir (2024).

Similar to industrial policy, most studies on export promotion estimate average treatment effects (Munch and Schaur, 2018; Hayakawa, Lee, and Park, 2014; Bagir, 2020; Volpe Martincus and Carballo, 2008), while some delve into specific mechanisms by examining which sectors and firm types benefit the 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, 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[16] – to $24,834 (2001).[17] Figure 1 highlights that export growth was particularly stark. Korean exports per capita in the 1950s were below 2% of the U.S. level. From 1960 on, however, exports increased rapidly and reached parity with the U.S. before the end of the century. Over this period, other developing countries saw little growth relative to the U.S.[18]

It is often argued that this growth is partly due to industrial policies conducted by a well-

[16]The countries with higher GDP per capita in 1962 in Sub-Saharan Africa in order of 2023 population: Nigeria, the Democratic Republic of the Congo, South Africa, Kenya, Ghana, Madagascar, Côte d'Ivoire, Cameroon, Niger, Zambia, Chad, Senegal, Zimbabwe, Guinea, Benin, Togo, Republic of the Congo, the Central African Republic, Liberia, Mauritania, Gambia, Namibia, Gabon, Mauritius, the Comoros, the Seychelles.

[17]Both numbers are in 2017 USD. Relative to the U.S., this corresponds to an increase from 1/16 of real GDP per capita in 1961 to 1/2 in 2001. Data from Penn World Tables.

[18]Even the growth in Chinese exports under Deng Xiaoping (from 1978) is moderate relative to South Korea's export growth between 1960 and the late 1980s.

functioning state.[19] On the other hand, the Korean state was described as aid-dependent and corrupt until the 1960s (Kim and Vogel, 2011).[20] These contradicting narratives, summarized in Section 2.1, motivate us to study how much a particular industrial policy's effect was due to the capacity with which it was implemented. We focus on overseas export promotion, a policy area highlighted in a representative survey of Korean manufacturers from 1976 as the one where government intervention improved the most under President Park Chung-hee (1961-1979) compared 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).[21]

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

[19] Wade (1990) and Cheng, Haggard, and Kang (1998) as cited by Besley, Burgess, Khan, and Xu (2022); Amsden (1989); Juhász, Lane, and Rodrik (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.

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

[21] Mason argues that this was particularly important for a government that intents to apply "discretionary command procedures" in addition to non-discretionary policies. Overseas export promotion may be considered as very discretionary, as the countries targeted, and the specific services supplied to which sector in a given country, are largely up to the decision of the bureaucrat assigned to the country.

instruments used for political purposes, but the staffing of the bureaucracy itself was an important form of patronage [Suh (sic), 1967]." (Cheng, Haggard, and Kang, 1998; Bark, 1967).[22] The limited capacity of the state 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." As over 90 percent 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). Figure 2 displays the countries where an overseas KOTRA office opened between 1962 and 2001. Most of these offices remained open from the indicated year until 2000, the end of the sample period, and beyond.[23]

While the South Korean government had many policies aimed at increasing exports, KOTRA's

[22] Cheng, Haggard, and Kang (1998) describe the Rhee bureaucracy in unfavorable terms except regarding their ability to extract long-term aid commitments from the U.S. Jones and Il (1980) also note that corruption was widespread during the Rhee presidency.

[23] Appendix Figure A.1 displays the number of countries with an overseas KOTRA office over time. By 1970, offices had opened in 32 countries. This number rose to 75 by 1981. Then the pace slowed.

overseas offices each targeted their efforts at exports *to a particular destination country* – making variation in the work of the overseas offices plausibly orthogonal to other industrial policies by the South Korean government, even if these targeted particular sectors or regions within South Korea.

The overseas offices contributed to three main functions of KOTRA that were maintained consistently from the early years of the organization's establishment throughout the period of study. First, KOTRA's "Investigation/Research" division investigated factors related to export supply and demand with a focus on (1) South Korea's capability to supply a product for exports and (2) the import demand in a particular foreign market. The overseas offices produced reports by product and country that were compiled and published by the head office.

Second, the overseas offices served a key role in the "Market Development" division by helping domestic producers and retailers find new trade partners in new and existing markets. This was achieved through direct matchmaking and by collecting and publicizing inquiries for imports from Korea by foreign firms, as well as export inquiries from domestic companies. These inquiries were published in KOTRA's Daily Market Newspaper. Business transactions were then mediated between the inquirers and respondents.

Third, the overseas offices helped the "Trade Fair" division with the organization of a South Korean pavilion at international trade fairs. KOTRA's reports attribute remarkably large effects to these: during a 5-day trade fair in London in 1978 the reported sales by the firms in the pavilion amounted to 3% of Korea's exports to the UK that year. The overseas offices coordinated logistics such as shipping the products that were to be exhibited. They also handled strategic tasks such as recruiting, selecting, and briefing exporters. Domestic companies were selected based on whether they were able to supply newly trending styles and designs that aligned with the market demands at the fair venues. Prior to a fair, the KOTRA office disseminated information about the exporters and their products to attract potential buyers. The bureaucrats ran ads, sent letters and made phone calls to promising exporters and foreign buyers, and reached out to trade associations.

Compared to other bureaucracies, KOTRA's overseas offices had a large degree of discretion regarding how to carry out the task of promoting exports. For this reason, we focus on KOTRA's

ultimate outcome of interest: exports. It appears difficult to centrally plan whether exports to a particular destination will benefit more from market reports or networking with potential importers, and whether networking should happen through fairs, phone calls, or some other channel. Instead, such a goal relies on the bureaucrats' knowledge, which may be both tacit and local, and requires substantial improvisation. Hence, rather than having a centrally mandated list of tasks to fulfill,[24] KOTRA office managers shared similarities with CEO-like bureaucrats running an entire geographic region in the Indian Administrative Service (Bertrand, Burgess, Chawla, and Xu, 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 exactly one policy with a clearly defined target, which can largely be summarized as exports during their appointments. The primary performance measure, as assessed by the head office, was whether export targets were met[25] – supported by Figure 8, which shows that bureaucrats saw fewer future appointments if exports underperformed during their first appointments. Thus, KOTRA bureaucrats were much less susceptible to the multi-tasking problem faced when evaluating the effectiveness of most bureaucrats with regional responsibilities. Finally, the outcome targeted by overseas offices, exports to the country, is an outcome of direct importance for economic growth and development.

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.[26]

[24] Or managing people who have a list of tasks to fulfill, as in Bandiera, Best, Khan, and Prat (2021); Fenizia (2022); Best, Hjort, and Szakonyi (2023).

[25] Archival records and interviews with former bureaucrats offer insight into some 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-based bonus could reach up to 500% of the monthly base pay. Additionally, a separate living allowance was provided to reflect 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. So, prior to a bureaucrat's first appointment as manager of a country office, it is possible for substantial uncertainty about their ability to persist.

[26] 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 most important countries (U.S., Japan) as well as geographically large ones (U.S., Canada, Australia). After our sample period, China joins both of these groups.

The assignment of bureaucrats to offices is determined by the head office's HR team. Our analysis and interviews we conducted with current and former KOTRA employees highlight several factors that influence 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.[27] 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 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 regular nature of these managers' appointments is highlighted by the fact that both the modal and median appointment duration is 36 months – three years. Appendix Figure A.2 plots the distribution of appointment durations. Between appointments, managers 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. Appendix Figure A.3 plots the distribution of gaps between appointments.

The rotation schedule provides the most direct reason for changes in the bureaucrat who manages a country office. Further, this rotation limits discretion in appointments as only a limited set of bureaucrats are available when a particular country is due to receive a new bureaucrat. More importantly, the appointment of bureaucrat b to country c in year t largely pre-determines that bureaucrat b will be replaced in year $t + 3$. Section 4.1 discusses two key implications of this that allow us to test the plausibility of our identifying assumption that bureaucrat appointments are orthogonal to underlying export trends in the country of appointment.

[27] Before becoming office head, bureaucrats usually see multiple appointments to overseas offices in subordinate roles.

2.4 KOTRA and Korea's Largest-Scale Industrial Policy

One reason for studying export promotion is the narrative of South Korea's development as being export-driven, as well as exports' prominent role in South Korean industrial policy. Korea's largest-scale industrial policy, the heavy and chemical industries drive (HCI), commenced in early 1973 and ended in October 1979.

Appendix Figure A.4 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.[28] 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 national sectoral industrial policies, such as the HCI drive, targeted particular sectors. As we study variation between destination countries, our variation should be orthogonal to policies targeting particular sectors or regions within South Korea. To account for the presence of such policies, and differential growth between sectors more broadly, our regressions include product-year fixed effects.

3 Data

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

3.1 Bureaucrat Appointments

The most complete source of bureaucrat appointments is given by contemporaneous reports on appointments to KOTRA's overseas offices published in major South Korean newspapers. In most

[28] To show the connection between KOTRA and HCI, we linked about 45,000 of the reports written by KOTRA's overseas offices between 1965 and 2001 to the products or sectors discussed by each report. When discussing whether a product was treated by HCI, we follow Lane (2025), who included those "listed in the enforcement decrees and national sectoral acts underlying HCI". HCI's six broadly defined target sectors included steel, nonferrous metals, shipbuilding, machinery, electronics, and petrochemicals.

years, there were two main dates at which appointments were announced, usually in January and July. The actual appointments most frequently started in April and October. Further, this information is usually reported in three major newspapers (*Dong Ah Ilbo*, *Choson Ilbo*, and *Kyonghyang Sinmun*). Because of these overlapping information sources, there are almost no rounds of announcements that were not reported by any newspaper. For almost all rounds of announcements, we were able to corroborate the information using at least two of these sources.

The newspaper announcements are further complemented and verified using a variety of KOTRA publications on the manager in charge of an office at a given point in time. We obtained and digitized the names of bureaucrats in (i) monthly publications aimed at non-Korean importers (1966-1971), (ii) a directory of KOTRA's network including all of its overseas office managers (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 that existed between 1962 and 2001, located in 87 distinct countries. 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 further in section 4.

Managers are identified using their names, which requires us to avoid two types of errors. First, we may erroneously code two bureaucrats as the same one, e.g., it may be that bureaucrats share names. A priori, this could have been a problem as 45% of bureaucrats in our sample share the last names *Kim*, *Lee*, and *Park*.[29] However, this is remedied by a great diversity in first names.[30] After a plethora of checks, it appears very unlikely that any bureaucrats in our data share the exact same full name. More challenging in practice, we have to determine whether slightly different names truly refer to distinct individuals. This task is complicated as over time our sources move

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

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

from Chinese to Korean characters to render the bureaucrats' names. In addition, in the few cases where names are given using romanizations, inconsistent romanization is used, e.g., both *rhee* and *lee* to render the same Korean syllable. We resolve this challenge in four steps: Identify incorrectly spelled or digitized names by (1) matching very unusual names to more common ones, (2) harmonizing the rendering of certain syllables, e.g., *rhee* and *lee*, (3) identifying offices with likely mistakes, e.g., the manager's name flips back and forth, and (4) re-creating the career of each bureaucrat and assessing patterns of overlap or missing years. Following these steps meticulously allowed us to create a consistent panel of unique bureaucrats covering all offices and all years.

3.2 Exports

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

3.3 Bureaucrat Output

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

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

4 Do Bureaucrats Drive Korean Exports?

This section finds that the effect of a South Korean overseas export promotion office differs substantially depending on the bureaucrat managing it: Increasing the ability of the bureaucrat 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 under study would have no effect if every bureaucrat's ability is reduced by one standard deviation.

4.1 Identifying Bureaucrat Fixed Effects

We adapt the two-way fixed effects framework (Abowd, Kramarz, and Margolis, 1999) to study how much bureaucrats matter in explaining South Korean exports. We proceed in two steps. In the current section, we obtain unbiased estimates of bureaucrat fixed effects and check that identifying assumptions are met. In Section 4.2, we use the estimated fixed effects to obtain measures of 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)[31] 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 ($\theta_{b(c,t)}$),[32] a country component (γ_c), and an error term (ϵ_{cpt}). As

[31] For intensive margin changes, the IHS behaves similarly to the natural logarithm. Contrary to the logarithm, the IHS is defined at zero. Export promotion affects the extensive margin, leading us to prefer the IHS transformation for our benchmark, while exploring robustness to alternative transformations. We show in Appendix C.2 that bureaucrat fixed effects based on the IHS are predictive of export changes along both the extensive and intensive margins. Furthermore, we show in Appendix C.3 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).

[32] 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

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.[33] Additionally, it requires that manager 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*. On the other hand, this orthogonality condition allows for manager assignment to offices on the basis of the permanent component of country effects γ_c or the permanent component of manager 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, would not violate this identifying assumption.

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. [34] The largest connected set thus contains 704 appointments of 397 bureaucrats, of whom 184 had appointments to multiple offices.[35]

The bureaucrat-country graph is interconnected enough such that 75 countries and 93% of appointments form a leave-one-out connected set, i.e., a set that remains connected if dropping one appointment from the data, no matter which one. The reason is that most country offices remain open for decades over which they are headed by many different office managers. Column 3 indicates that in the leave-one-out connected set, 72 offices are managed by four or more distinct bureaucrats, 61 offices have six or more bureaucrats, and 49 offices even have eight or more distinct

to nine months after their successor has been appointed.

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

[34] Only Zimbabwe is outside the largest connected set. The data only contains one appointment to Zimbabwe because its office opening occurred shortly before the end of our sample period.

[35] 184 happens to be the same number as in the balanced analysis sample of Fenizia (2022). In Fenizia (2022), these moves result in 143 distinct connected sets – while in our case, the moves connect all offices in one single connected set (except for Zimbabwe).

bureaucrats.

Our preferred estimation uses only the appointments in the largest leave-one-out connected set. The leave-one-out connected set allows us to apply Kline, Saggio, and Sølvsten (2020)'s correction method for limited mobility bias in our variance decomposition exercise, which allows us to distinguish bureaucrat effects from noise as corroborated by our placebo exercise of reshuffling bureaucrats.[36]

Identifying Assumption: Appointments are Orthogonal to Export Trends

This section discusses how factors influencing bureaucrat appointments relate to the identifying assumption that appointments are orthogonal to export trends.

The central factor driving the movement of bureaucrats is their three-year rotation schedule. As highlighted in Appendix Figure A.2, a new appointment to country c in year t usually occurs if the previous bureaucrat's appointment to country c occurred in year $t-3$. This has two important implications that allow us to investigate the assumption that bureaucrat appointments are as good as random with respect to export trends.

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 timing all such appointments to precisely the years when exports start to rise. Then, the increasing trend would have materialized before some appointments, once again leading to differential pre-trends by bureaucrat ability.[37] Section 4.4 tests this prediction in an event-study setup estimating the effect on exports of switches between bureaucrats. We find parallel pre-trends, alleviating concerns that KOTRA may appoint good bureaucrats to countries with increasing export trends.

Second, while KOTRA has some discretion in new appointments, the rigid rotation schedule

[36] Kline, Saggio, and Sølvsten (2020) also show that restricting attention to the leave-one-out connected set substantially reduces limited mobility bias directly.

[37] This would be the case even if there are exceptions to the rotation schedule – as the rotation schedule is binding for most rotations.

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 compared to the decrease when losing a bureaucrat of the same ability. Importantly, this prediction only requires that discretion is greater at the end 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 assuages 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 identifying 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 the reduced discretion regarding when a bureaucrat's term ends.

In addition to these testable implications, the rotation schedule also constrains the pool of bureaucrats available when a country office needs a new manager, restricting the potential for strategic appointments. From qualitative interviews with KOTRA employees and our analysis, several constraints emerge: (1) Bureaucrats are more likely to be appointed to a country when they speak the local language. (2) Bureaucrats prefer being appointed to high-income, English-speaking countries. (3) Because these preferences are commonly 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, further discussed in section 4.6. In most cases, a country's language and its income, importance, and desirability, both relative to other countries, change little over time. Therefore, while appointments may be correlated with country fixed effects,[38] the scope for appointing bureaucrats based on (anticipated) export trends is very limited.

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

Despite these constraints, one may still wonder why KOTRA does not strategically appoint good bureaucrats to countries with growing import demand. An important reason may be that KOTRA prioritized time-invariant country characteristics – which are very large – over trends in gauging export potential, much like it did for office openings.[39]

4.2 The Share of Variation in Exports Explained by Bureaucrats

This section explains how we decompose the variance in South Korean exports to estimate how much of it is due 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).[40] The remaining variation in exports can be 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 in the estimation of the bureaucrat or country fixed effects, we take the spell-level averages of the residualized exports as the total variation.[41]

Our primary object of interest is the variation explained by the bureaucrats: $\mathsf{Var}(\theta_{b(c,t)})$. The challenge in obtaining an estimate for $\mathsf{Var}(\theta_{b(c,t)})$ 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

[39] The rollout of export promotion offices mainly followed pre-determined gravity variables, as reported in Figure 12 and discussed in Section 5.3.

[40] The calculation of $(\text{exports}|pt)_{cpt}$ follows Chetty, Friedman, and Rockoff (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.

[41] The bureaucrat and country fixed effects estimated from the collapsed data equal those estimated from the uncollapsed, raw data. The variance of the raw (i.e., country×product×year-level) residualized exports is also reported in Table 2 for reference.

move countries too infrequently. Such *limited mobility bias* in our setting is illustrated in Appendix Figure A.5.

Our preferred variance decomposition approach follows the bias correction method of Kline, Saggio, and Sølvsten (2020). Kline, Saggio, and Sølvsten (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.[42]

One downside of the analysis based on Kline, Saggio, and Sølvsten (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 alternative approach involves shrinking each raw fixed effect to minimize its mean-squared-error.[43] We use covariance shrinkage, developed by Best, Hjort, and Szakonyi (2023), which extends the standard shrinkage method (Chetty, Friedman, and Rockoff, 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[44] 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

[42]The Kline, Saggio, and Sølvsten (2020) correction method can only be performed on the leave-one-out connected set which covers 75 countries and 380 bureaucrats, 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, Saggio, and Sølvsten (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.

[43]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.

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

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 E for details). The bootstrap bias correction allows us to relax an assumption in Kline, Saggio, and Sølvsten (2020) that is used to impute the variance of the sampling error for single-appointment bureaucrats,[45] since bootstrapping equation (2) yields variance estimates for all bureaucrats.

4.3 Result: Bureaucrats Drive Korean Exports

Figure 3 plots the cumulative distribution function of raw bureaucrat fixed effects, showing substantial dispersion.[46]

Table 2 reports the variance decomposition results, correcting for bias using the Kline, Saggio, and Sølvsten (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 variance due to bureaucrats is remarkably similar. The fact that both correction methods 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 appointment level. Furthermore, column 1 shows that one standard deviation of bureaucrat ability is estimated to be 0.318,[47] implying an

[45] 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, Saggio, and Sølvsten (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.5 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.

[46] While raw fixed effects are unbiased under the identifying assumptions, sampling error would lead to greater dispersion than that present in underlying abilities. Appendix Figure A.6 compares the distributions of the raw fixed effects and the covariance-shrunk fixed effects, highlighting that much of the dispersion in Figure 3 is due to differences between bureaucrats – it remains even after shrinkage to account for sampling error. Appendix Figure A.6, for example, shows that covariance shrinkage reduces the difference between the 25th and 50th percentile by 30% and between the 50th and 75th percentile by 18%.

[47] $0.318 = \sqrt{0.101}$

increase in the dollar value of exports of 37%.[48] 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.[49]

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.[50] 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 co-

[48] We approximate the percentage change using $(e^{0.318} - 1) \times 100 \approx 37$.

[49] Under the simplification that the estimated effect of office opening reflects the office's true effect under an average bureaucrat.

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

variance is still negative but is smaller and statistically insignificant. The negative covariance is almost exclusively driven by single-appointment bureaucrats.[51] For multi-appointment bureaucrats, 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 evidence supporting the identifying assumption, discussed in detail in Section 4.1. When this assumption holds, our bureaucrat fixed effect estimates are unbiased, but they 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 C, corroborating the additive separability assumption of bureaucrat and country effects, as well as the robustness to different transformations of export value.

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, as well as asymmetric effects of gaining and losing a bureaucrat.

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

$$y_{ept} = \eta_{ep} + \lambda_{pt} + \sum_{k \neq -2} \Big(\alpha_k + \beta_k \cdot \mathbb{1}\{\Delta\hat{\theta}_e \text{ in top tercile}\} + \delta_k \cdot \mathbb{1}\{\Delta\hat{\theta}_e \text{ in mid tercile}\} \Big) \mathbb{1}\{t = T + k\} + \epsilon_{ept} \tag{4}$$

First, to test for differential pre-trends, we estimate equation (4), which models how exports change around an event (e), a switch from one bureaucrat to another. An event e is uniquely defined by the country (c) and the year of the event (T), defined as the first full year that the new bureaucrat is appointed to country c. We estimate how exports change for different terciles of the change in bureaucrat ability ($\Delta\hat{\theta}_e$).[52] The key parameters of interest are β_k and δ_k. β_k (δ_k) is the double-difference of exports in year $T + k$ relative to year $T - 2$, for a change in bureaucrat ability in the top (middle) tercile relative to a change in the bottom tercile. 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}).

Figure 4 plots the estimates $\hat{\beta}_k$ and $\hat{\delta}_k$. Prior to the bureaucrat switch, there are no discernible trends. The point estimates are close to zero and statistically indistinguishable from zero up to the year of the bureaucrat switch. The lack of differential pre-trends corroborates the identifying assumption. At the same time, we do see jumps, over 1-2 years, in exports upon the appointment of the new bureaucrat, and more so for the top transition (30%) than the middle transition (11%), adding confidence that our bureaucrat fixed effects meaningfully capture the differences in bureaucrat abilities.

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

Second, to test whether gaining a bureaucrat coincides with larger increases in exports than the decreases associated with losing one, we estimate equation (5). It models exports as a time-varying function of the fixed effects of the new bureaucrat ($\hat{\theta}_e^{\text{new}}$) and the old bureaucrat ($\hat{\theta}_e^{\text{old}}$). Other than

[52] We consider terciles, following the literature (Fenizia, 2022; Otero and Muñoz, 2025).

distinguishing between $\hat{\theta}_e^{\text{new}}$ and $\hat{\theta}_e^{\text{old}}$, the specification follows equation (4). As argued in section 4.1, the rotation schedule means we would expect stronger "effects" in absolute terms due to the ability of a new bureaucrat than due to the old bureaucrat.

Figure 5 plots the event-study estimates, $\hat{\beta}_k$ and $\hat{\delta}_k$, obtained from equation (5). 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.[53] In addition, the pre-trends in this specification, like in Figure 4, 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 more than 35% annually[54] 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.

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.

[53] Fenizia (2022) notes the possibility that idiosyncratic shocks to exports in ϵ_{ept} could also appear as finite-sample estimation error of the bureaucrat fixed effects, creating a spurious correlation in the post-period even in the absence of a causal relationship. While this is less concerning regarding the pre-trends, it may be relevant for comparing changes to exports after the event. Fenizia addresses the problem by estimating fixed effects for each value of k using data that excludes k. In Appendix Figure A.7, 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, reported in Appendix Figure A.7, continue to support the identifying assumption. The increase in exports when gaining a bureaucrat is not larger than the decrease when losing them.

[54] Between 1962 and 1981, South Korean exports increased from 55 million to 21 billion U.S. Dollars, implying an annual growth rate of 36.7%.

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.[55] 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.[56] 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, Saggio, and Sølvsten (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 fixed effects. An implication of overfitting is that the fixed effect estimates would have low predictive power outside the sample. We directly test and confirm that our estimated bureaucrat effects are, indeed, predictive out of sample.[57] 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

[55] 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.

[56] 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.

[57] 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.

large connected set.

Figure 6 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.[58] Furthermore, we replicate the event-study regression depicted in Figure 5 using the out-of-sample fixed effects. Appendix Figure A.7 shows that upon a switch of bureaucrats, new and old bureaucrat ability still statistically significantly predict exports in the expected way 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.[59]

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 the Korean ambassadors – who also follow rotational assignments – to each country with a KOTRA overseas office.[60] 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

[58] This value is very close to the relationship between a retail store manager's pre-Covid and Covid performance (in the same store) found by Metcalfe, Sollaci, and Syverson (2023), who study managers of retail stores.

[59] In addition to addressing overfitting concerns, Appendix Figure A.7 also displays a lack of pre-trends, in support of the identifying assumption of the bureaucrat fixed effects.

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

the causal impact on exports of an individual bureaucrat managing a KOTRA office.

4.5 Mechanism: Good Bureaucrats Tap into Import Demand

In this section, we investigate how a high-ability bureaucrat increases exports compared to other bureaucrats. We find that they raise the reactivity of Korean exports to market conditions. Upon the switch to a more effective bureaucrat, Korean exports increase more strongly for products that see increasing import demand in a given country-year. They also increase more strongly for products that see increasing export supply from Korea to other destination countries. Our findings suggest that much of the effect of high-ability bureaucrats stems from their ability to better exploit market conditions, for instance by relaying timely information about destination market demand.

$$
\begin{aligned}
y_{ecpt} =& \eta_{ep} + \lambda_{pt} + \psi^D \text{demand}_{cpt} + \psi^S \text{supply}_{cpt} \\
&+ \left(\beta^D \text{demand}_{cpt} + \beta^S \text{supply}_{cpt}\right) \times \mathbb{1}\{\Delta\hat{\theta}_e \text{ in top tercile}\} \\
&+ \left(\delta^D \text{demand}_{cpt} + \delta^S \text{supply}_{cpt}\right) \times \mathbb{1}\{\Delta\hat{\theta}_e \text{ in mid tercile}\} \\
&+ \sum_{k \neq -2} \Big[\alpha_k + \psi_k^D \text{demand}_{cpt} + \psi_k^S \text{supply}_{cpt} \\
&+ \left(\beta_k + \beta_k^D \text{demand}_{cpt} + \beta_k^S \text{supply}_{cpt}\right) \times \mathbb{1}\{\Delta\hat{\theta}_e \text{ in top tercile}\} \\
&+ \left(\delta_k + \delta_k^D \text{demand}_{cpt} + \delta_k^S \text{supply}_{cpt}\right) \times \mathbb{1}\{\Delta\hat{\theta}_e \text{in mid tercile}\}\Big] \mathbb{1}\{t = T + k\} + \epsilon_{ecpt}
\end{aligned}
\tag{6}
$$

We estimate equation (6), an event-study regression around the event of bureaucrat switches. It augments equation (4) with the main and interacted effects of import "demand" and export "supply." demand_{cpt} equals the IHS of other countries' exports of the same product p to the same destination country c in year t, while supply_{cpt} equals the IHS of Korean exports of the same product p to destination countries other than c in year t. ψ^D and ψ^S estimate the elasticity of Korean exports with respect to market conditions in the base year, $T-2$. β^D, β^S (δ^D, δ^S) estimate the additional elasticity when the change in bureaucrat ability falls in the top (middle) tercile, relative to the bottom tercile, in year $T-2$. ψ_k^D and ψ_k^S allow the elasticity to market conditions to differ by event year, and $\alpha_k, \beta_k, \delta_k$ allow exports to differ by event year, for each tercile. The new

parameters of interest are $\beta_k^D, \beta_k^S, \delta_k^D$, and δ_k^S. They capture the double-difference of the elasticity of Korean exports to market conditions in year $T + k$ relative to year $T - 2$, for a change in bureaucrat ability in the top (or middle) tercile relative to a change in the bottom tercile.

Figure 7 plots the estimates $\hat{\beta}_k^D$ and $\hat{\delta}_k^D$ in the left panel and $\hat{\beta}_k^S$ and $\hat{\delta}_k^S$ in the right panel. We find that upon the arrival of a new bureaucrat, the elasticity of Korean exports to market conditions rises more when the increase in bureaucrat ability is greater. The patterns are mostly similar whether we consider the elasticity to import demand or to export supply. Thus, high-ability bureaucrats appear to conduct a more effective transmission of market information.

Lastly, the estimated coefficients in the pre-period are statistically insignificant and much smaller in magnitude than those in the post-period. The parallel pre-trends in this figure, similar to Figure 4, support the orthogonality of bureaucrat appointments and export trends.

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. Then, selecting good talent becomes crucial for strengthening bureaucratic capacity. What could be a viable way for the state to select high-ability bureaucrats, when there is incomplete information on their ability, perhaps 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 8 presents a visual illustration. It plots the probability density function of (a) residualized exports by appointment, and (b) bureaucrat fixed effects (covariance-shrunk), 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.[61]

[61] We use covariance-shrunk bureaucrat fixed effects to correct for the fact that the raw fixed effects of single-

This visual observation is buttressed by a regression of bureaucrat's total number of appointments on residualized exports during their first appointment. We include dummies for the year of a bureaucrat's first appointment to rule out various omitted variable biases, such as the number of appointments mechanically being smaller for bureaucrats whose first appointments occur near the end of our sample, and systematic differences in the characteristics of bureaucrats by cohort. We find a positive and statistically significant correlation, with a regression coefficient of 0.240 (standard error: 0.112). The positive correlation survives alternative specifications. For instance, the coefficient increases to 0.430 (standard error: 0.109) when we replace first-appointment residualized exports with a dummy indicating whether they exceed 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 9 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 and a country's fixed effect in equation (1). Panel (a) shows that first appointments make up the largest share of appointments to offices that opened after 1975, i.e., late or less important offices. On the other hand, third appointments constitute the lion's share of appointments to offices that opened before 1968, i.e., early or important offices. Second appointments form an intermediate case. Panel (b) paints the same 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 those first appointments to select out low performers from getting reappointed, particularly to important countries.[62] An appointment as a manager of a small

appointment bureaucrats are noisier. Furthermore, Panel (a) shows that among multi-appointment bureaucrats, there is a very close resemblance between the distributions of residualized exports in the first appointment (solid blue) and in subsequent appointments (dashed blue). This strengthens the case that the difference in single- and multi-appointment bureaucrats reflect differences in true ability rather than noise.

[62]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

country office is plausibly much more informative about a bureaucrat's ability than a subordinate role in a more important country. Therefore, assigning bureaucrats to manage small countries may be justified, even if placing them in subordinate roles in key markets could have a greater impact on overall Korean 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 3 provides three simple counterfactuals to understand 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 might be interpreted as the prevailing "screening ability" of KOTRA.

To investigate the effect of the current degree of screening, counterfactual (1) simulates the change in South Korean exports in the absence of screening. Specifically, we remove the difference in ability between single- and multi-appointment bureaucrats by setting both to the average ability across all bureaucrats. In this counterfactual, the dollar value of exports for appointments led by single-appointment bureaucrats increases by 11%, 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. This scenario may be relevant if KOTRA could operate a new set of offices in smaller countries to learn about bureaucrats before assigning them to the countries in our data. The estimated difference in ability implies a 25% increase in export value for appointments headed by single-appointment bureaucrats, from replacing them with the average multi-appointment bureaucrat. As around 30%

Resources considerations.

of appointments are headed by single-appointment bureaucrats, the total export value across all appointments would rise by 6.8%, again a substantial effect. 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 observed screening ability.

Counterfactual (3) replaces bureaucrats in the bottom quartile of ability with the median bureaucrat. This counterfactual is based on individual points in the distribution of bureaucrat fixed effects, as opposed to the average ability of groups of bureaucrats. We use covariance-shrunk fixed effects to account for limited mobility bias since different bureaucrats' fixed effects are estimated with different degrees of sampling error. We find that the value of exports would increase by 8.3%. So the gains from KOTRA's screening ability – highlighted by counterfactuals (1) and (2) – are more than 60% compared to this last counterfactual that is not based on observed screening ability.

These counterfactuals highlight the importance of screening processes in strengthening bureaucratic capacity and suggest that improvements in bureaucratic selection could yield substantial gains in policy effectiveness.

5 What is the Policy's Average Effect?

In this section, we use the staggered rollout of overseas export promotion offices to identify the causal effect of the offices on Korean exports to the respective destination countries.[63] The average effect of export promotion offices is interesting and policy-relevant on its own. More importantly for this paper's main research question, it serves as a natural benchmark to compare the effect of individual bureaucrats, estimated in Section 4, against. We estimate a 38% increase in the value of exports 9-11 years after an office opens, comparable to the effect of a one 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, it follows that an office with a bureaucrat one standard deviation below average would have no effect on exports.

[63] For multi-office countries, we focus on the opening of the first office.

Our setting is exceptional in enabling this natural comparison between a policy's average effect and how this effect changes due to implementation by individual bureaucrats.[64] The comparison is possible because (1) we observe a sufficient number of office openings, and (2) South Korean exports to a country are a well-defined variable even without an office.

The staggered rollout of offices is depicted in Appendix Figure A.1, which plots the cumulative number of offices by year. KOTRA was established in 1962. Over the next two decades, offices opened in 75 countries – nearly four new countries every year. After the intensive initial rollout, KOTRA's expansion came to an abrupt halt, with only three new countries with a first office opening during the next seven years.[65] The empirical analysis in this section focuses on the initial office openings (1962-1981). Figure 2 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 and not others. As this is not feasible, the analysis here will use the staggered rollout of offices to countries. We allow for dynamic effects, since it may take some time for the offices to consolidate.

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

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.[66] 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

[64] Settings from related papers do not lend themselves to obtaining such a benchmark (Fenizia, 2022; Best, Hjort, and Szakonyi, 2023; Otero and Muñoz, 2025; Metcalfe, Sollaci, and Syverson, 2023).

[65] Only 13 openings in total over the next twenty-year period (1982–2001).

[66] We prefer country-product fixed effects to simpler country fixed effects because the former allow us to account for additional variation unrelated to the office opening's effect. In practice, estimates are largely unaffected by this choice.

any controls.

Two central assumptions need to be met for $\hat{\pi}_k$ to be unbiased estimates of the causal effect of the office, k years since opening. The first is a parallel trends assumption. We assume that counterfactual export trends, in the absence of an office, do not differ in periods $g + k$ (where $k > 0$) between countries treated in year g and the control.[67] There is little indication of differential pre-trends (discussed in Section 5.2), and offices do not affect non-Korean exports to the country after they have opened (discussed in Section 5.3), supporting the parallel trends assumption. 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. Yet, we show that the rollout order of European offices is almost fully explained by pre-determined (1962) import market size, which suggests that there was little room to time office openings in this way, 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 it is located in and has no spillovers on exports to other countries. If this assumption is violated, the estimated effect may reflect reallocations of export flows across destination countries, rather than isolating the absolute effects on exports to a given country. Concerns regarding this assumption are somewhat alleviated by the fact that we find similar effects whether we use a never-treated or not-yet-treated control group, as SUTVA violations would likely affect these two groups differently. Further, SUTVA violations would be most concerning if they caused an upward bias in the estimated effect of offices. Alfaro-Ureña, Castro-Vincenzi, Fanelli, and Morales (2023) provide some justification to believe SUTVA violations, if they exist, would bias our results downwards, not upwards. They find that exports to different countries are complements. In that case, violations of SUTVA would more likely lead to underestimates.

A third assumption, no anticipation, is required for identification. This assumption would be

[67] 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).

violated if office openings had a causal effect at $k < 0$. Negative anticipation – an Ashenfelter (1978) dip – would lead to overestimates. We find no indication of anticipation.[68]

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 (de Chaisemartin and D'Haultfœuille, 2020; Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; Sun and Abraham, 2021; Borusyak, Jaravel, and Spiess, 2024; Rambachan and Roth, 2023). For example, the specification would produce biased estimates if there are dynamic treatment effects and already-treated units are included in the control group. For this reason, 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.[69] Figure 10 Panels (a)-(b) shows that the pre-trends remain near-zero and effect sizes remain comparable whether 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.

For countries that experienced an office opening between 1967 (1964) and 1981, we include all observations from five (two) years before to eleven years after the office opening. Hence, the

[68] Positive anticipation may occur if firms decide to export to a market today because KOTRA will open an office there next year. Under the parallel trends assumption, positive anticipation is part of KOTRA's causal effect and would give rise to an underestimate. In all specifications with the never-treated control group, there is no indication of any positive anticipation.

[69] Territories with first office opening in 1962: U.S., Thailand, Taiwan. 1963: none. 1964: Japan, Singapore, Indonesia, South Vietnam. 1965: Philippines, Peru, Kenya, Hong Kong, Iran, UK. 1966: Italy, Netherlands, Panama, Nigeria.

earliest start year for a treated country's event horizon is 1962 while the latest end year is 1992. For the never-treated control, we include all observations between 1962 and 1992. λ_{pt} is estimated from the product-year variation in the control group.[70]

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,[71] adds support that the never-treated group is comparable. Furthermore, Appendix G.1 finds that the coefficients remain similar with an alternative control group, the not-yet-treated countries, using the estimator proposed by Callaway and Sant'Anna (2021).

5.2 Results: Effect of Office Opening on Exports

Figure 10 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 coefficients in the pre-periods are economically small and not statistically different from zero, mitigating concerns about a violation of 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%[72] higher export value relative to the control group in the higher-powered regression (Panel (b)).

The effect size is large. To put it into perspective, we again use the effect of distance on trade as a comparison point. With an elasticity of trade to distance of -1 (Anderson, 2011; Head and Mayer,

[70] To ensure a balanced panel, we exclude never-treated countries with gaps in the export data between 1962 and 1992.

[71] Besides the main outcome variable of Korean exports, we also consider the following outcomes: non-Korean exports (as a placebo), a dummy indicating positive Korean export value, the number of KOTRA market reports, and the number of import inquiries.

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

2014), an office opening corresponds to the effect of reducing trade distance from London-Seoul (8,900km/5,500 miles) to Mumbai-Seoul (5,600km/3,500 miles).

The opening of an office increases KOTRA activity specific to the relevant country. Appendix Figure A.9 reports the results from estimating equation (7) as before but replacing the outcome variable with three measures of 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. Like exports, they are each transformed by the IHS to account for zeros in the data. We collected these data by digitizing KOTRA's Daily Market Newspaper. We confirm that the pre-treatment coefficients are statistically indistinguishable from zero for all panels, again supporting the parallel trends assumption. Panel (a) shows that from the year of the opening, the IHS of the number of reports on the country increases, and is already stable at around 1.2 from one year after the opening, implying a 3-fold increase in the number of reports. Panel (b), showing the results for the subset of reports that are product-specific,[73] are very similar to Panel (a). Panels (c) and (d) plot the results for import inquiries, which list the contact details and product category of interest of potential importers in the destination country. Panel (c), with just one pre-period, allows for a wider treatment window (openings between 1975 and 1981). It shows an increase in the number of inquiries from importers of the country of the office. Panel (d), covering openings between 1978 and 1981, also shows an increase overall, although the coefficients are noisier and less statistically significant individually.

5.3 Robustness and Validity Checks

This section discusses the possibility of violations of the parallel trends assumption, beyond checking for the absence of differential pre-trends.

Other robustness checks are performed in Appendix G. We discuss in detail the robustness of the results to (1) using a not-yet-treated control group following Callaway and Sant'Anna (2021),

[73] Product-specific and general reports on a country are both informative to potential Korean exporters. Example of a product-specific report title: "TV/Video Recorder/Stereo Market Profitability in West Germany is Worsening." Example of a general report title: "Argentina Abolishes the Import Quota System."

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 strategically be opened at a time when import demand in that country is expected to increase, violating 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 Figure 11.[74] The point estimates are close to zero both before and after an office opening, with wide confidence intervals, implying that office openings do not coincide with increases in import demand. Second, we use non-Korean exports as a control variable while maintaining Korean exports as the outcome. Panels (a) and (b) of Appendix Figure A.8 show that the estimates from this specification are largely unchanged compared to the baseline. Given South Korea's rapid economic growth, it may be that the relationship between Korean and non-Korean exports changes over time.[75] Panels (c) and (d) show that estimates remain similar when the effect of non-Korean exports is allowed to vary by year, by controlling for non-Korean exports interacted with year dummies.

Rollout Follows Pre-Determined Gravity Variables

This section shows that the year in which a country's first office opened was largely pre-determined by time-invariant factors. If the effect of these factors on exports is also time-invariant, they are absorbed in the country fixed effects (γ_c). Even if the effect of these factors is not time-invariant, the pre-determined order of the roll-out makes it unlikely that office openings were timed to coincide with counterfactual increases in exports.

To predict office openings, we draw on insights from the gravity equation, where distance and

[74] The coefficients are robust to alternative event horizon lengths, which alters the sample to include the earliest office openings (Panels (e)-(f) of Appendix Figure A.8).

[75] While non-Korean exports are positively associated with Korean exports on average, controlling for them does not substantially change the coefficients as non-Korean exports do not change systematically following an office opening.

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.[76]

Within Europe, the distance from South Korea does not vary much between countries,[77] so the main predictor would be market size.[78] 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. Figure 12 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 roll-out 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 rank correlation between 1962 imports and year of office opening is very strong, at 0.87. The rigidity in the roll-out schedule counters concerns that KOTRA timed openings based on export trends, which would have violated the parallel trends assumption.

To sum up, the evidence of no increase in non-Korean exports upon office opening, combined with evidence that the timing was determined by stable factors rather than anticipation of export increases or market-specific shocks, counters concerns that KOTRA timed openings based on export trends. These findings support the causal interpretation of our estimated office effects.

6 Conclusion

This paper shows that bureaucratic capacity can fundamentally change the effect of an industrial policy on exports. 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 state capacity and the East Asian miracle, advancing our understanding of how state involvement might influence macro development trajectories.

[76] North Korea, China, the USSR, and North Vietnam were ideological opponents of South Korea.
[77] The distance between Athens and Seoul is 96% of that between London and Seoul.
[78] Other determinants of trade, such as language or cultural distance, also vary little within Europe.

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, Lane, Oehlsen, and Pérez, 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 significantly on concurrent investments in bureaucratic capacity, such as selection mechanisms that identify capable individuals and organizational structures that empower their decision-making while maintaining accountability.

Figure 1: Growth in Korean Exports

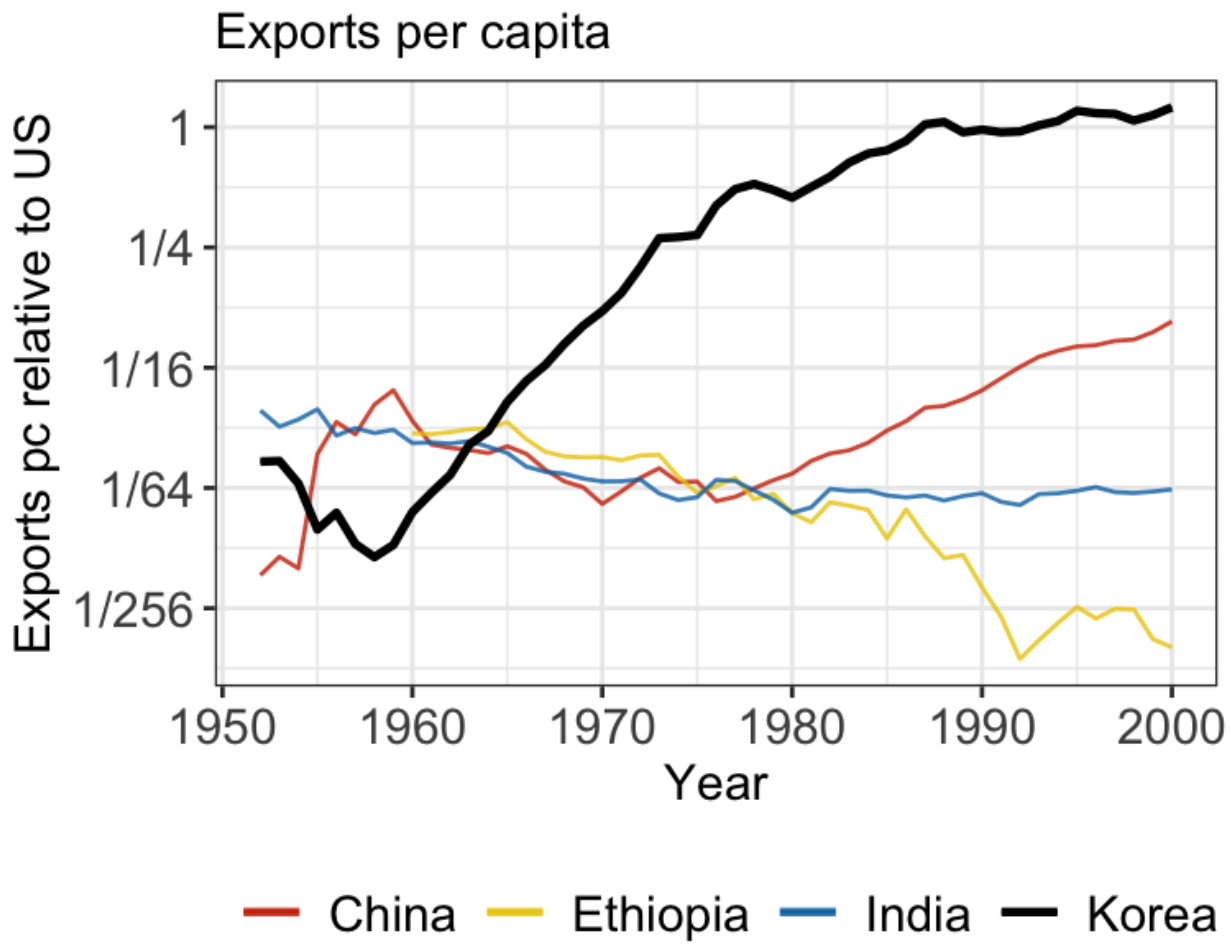

Notes: The figure displays exports per capita relative to the U.S. from 1952 to 2000, for South Korea and a selected group of other countries. Data on exports and population are obtained from International Monetary Fund (2023): Direction of Trade Statistics. Back to Section 2.

Figure 2: Korean Export Promotion Offices

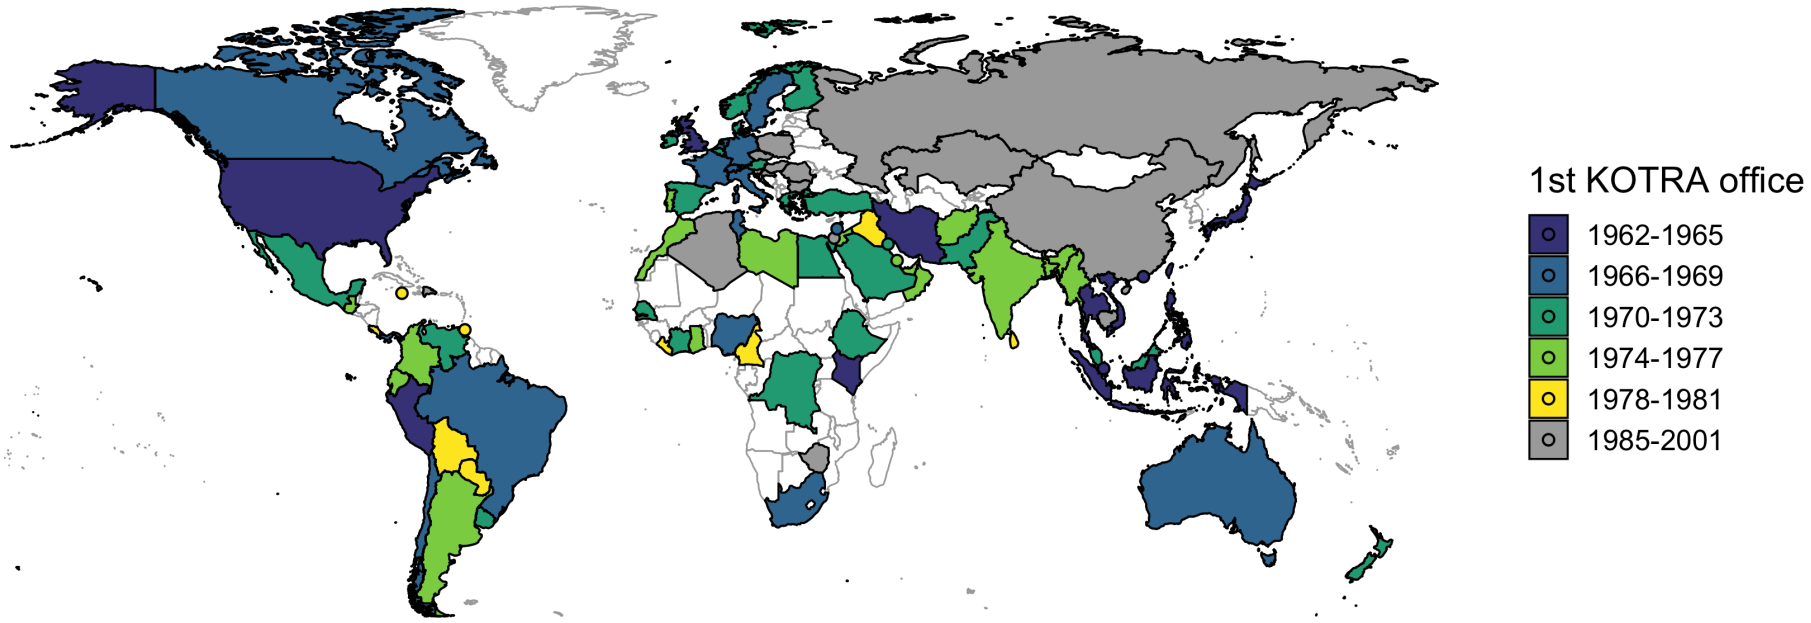

Notes: Colors indicate the opening year of the first Korean export promotion office. Most offices remain open from the opening year for the entire period of study. White indicates territories which have not received an office by 2001. Back to Sections 2.2 and 5.

Figure 3: Distribution of Raw Fixed Effects of Managers of Country Offices

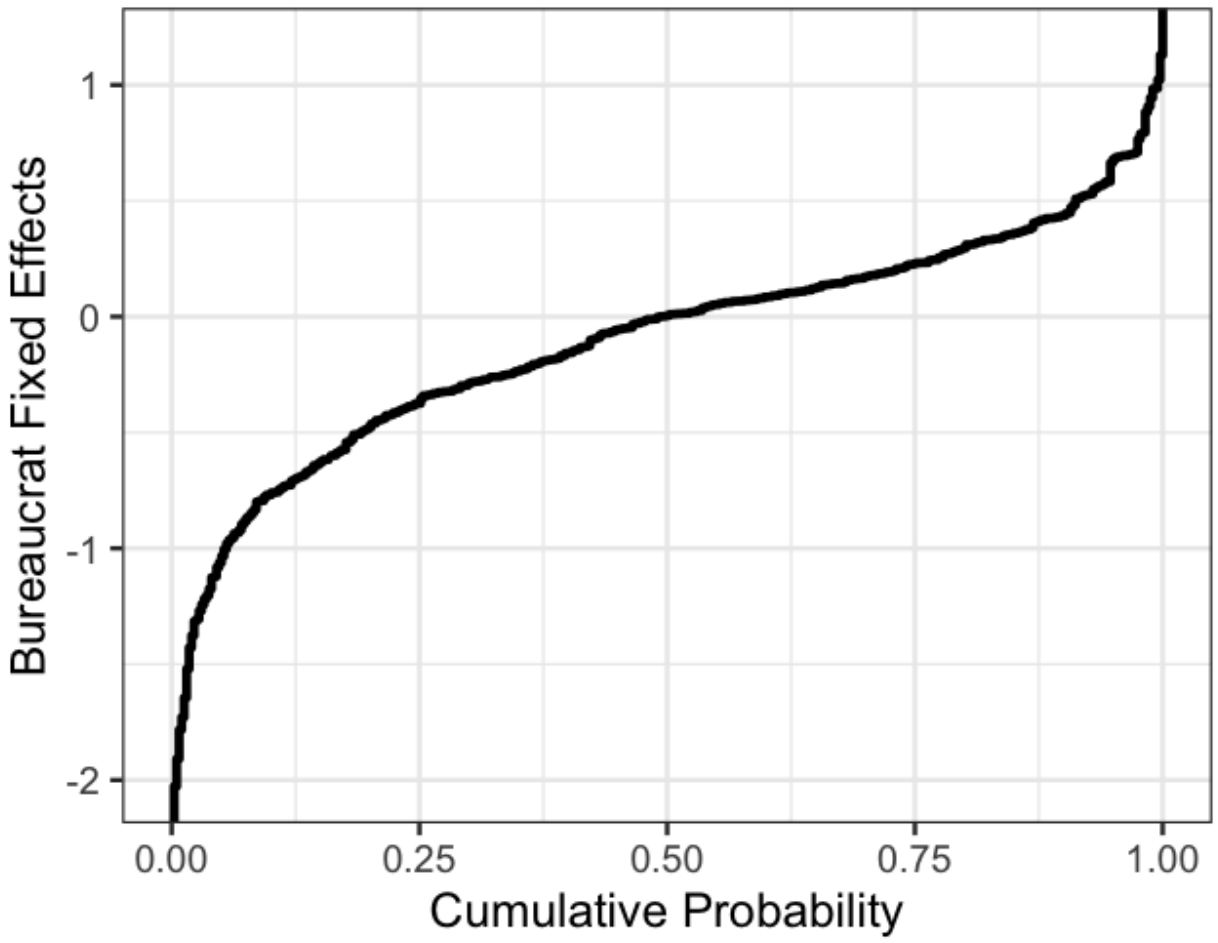

Notes: The figure shows the cumulative distribution function of bureaucrat fixed effects, estimated from equation (2). Back to Section 4.3.

Figure 4: Switches to Better Bureaucrats Are Not Preceded by Differential Trends

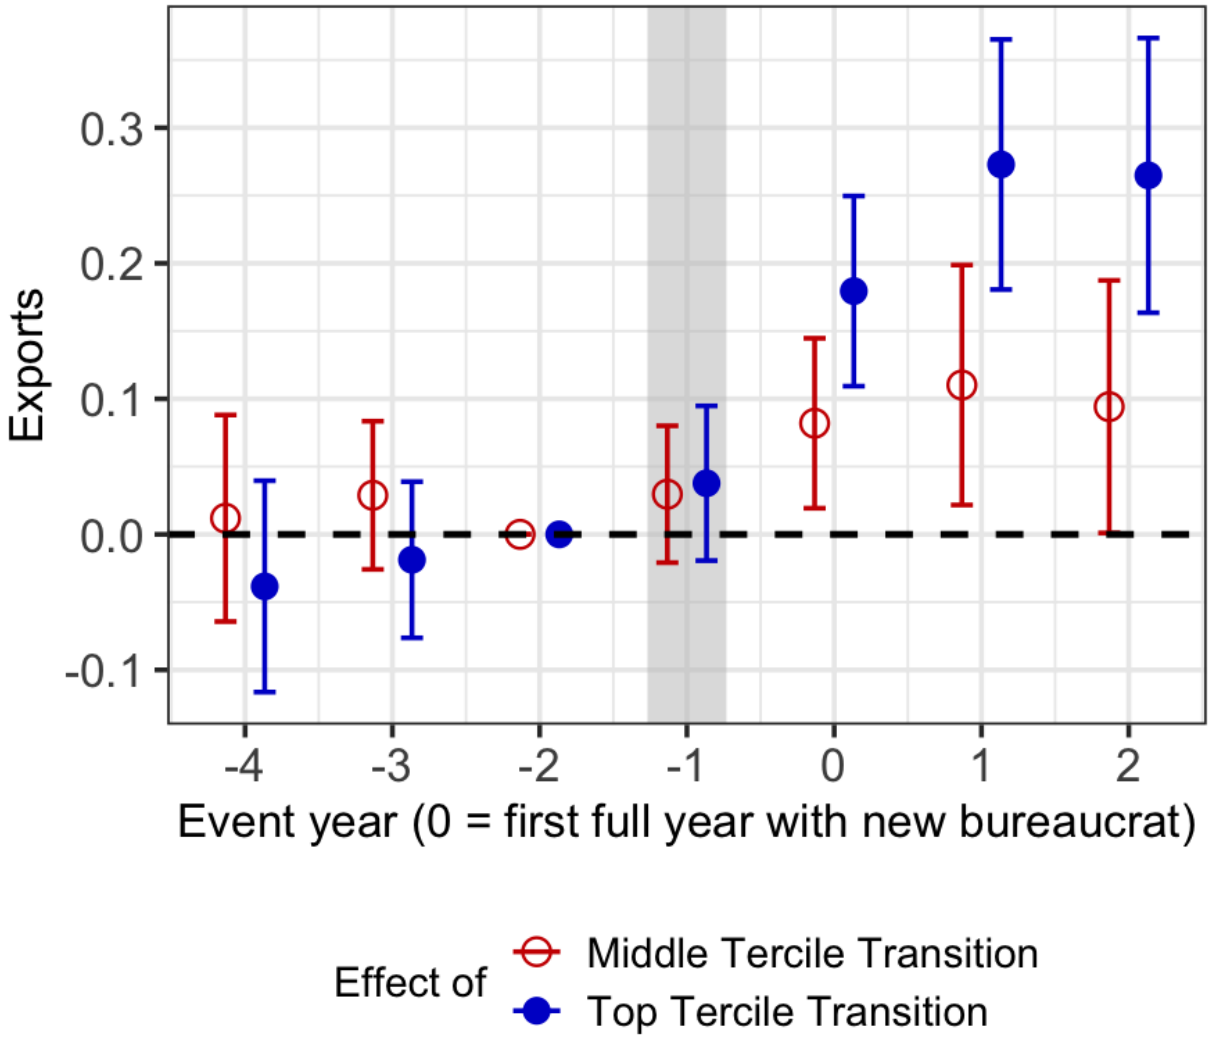

Notes: The figure shows the effect of the change in bureaucrat fixed effects on exports, around the time that the bureaucrat managing a country office changes. Specifically, we plot $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (4). The outcome variable is the inverse hyperbolic sine of export value to the country of the bureaucrat switch. Transitions are categorized into terciles depending on the size of the change in bureaucrat fixed effects induced by the switch, where the fixed effects are estimated separately from equation (2). The omitted category is a transition in the bottom tercile. 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 5: Symmetric Effects from Gaining and Losing a Bureaucrat, No Differential Pre-Trends

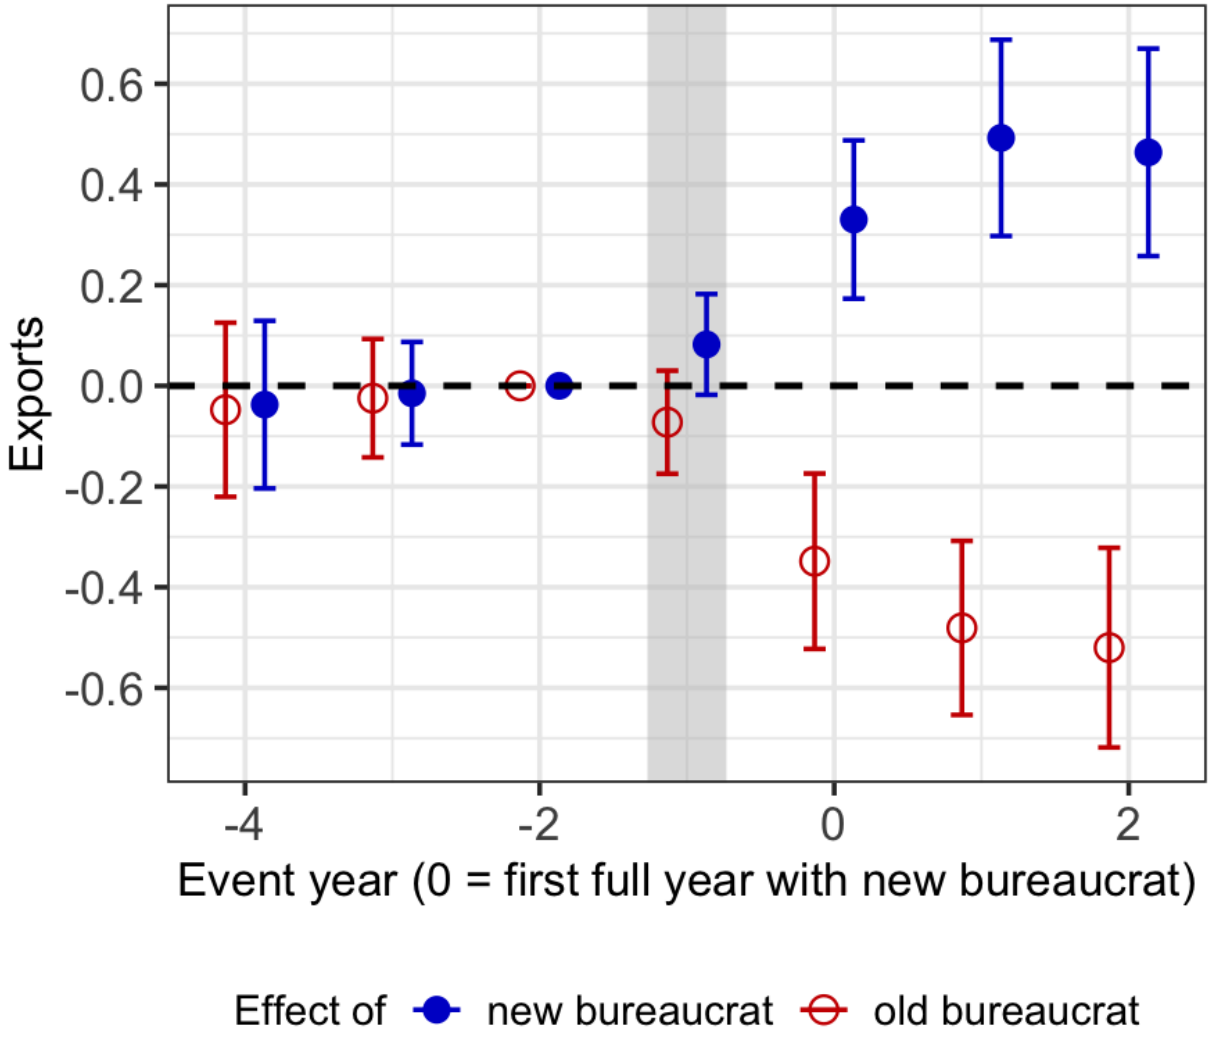

Notes: The figure shows the effect of the incoming and outgoing bureaucrats' fixed effects on exports, around the time that the bureaucrat managing a country office changes. Specifically, we plot $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (5). The outcome variable is the inverse hyperbolic sine of export value to the country of the switch between bureaucrats. 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 (5) 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 6: Bureaucrat Fixed Effects Predict Exports Out of Sample

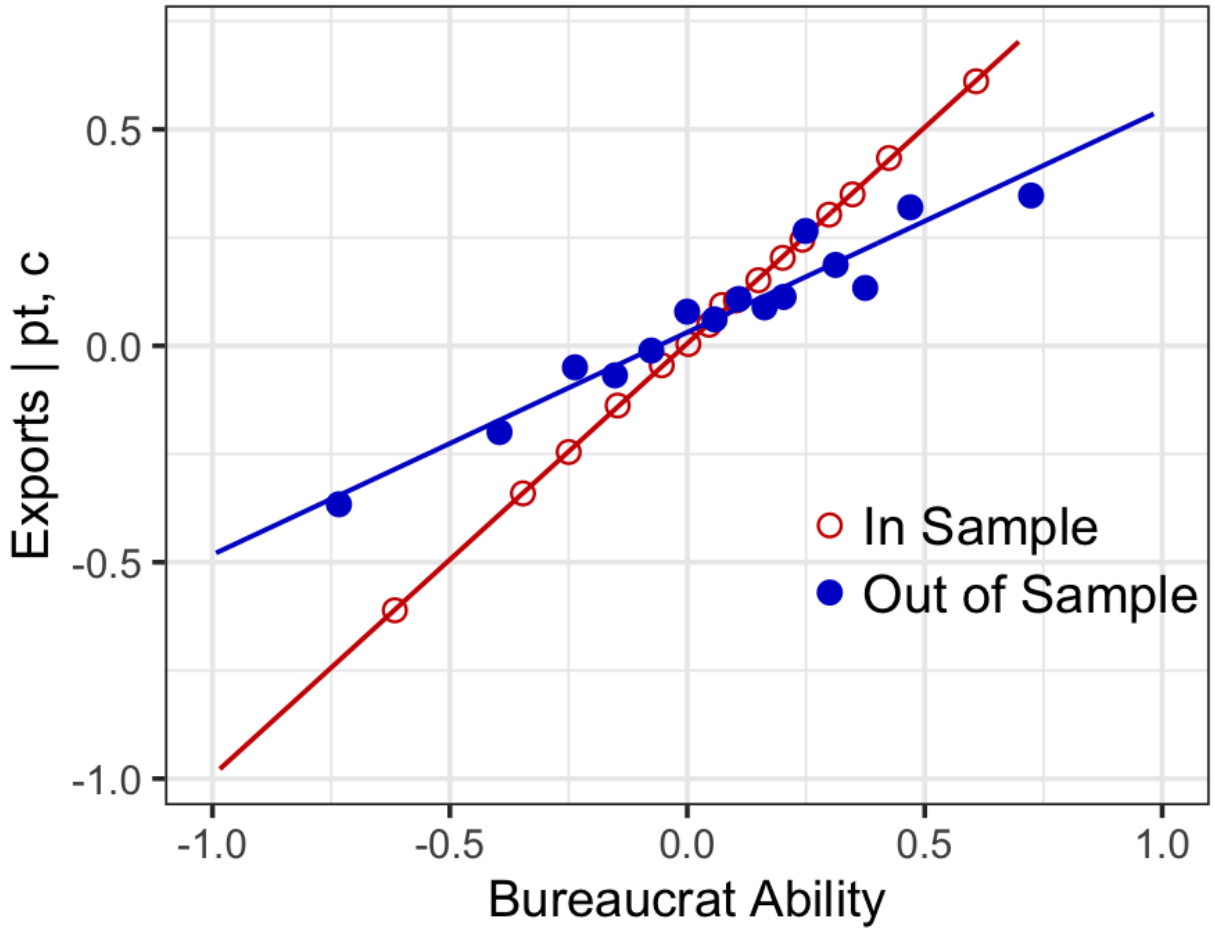

Notes: The figure displays a binned scatterplot of residualized exports on in-sample (i.e., baseline) and out-of-sample bureaucrat fixed effects. On the vertical axis, we plot exports, i.e., the inverse hyperbolic sine of export value, residualized for product-year fixed effects ($\hat{\lambda}_{pt}$) and country fixed effects ($\hat{\gamma}_c$). These fixed effects, as well as the in-sample bureaucrat fixed effects, are estimated from equation (2) using data from all country-years. Hence, by construction, each in-sample dot lies on the 45-degree line, such that 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. Back to Section 4.4.

Figure 7: Good Bureaucrats Increase Exports for Products with Demand (Supply) Growth

(a) Bureaucrat Ability and Effect of Product Demand (b) Bureaucrat Ability and Effect of Product Supply

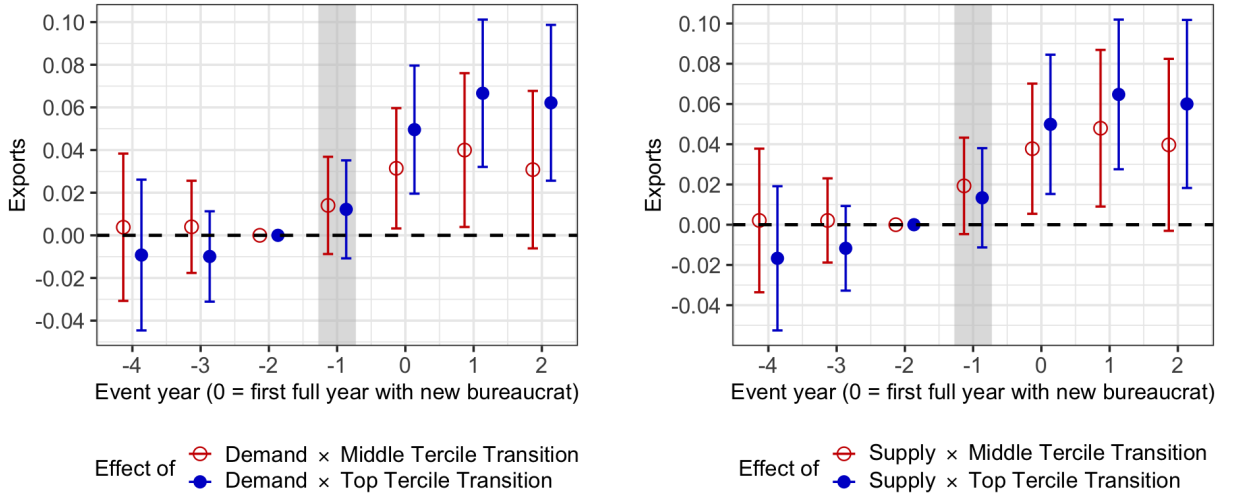

Notes: The figure plots the effect on exports of the change in bureaucrat ability when interacted with two kinds of shocks, estimated from equation (6). The outcome variable is the inverse hyperbolic sine of export value to the country of the switch between bureaucrats. The left panel depicts $\hat{\beta}_k^D$, $\hat{\delta}_k^D$, the interaction effects with the destination country's product-specific "demand," proxied by exports of the same product to the same destination by countries other than Korea. The right panel depicts $\hat{\beta}_k^S$, $\hat{\delta}_k^S$, the interaction effects with South Korea's product-specific "supply," proxied by Korean exports of the same product to other destinations. The horizontal axis indicates k, the years relative to a bureaucrat switch. 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. As the regressions include product-event fixed effects, both of these interactions capture the differential effect on Korean exports of *changes* in market conditions for product p in country c relative to year -2. Since equation (6) 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 4.5.

Figure 8: Single-Appointment Bureaucrats Are Less Effective Than Multi-Appointment Bureaucrats

(a) Residualized exports by appointment

(b) Bureaucrat fixed effects

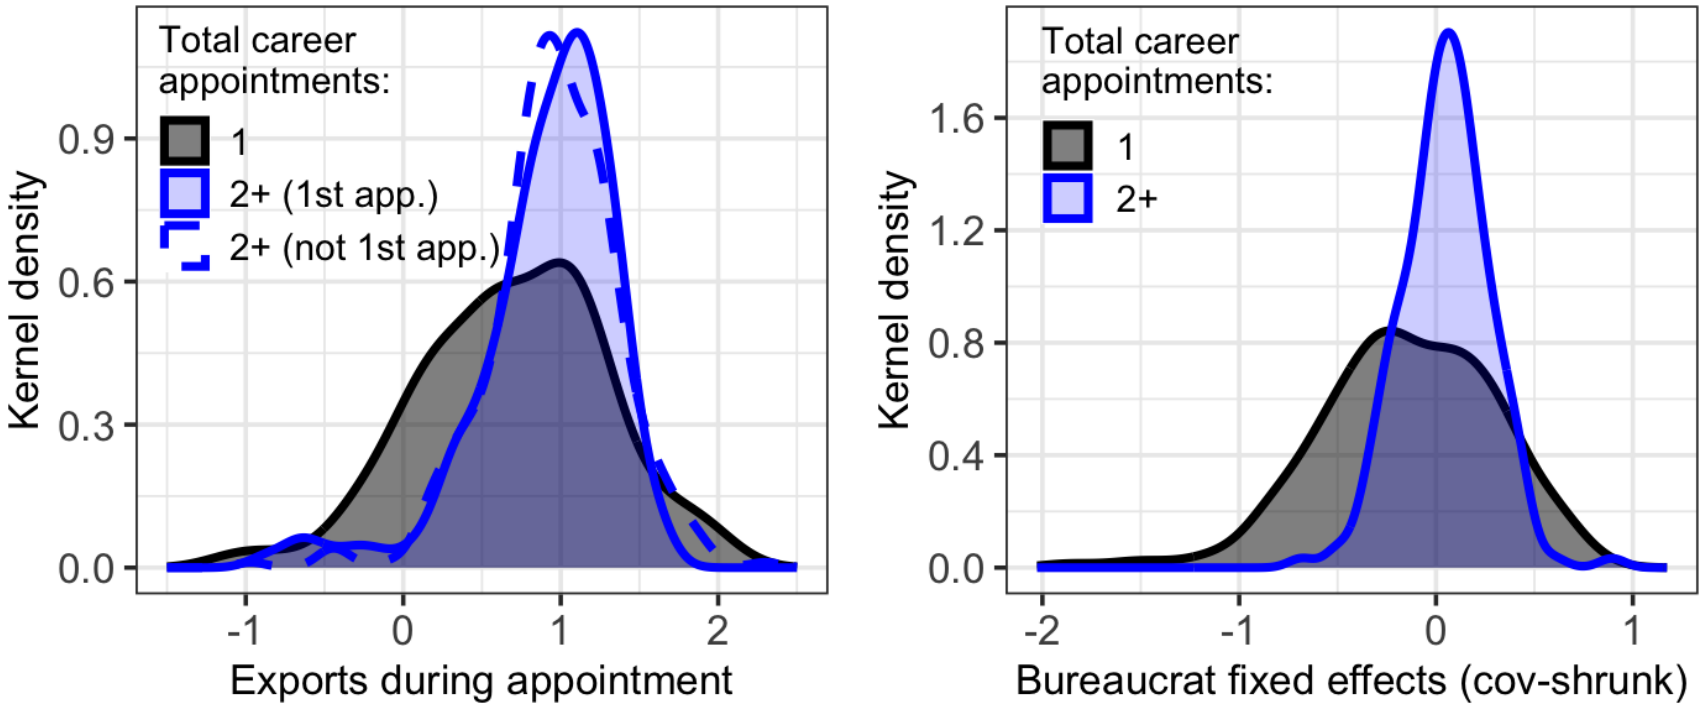

Notes: The figure shows the kernel density of (a) residualized exports during a bureaucrat's appointment as country office manager and (b) bureaucrat fixed effects (covariance-shrunk). It does so separately for bureaucrats with one appointment (black) and for those with two or more appointments (blue) over the course of their careers. In both panels, the distributions for single-appointment bureaucrats have much fatter left tails than for multi-appointment bureaucrats. 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. Back to Sections 2.2 and 4.6.

Figure 9: As Bureaucrats' Careers Progress, They Are Appointed to More Important Countries

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

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

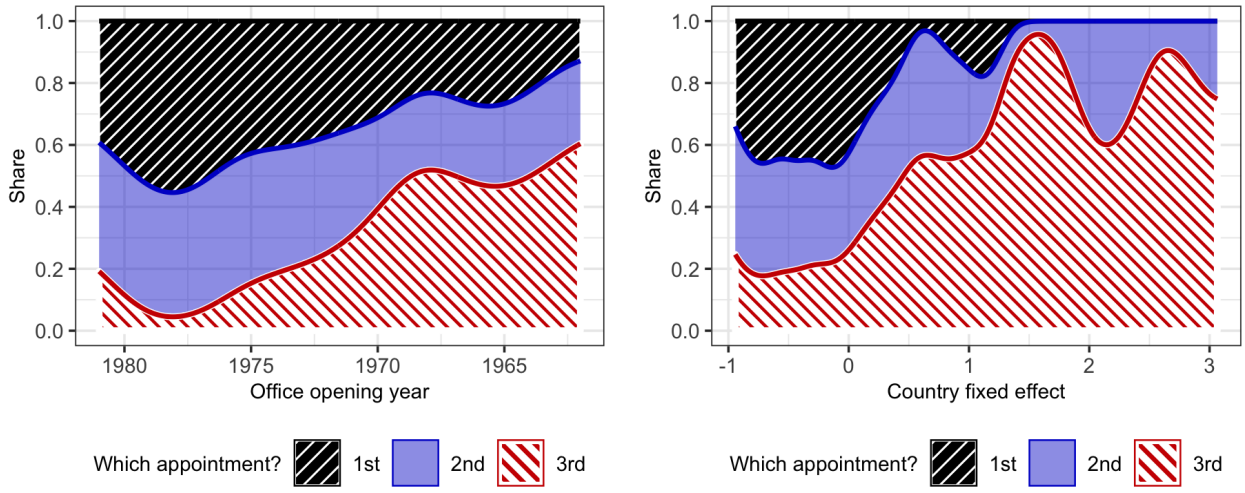

Notes: The figure shows the (smoothed) shares of appointments by (1) country importance and (2) whether the appointment is the 1st, 2nd, or 3rd appointment of a bureaucrat. Panel (a) depicts the shares in terms of the office opening year of a country. Panel (b) depicts the shares in terms of the value of the country fixed effect. For both plots, countries increase in importance as South Korea's trade parter from left to right, reflected in earlier openings and larger fixed effects. We consider offices that were opened during the main rollout period (1962-1981). 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. More important countries tend to be filled by bureaucrats in their 2nd, and especially 3rd, appointments. The pattern remains, slightly less strong, when all bureaucrats are included (not shown). Back to Section 4.6.

Figure 10: Opening an Export Promotion Office Increases Korean Exports to That Country

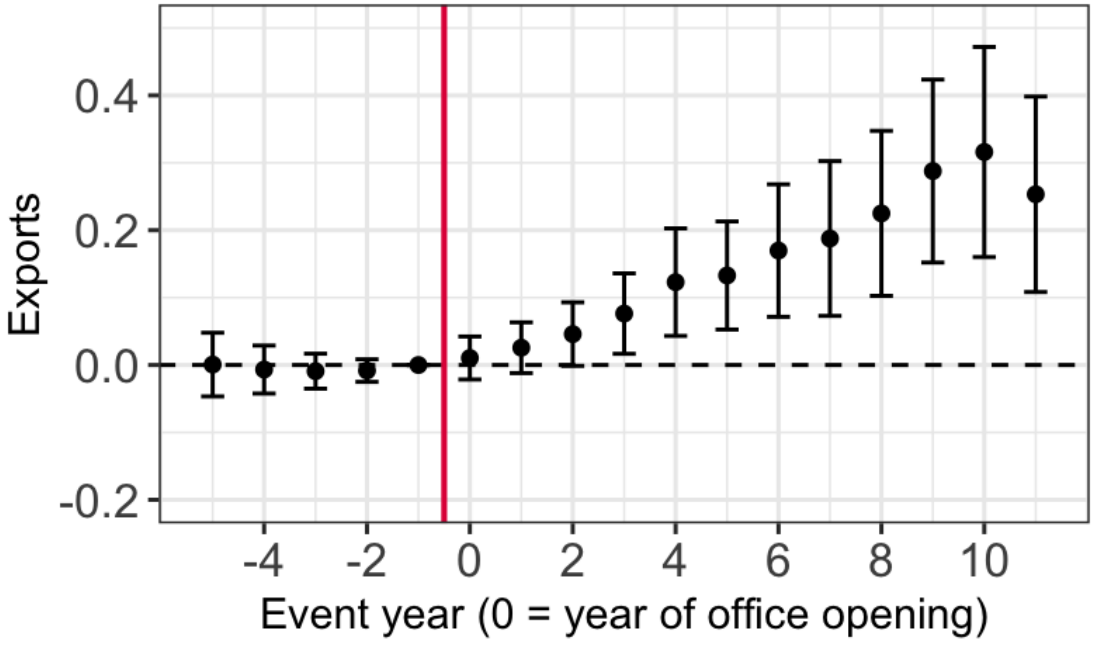

(a) Sample includes openings from 1967 to 1981

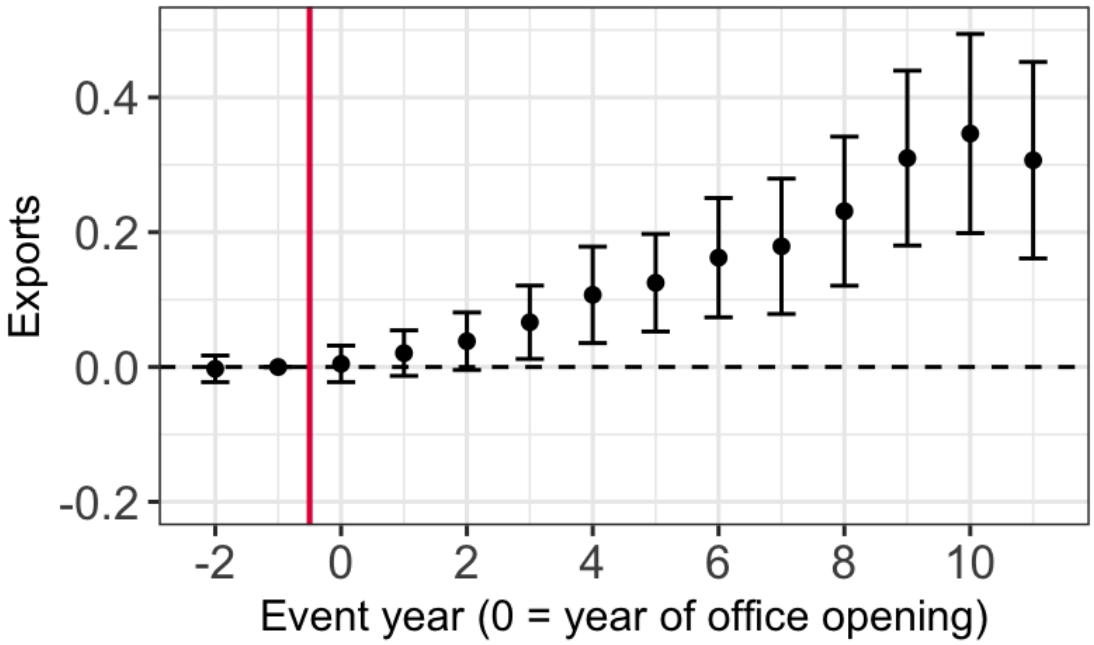

(b) Sample includes openings from 1964 to 1981

Notes: This figure plots the event-study coefficients ($\hat{\pi}_k$) of the effect of opening export promotion offices on Korean exports, transformed using the inverse hyperbolic sine, estimated from equation (7) with a never-treated control group. Standard errors are clustered at the country level. Back to Section 5.2.

Figure 11: Office Opening Does Not Coincide with
Increases in Import Demand of That Country

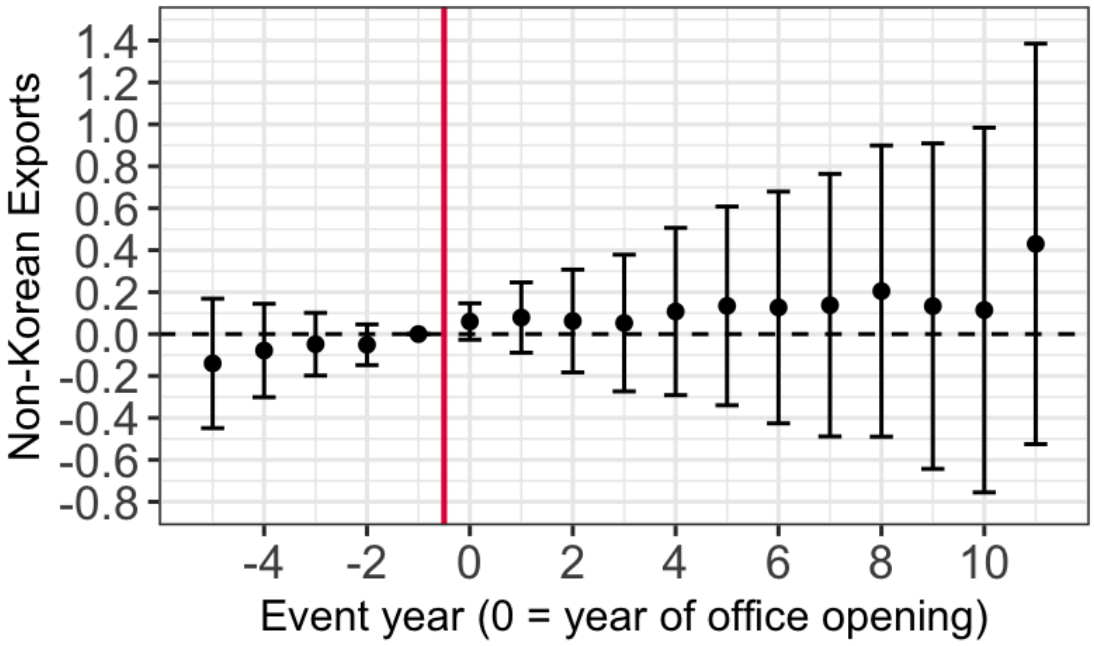

(a) Non-Korean exports as outcome. Sample includes openings from 1967 to 1981.

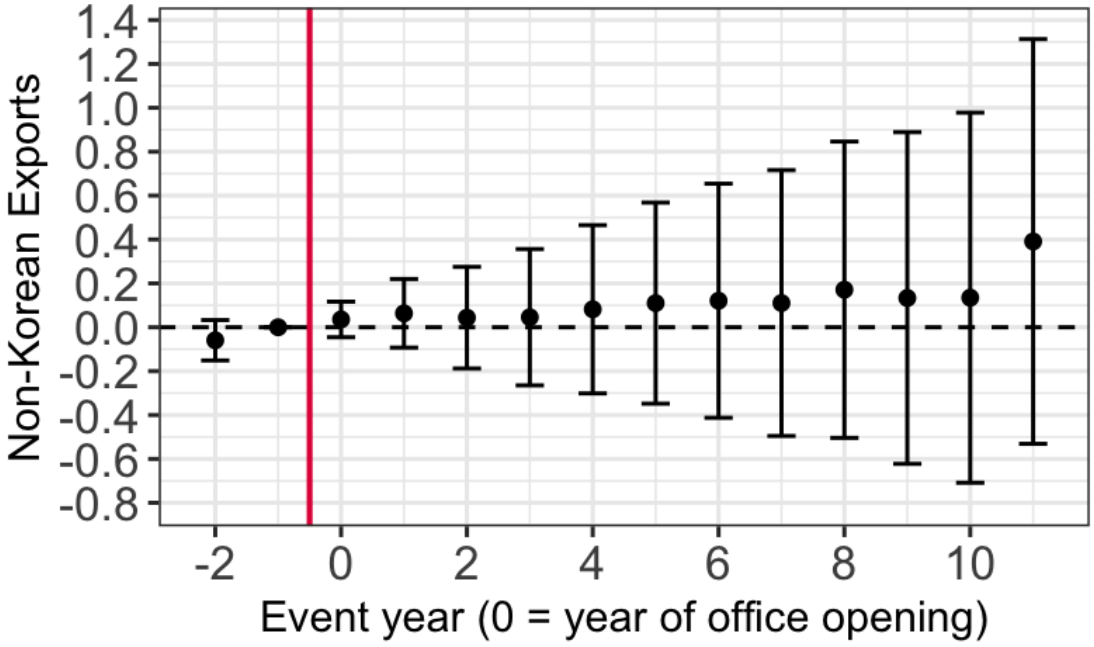

(b) Non-Korean exports as outcome. Sample includes openings from 1964 to 1981.

Notes: This figure plots the event-study coefficients ($\hat{\pi}_k$) of the effect of opening export promotion offices on *non-Korean* exports, transformed using the inverse hyperbolic sine, estimated from equation (7) with a never-treated control group. Standard errors are clustered at the country level. Back to Section 5.3.

Figure 12: Order of Office Openings in Europe Follows Pre-Determined Market Size

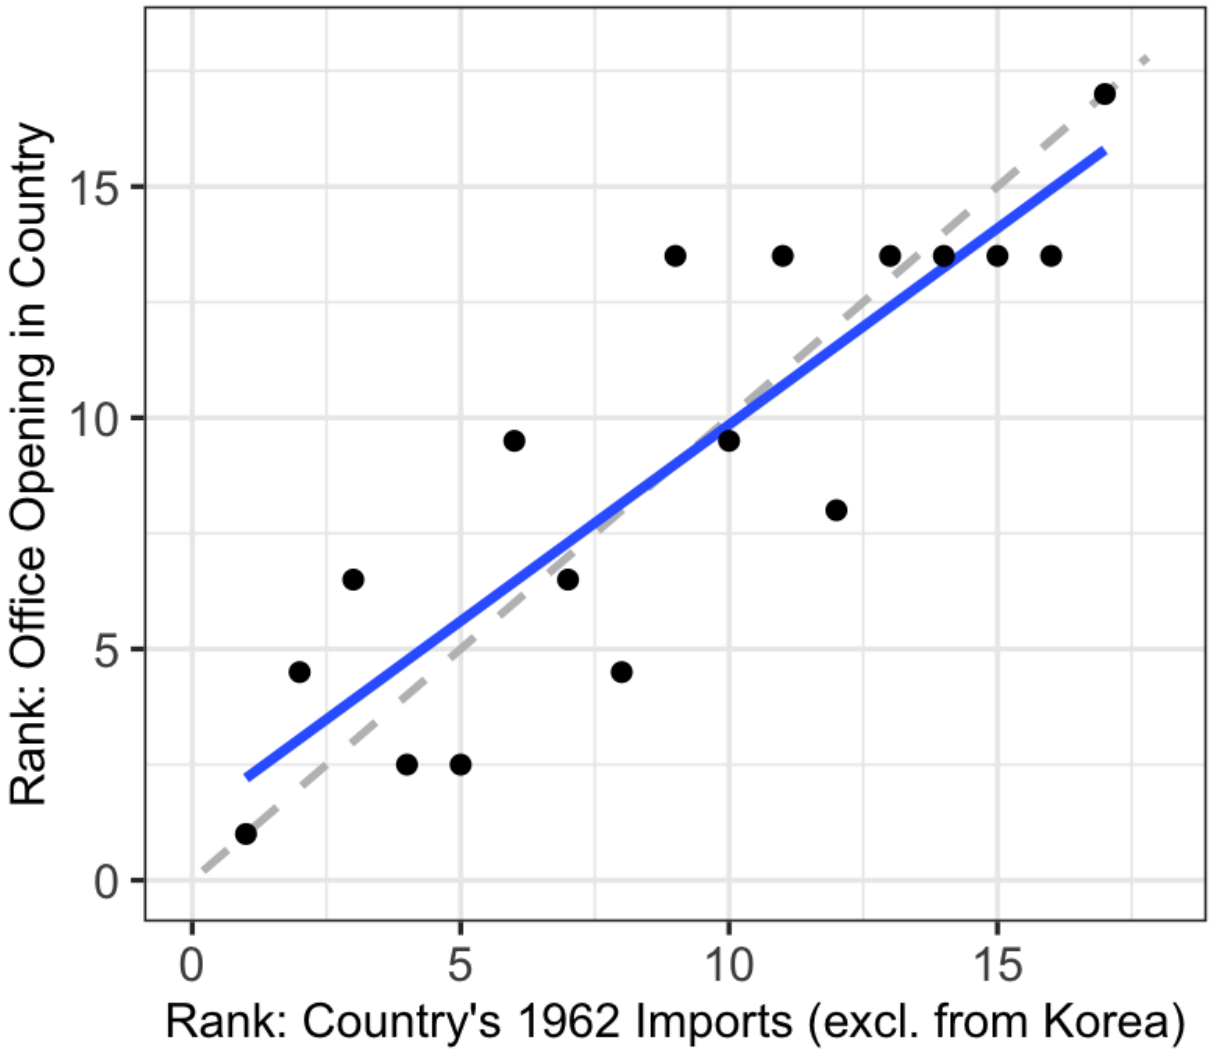

Notes: Each dot corresponds to a European country that received a KOTRA office during the main roll-out of offices (1962-1981). The horizontal axis gives each country's rank in terms of 1962 imports, excluding imports from Korea. The vertical axis gives each country's rank in terms of the order of their office openings. The solid blue line gives the linear fit using 1962 market size to predict the order of office openings. The rank correlation between 1962 imports and office opening is 0.87. The dashed gray line gives the 45-degree line, where the two ranks are exactly equal. This is the case for the UK (rank 1) and Portugal (rank 17). When multiple countries have the same opening year, we assign the average rank to them. For example, Italy and the Netherlands get the second and third offices. As these openings occur in the same year, both have rank 2.5. Back to Section 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 (1)	% Share (2)	Component (3)	% Share (4)	Component (5)	% Share (6)	Component (7)	% Share (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, Saggio, and Sølvsten (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, for each bureaucrat the number of appointment spells. Columns 7-8 include fixed effect for Korean ambassadors instead of KOTRA office managers. The Kline, Saggio, and Sølvsten (2020) bias correction method is applied to columns 5-8 as well. Since the algorithm is based on numerical approximations of the traces of large matrix inverses, there is a small degree of randomness in the decomposition results. There is also additional randomness in columns 5-6 arising from the random reshuffling of bureaucrats. Thus, the variances and covariances we report in all columns are averages of 100 iterations of the algorithm (holding the original sample fixed). The number of observations, bureaucrats, and countries are also averages of 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. Again, for each bootstrap sample, we take the average of 100 iterations. Back to Sections 4.3 and 4.4.

Table 3: 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 the change in exports (inverse hyperbolic sine of export value) and the percentage change in export value (approximated using $(e^{\Delta \text{exports}} - 1) \times 100$) for 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.

Bibliography

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

Adhvaryu, A., N. Kala, and A. Nyshadham (2022). Management and Shocks to Worker Productivity. *Journal of Political Economy 130*(1), 1–47.

Adhvaryu, A., A. Nyshadham, and J. Tamayo (2022). Managerial Quality and Productivity Dynamics. *The Review of Economic Studies 90*(4), 1569–1607.

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.

Akhtari, M., D. Moreira, and L. Trucco (2022). Political Turnover, Bureaucratic Turnover, and the Quality of Public Services. *American Economic Review 112*(2), 442–93.

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.

Aman-Rana, S. (2025). Meritocracy in a Bureaucracy. *Journal of Development Economics 175*.

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.

Angrist, J. D. and J.-S. Pischke (2010). The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics. *Journal of Economic Perspectives 24*(2), 3–30.

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

Ashraf, N. and O. Bandiera (2018). Social Incentives in Organizations. *Annual Review of Economics 10*(1), 439–463.

Atkin, D., A. Costinot, and M. Fukui (2025). Globalization and the Ladder of Development: Pushed to the Top or Held at the Bottom? *Review of Economic Studies*.

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.

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

Bandiera, O., A. Prat, S. Hansen, and R. Sadun (2020). CEO Behavior and Firm Performance. *Journal of Political Economy 128*(4), 1325 – 1369.

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

Barwick, P. J., K. Myrto, and Z. N. Bin (2025). Industrial Policy Implementation: Empirical Evidence from China's Shipbuilding Industry. *Review of Economic Studies*.

Bergeron, A., G. Tourek, and J. L. Weigel (2024). The State Capacity Ceiling on Tax Rates: Evidence From Randomized Tax Abatements in the DRC. *Econometrica 92*(4), 1163–1193.

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.

Borusyak, K., X. Jaravel, and J. Spiess (2024). Revisiting Event Study Designs: Robust and Efficient Estimation. *Review of Economic Studies 91*(6), 3253–3285.

Brown, C. and T. Andrabi (2021). Inducing Positive Sorting through Performance Pay: Experimental Evidence from Pakistani Schools. Working paper.

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 (2024). Industrialization and the Big Push: Theory and Evidence from South Korea. Working Paper 2024/259, IMF.

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.

de Chaisemartin, C. and X. D'Haultfœuille (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. *American Economic Review 110*(9), 2964–96.

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.

Frederiksen, A., L. B. Kahn, and F. Lange (2020). Supervisors and Performance Management Systems. *Journal of Political Economy 128*(6), 2123–2187.

Friebel, G., M. Heinz, and N. Zubanov (2022). Middle Managers, Personnel Turnover, and Performance: A Long-Term Field Experiment in a Retail Chain. *Management Science 68*(1), 211–229.

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.

Goodman-Bacon, A. (2021). Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics 225*(2).

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.

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

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.

Hoffman, M. and S. Tadelis (2021). People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis. *Journal of Political Economy 129*(1), 243 – 285.

Janke, K., C. Propper, and R. Sadun (2019). "The Role of Top Managers in the Public Sector: Evidence from the English NHS". Working Paper 25853, National Bureau of Economic Research.

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.

Khan, A. Q., A. I. Khwaja, and B. A. Olken (2019). Making Moves Matter: Experimental Evidence on Incentivizing Bureaucrats through Performance-Based Postings. *American Economic Review 109*(1), 237–370.

Khan, M. Y. (2025). Mission Motivation and Public Sector Performance: Experimental Evidence from Pakistan. Working paper.

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*.

Lazear, E., K. L. Shaw, and C. T. Stanton (2015). The Value of Bosses. *Journal of Labor Economics 33*(4), 823 – 861.

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.

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*.

Mehmood, S. (2022). The Impact of Presidential Appointment of Judges: Montesquieu or the Federalists? *American Economic Journal: Applied Economics 14*(4), 411–45.

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.

Mollick, E. (2012). People and process, suits and innovators: the role of individuals in firm performance. *Strategic Management Journal 33*(9), 1001–1015.

Munch, J. and G. Schaur (2018). The Effect of Export Promotion on Firm-Level Performance. *American Economic Journal: Economic Policy 10*(1), 357–87.

Muñoz, P. and M. Prem (2024). Managers' Productivity and Recruitment in the Public Sector. *American Economic Journal: Economic Policy 16*(4), 223–53.

Otero, C. and P. Muñoz (2025). Managers and Public Hospital Performance. Working Paper.

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

Patault, B. and C. Lenoir (2024). Customer Capital Spillovers: Evidence from Sales Managers in International Markets. *American Economic Journal: Applied Economics 16*(4), 404–443.

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

Rasul, I. and D. Rogger (2017). Management of bureaucrats and public service delivery: Evidence from the nigerian civil service. *The Economic Journal 128*(608), 413–446.

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.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199. Themed Issue: Treatment Effect 1.

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.

Online Appendix

A Appendix Figures

Figure A.1: Policy Roll-out: Number of Countries with Export Promotion Office

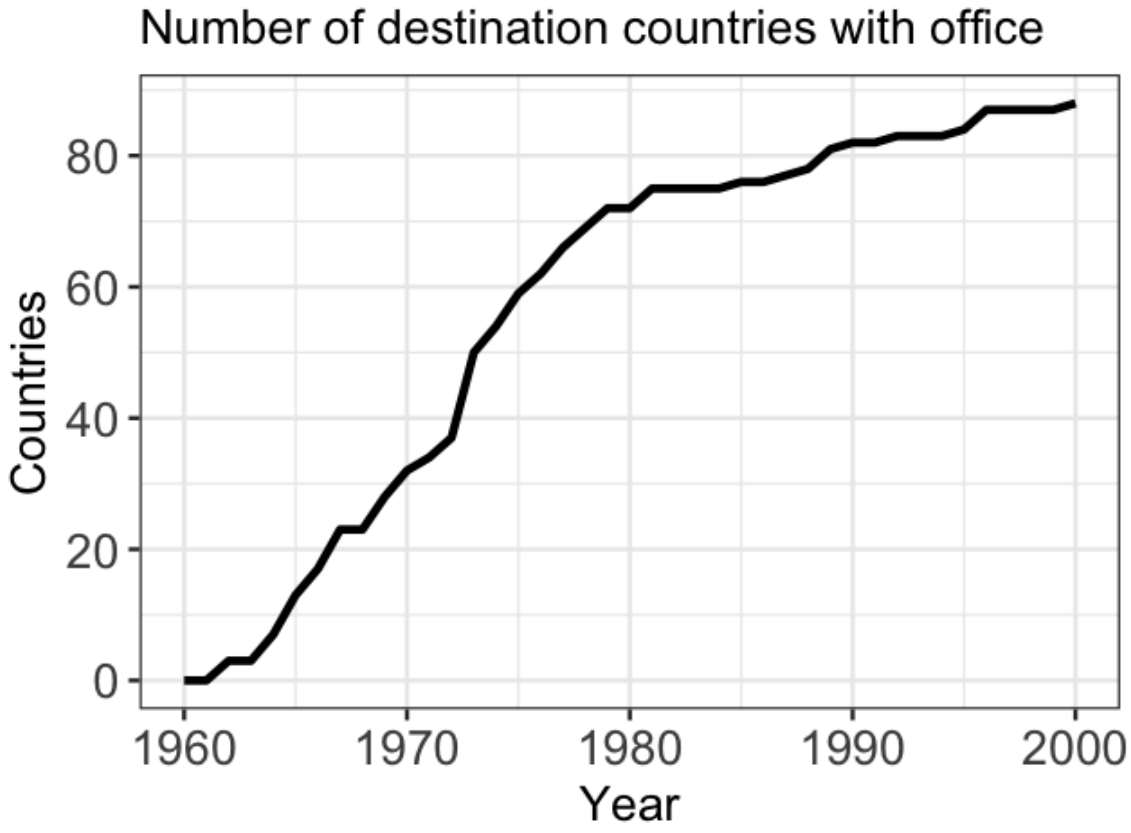

Notes: This figure presents the cumulative number of countries with KOTRA offices by year. Back to Sections 2.2 and 5.

Figure A.2: Distribution of Appointment Durations.
Median and Modal Duration: 36 Months.

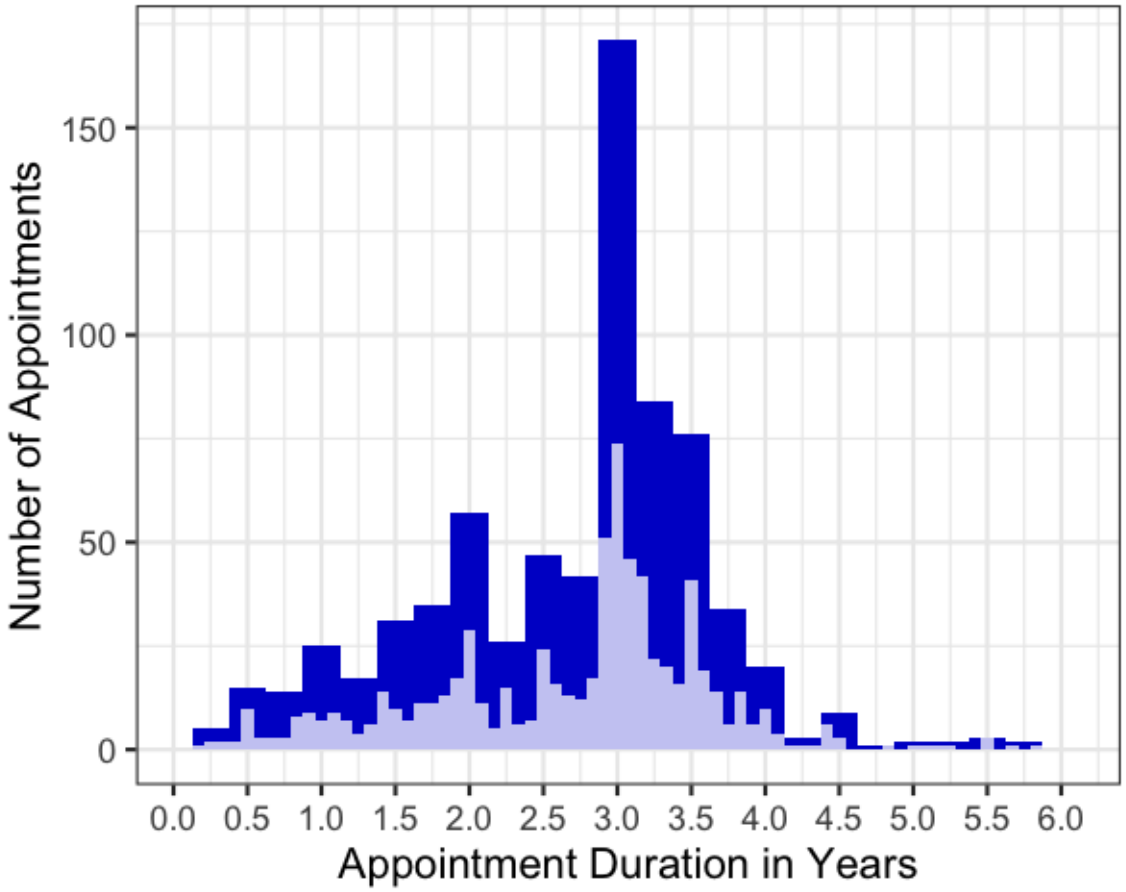

Notes: This figure represents the distribution of appointment durations. The blue (light blue) bars indicate the number of appointments by quarterly (monthly) duration. As each quarter contains multiple months, the blue bars always weakly exceed the light blue ones. Back to Sections 2.3 and 4.1.

Figure A.3: Distribution of Gap Lengths.
Median: 29 Months. Mode: 30 Months.

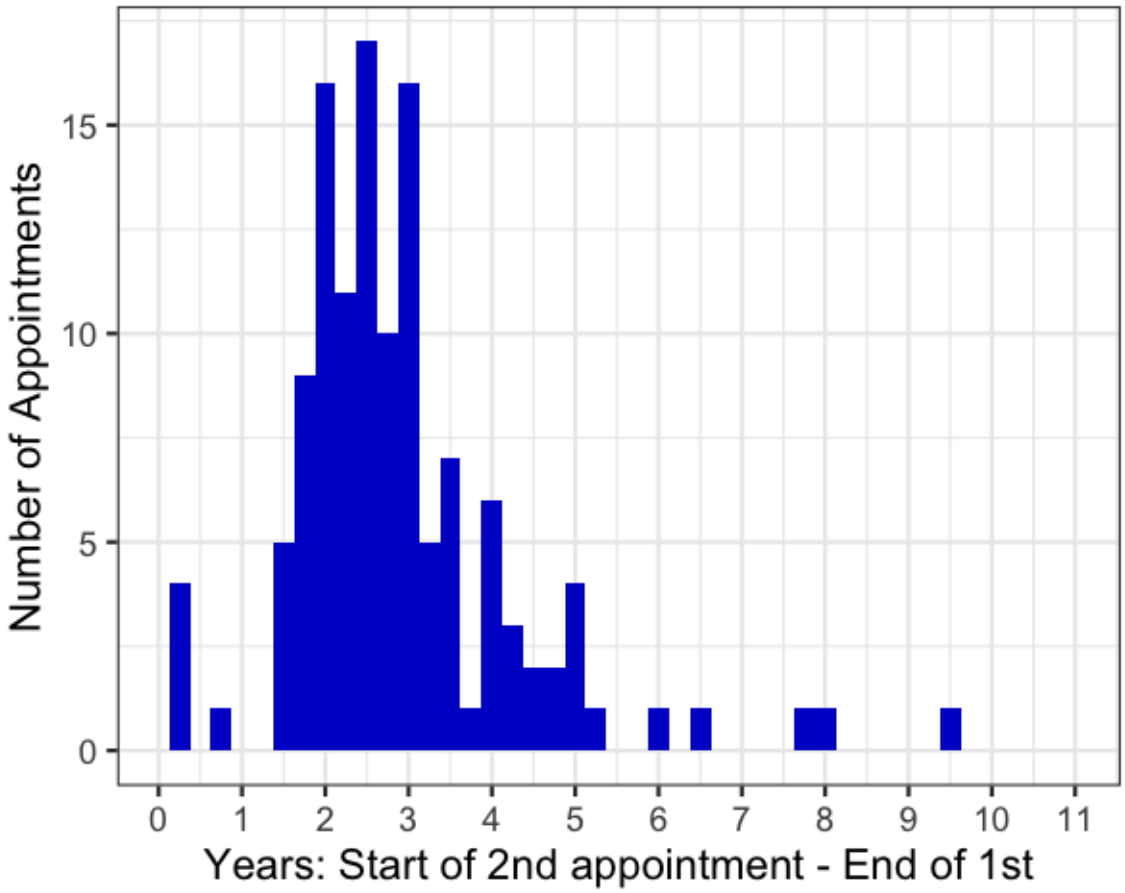

Notes: This figure represents the distribution of the duration of gaps between appointments. The blue bars indicate the number of gaps by quarterly duration. Back to Section 2.3.

Figure A.4: Targeting of Export Promotion Activity by Product.
Export Promotion Activity Moves in Parallel with Heavy-Chemical Industry Drive

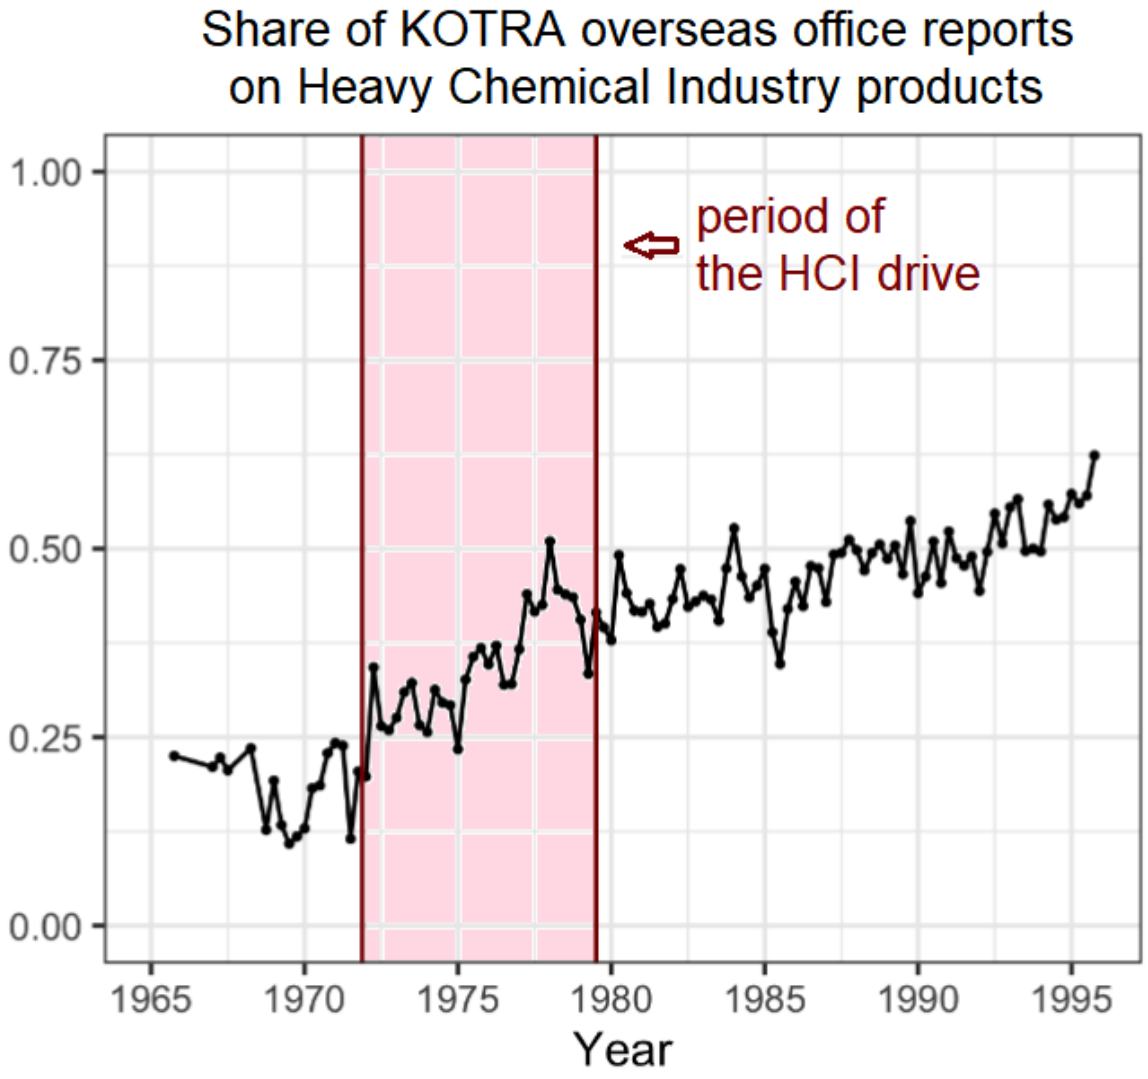

Notes: Targeting of export promotion activity by product. For each quarter, the vertical axis presents the share of overseas office reports that could be linked to the target products of the Heavy-Chemical Industry (HCI) Drive, relative to the number of reports that could be linked to any product. Back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

Figure A.5: Demonstration of Limited Mobility Bias: Variances and Covariances of Sampling Errors Converge to Zero As Mobility Increases

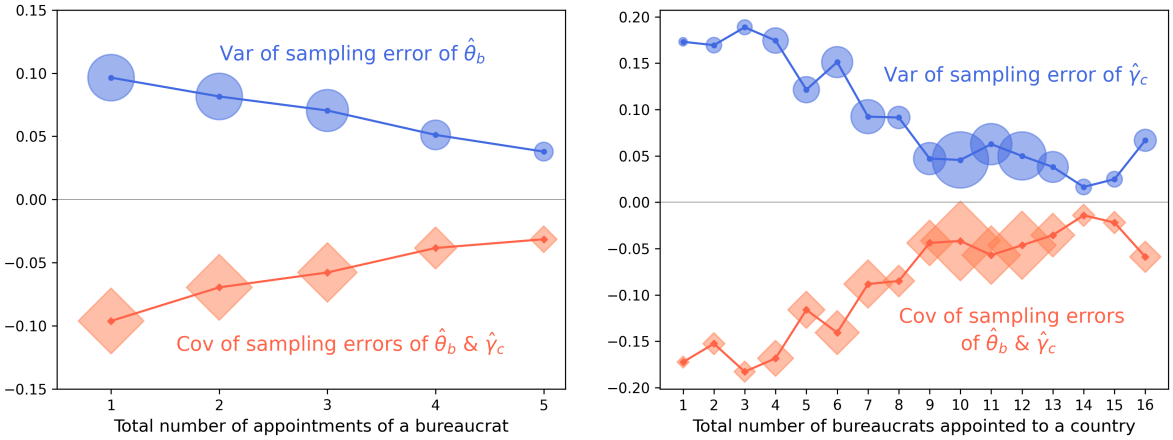

(a) Sampling errors of bureaucrat fixed effects

(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 E 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.6: Distribution of Bureacrat Fixed Effects, Raw vs. Covariance-Shrunk

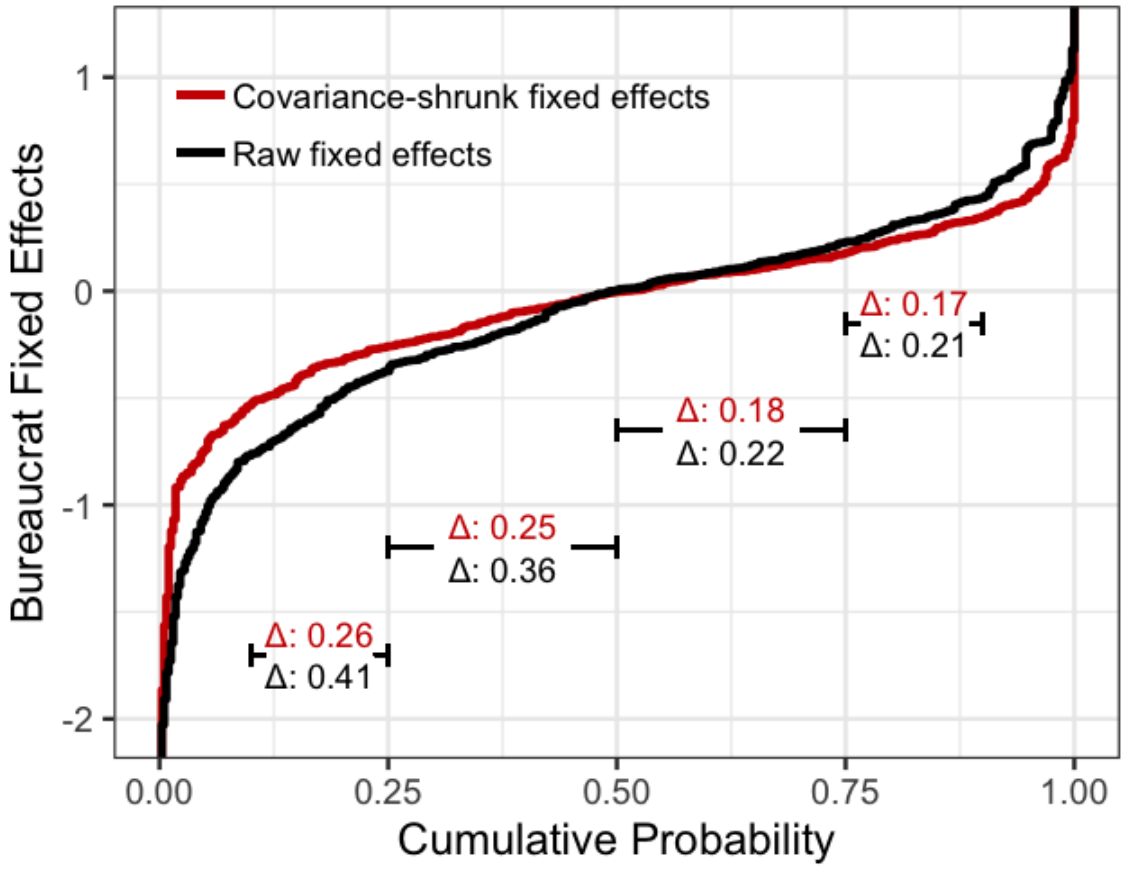

Notes: The figure shows the cumulative distribution function of the raw bureaucrat fixed effects estimated from equation (3), as well as covariance-shrunk bureaucrat fixed effects. Moving from the 25th to the 50th percentile would imply an increase in the inverse hyperbolic sine of export value to a country 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.7: Robustness: Event Study With Out-of-Sample Bureaucrat Fixed Effects

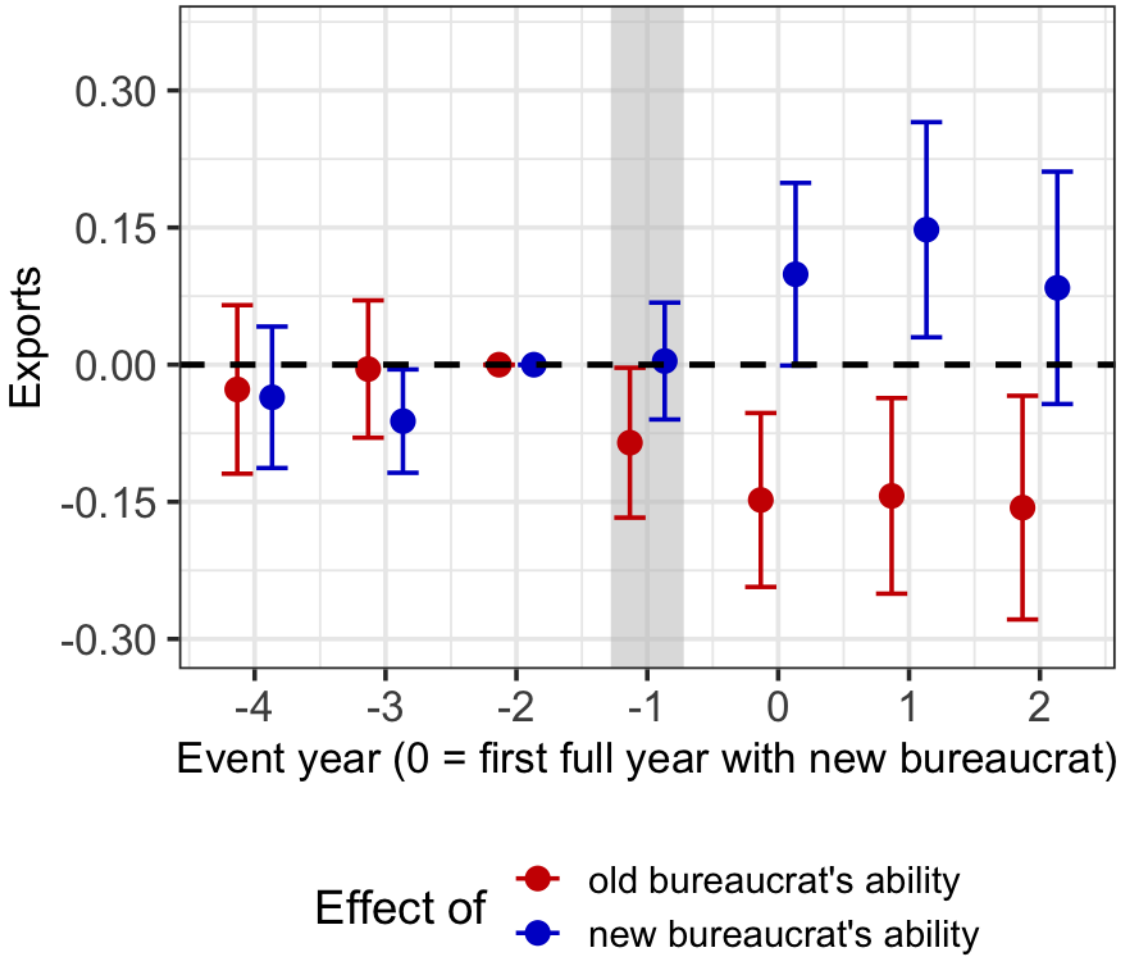

Notes: The figure shows the effect on exports of the incoming and outgoing bureaucrats' fixed effects, estimated *out of sample*, around the time that the bureaucrat managing a country office changes. Specifically, we plot $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (5). The outcome variable is the inverse hyperbolic sine of export value to the country of the switch between bureaucrats. Out-of-sample estimation means that, for example, for event studies of bureaucrat switches in the UK office, we use fixed effects estimated using data excluding the UK. As out-of-sample fixed effects are not available for single-appointment bureaucrats, we report coefficients from two different models to maximize power. First, we estimate equation (5) using *out-of-sample* estimates for outgoing bureaucrats and *in-sample* estimates for incoming bureaucrats. Second, we estimate equation (5) using *in-sample* estimates for outgoing bureaucrats and *out-of-sample* estimates for incoming bureaucrats. For each model, we only report the out-of-sample coefficients because they are the ones of interest. Though unreported, the in-sample coefficients in each model are almost symmetric to the out-of-sample ones. 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. Standard errors are clustered by country. Back to Section 4.4.

Figure A.8: Robustness of Office Opening Results to Controls

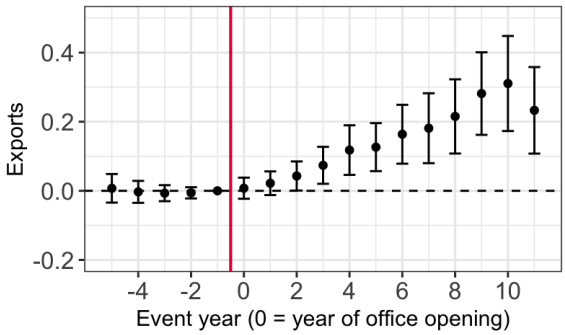

(a) Control for non-Korean exports. Sample includes openings from 1967 to 1981

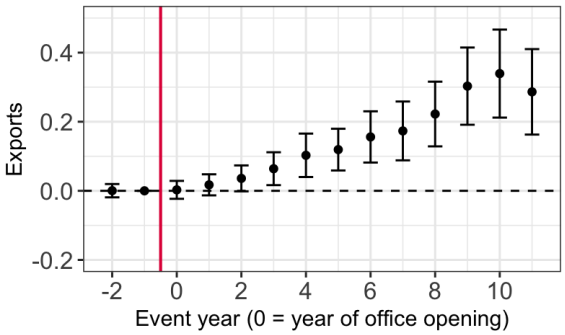

(b) Control for non-Korean exports. Sample includes openings from 1964 to 1981

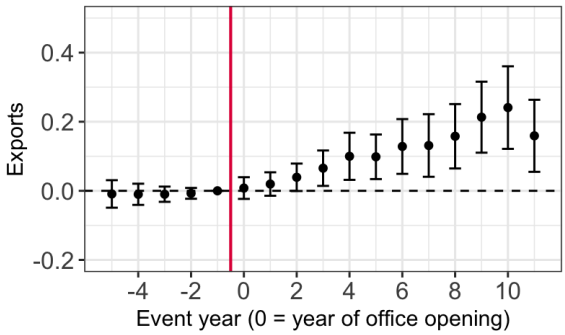

(c) Control for non-Korean exports × year dummies. Openings from 1967 to 1981

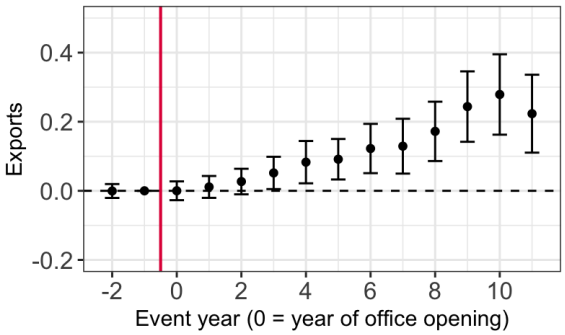

(d) Control for non-Korean exports × year dummies. Openings from 1964 to 1981

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 export value to the country of the office opening. Panels (a)-(d) perform robustness checks on the results in Figure 10 by controlling for non-Korean exports for the two samples (allowing for different numbers of pre-periods). Standard errors are clustered by country. Back to Section 5.3.

Figure A.9: Opening an Office Increases KOTRA Activity on That Country

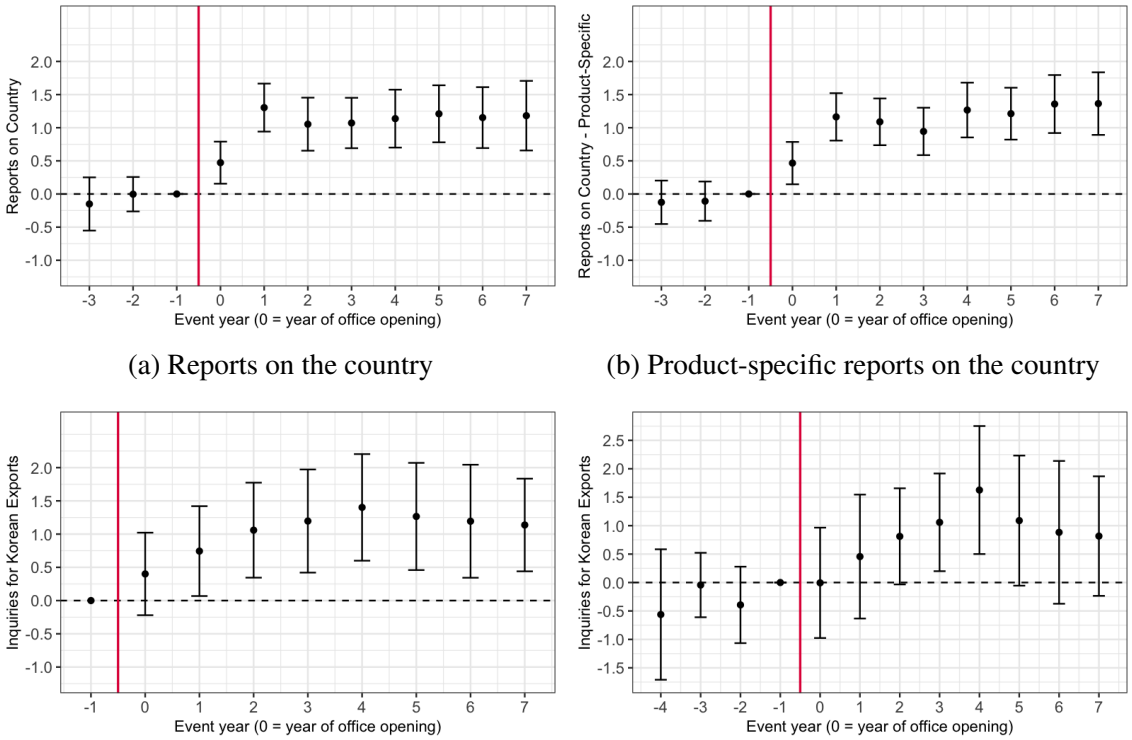

(a) Reports on the country

(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 – which may be more specific or informative, and (3) the number of inquiries from potential importers in that country interested in trade with Korea, each transformed by the inverse hyperbolic sine. The data on reports covers the years 1965 to 2001, so we exclude offices that opened before 1968 from the analysis in Panels (a) and (b). The data on inquiries covers the years 1974 to 1997, so we exclude offices that opened before 1974 from the analysis in Panel (c). Including openings from 1975 comes at the cost of being able to have only 1 pre-period in Panel (c). Panel (d) takes an alternative approach of including 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: The Effect of Export Promotion Depends on the Individual Bureaucrat. Bureaucrat Effects Do Not Differ Between Appointments

	Exports			
	(1)	(2)	(3)	(4)
Share of Variation explained by FE				
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 of Korean export value 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 C.1.

C Further Diagnostics of Variation Explained by Bureaucrats

C.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 C.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.[79] Appendix Table B.1 reports the changes in the R^2 when we regress exports on different combinations of fixed effects.[80] 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.[81] To understand 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 C.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, Heining, and Kline (2013), Card, Cardoso, and Kline (2016) and Best, Hjort, and Szakonyi (2023). Appendix Figure C.2 shows that the main takeaways from Figures 4 and 5 are present for transitions across all terciles of old and new bureau-

[79] 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).

[80] While informative, the R^2 analysis is subject to limited mobility bias. In the formal variance decomposition exercise of Table 2 in the main text, the issue is addressed by adopting the Kline, Saggio, and Sølvsten (2020) bias correction method.

[81] 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 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).

crats. 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.[82]

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

[82]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 C.1: Residuals By Estimated Bureaucrat and Country Effects.
No Clear Pattern Suggesting Misspecification.

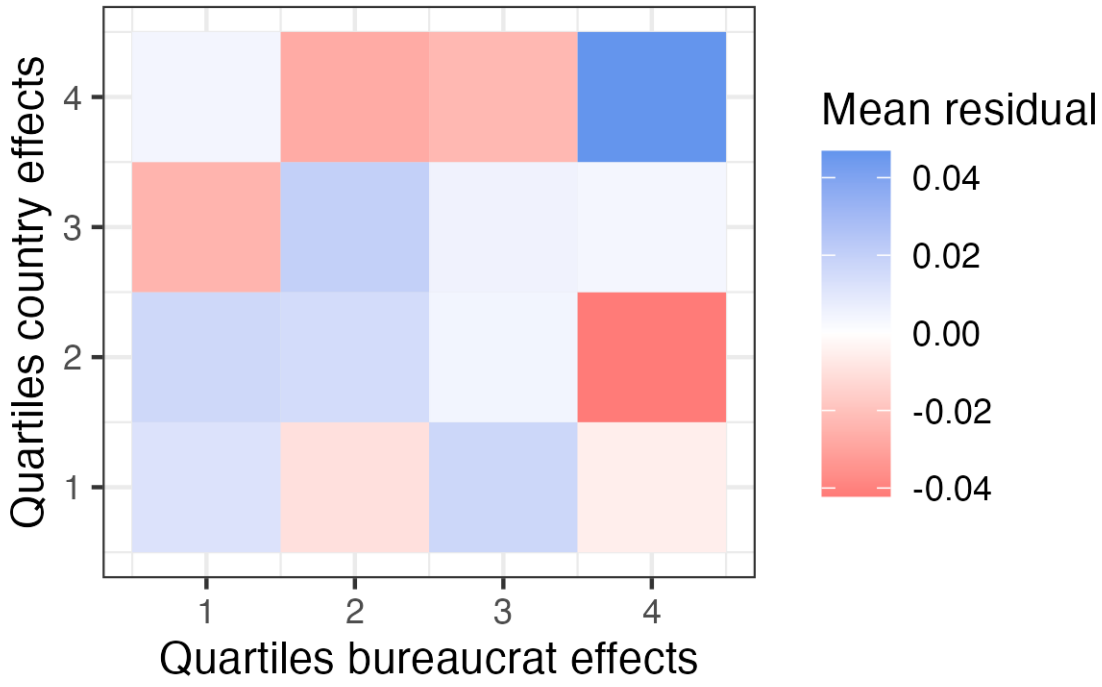

Notes: This figure shows mean residuals from equation (2) with cells defined by quartiles of estimated bureaucrat and country effects.

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

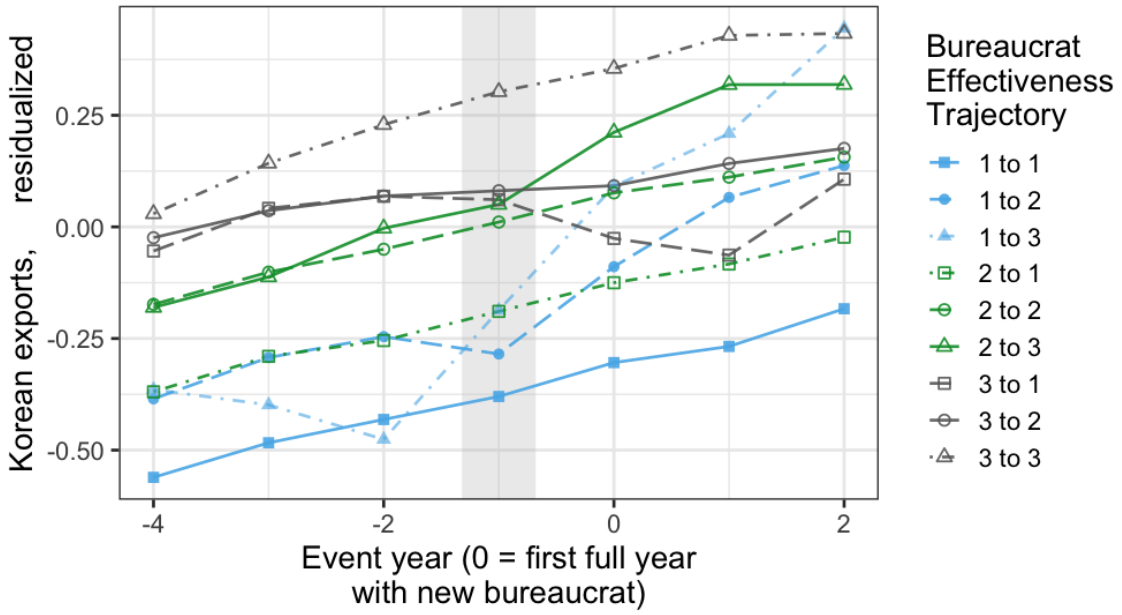

Notes: The figure shows time trends in exports around the time of bureaucrat switches. The switch occurs in year -1 so year 0 is the first full year that the new bureaucrat manages the country office. Year -2 is the last full year with the old bureaucrat. The vertical axis measures the average residualized 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).

C.2 How Estimated Ability Affects the Extensive & Intensive Margin

This section unpacks the effect on the inverse hyperbolic sine (IHS) of exports into the extensive and intensive margin. We find that bureaucrat effects cause increases along both margins. The results are depicted in Appendix Figure C.3.

Panel (a) reports the event study estimates of bureaucrat effects estimated from equation (5) with the outcome variable replaced by a dummy indicating whether the value of Korean exports of a particular product to this country-year exceeds 0. There is no indication of differential pre-trends. In event years 0 and 1, the new bureaucrat's ability increases the likelihood of exporting a given product by 5-7 percentage points, a sizable effect. The old bureaucrat's ability decreases the likelihood by the same amount, suggesting that losing bureaucrat ability has symmetric effects to gaining such ability.

Panel (b) reports the estimates using only the sample of products with changes in the extensive margin during the event horizon. For this sample, the new bureaucrat's ability increases exports by 22-31 percentage points – a very large effect. Losing a bureaucrat has a symmetric effect. There are again no differential pre-trends, especially regarding the effect of the new bureaucrat's ability.

Panel (c) explores the effect along the intensive margin. It replicates Figure 5 using data on only these products for which only the intensive margin matters. As expected, the estimates become noisier. However, pre-trends remain absent, the point estimates go in the expected direction, and are quantitatively similar to Figure 5. Due to the decreased statistical power, only the coefficients on the old bureaucrat's effect remain statistically significant.

Overall, this section shows that a bureaucrat estimated to be high ability – using the IHS of export value – increases both the intensive and extensive margin of exports.

Figure C.3: Unpacking the Export Response to Bureaucrat Switches into Extensive and Intensive Margins

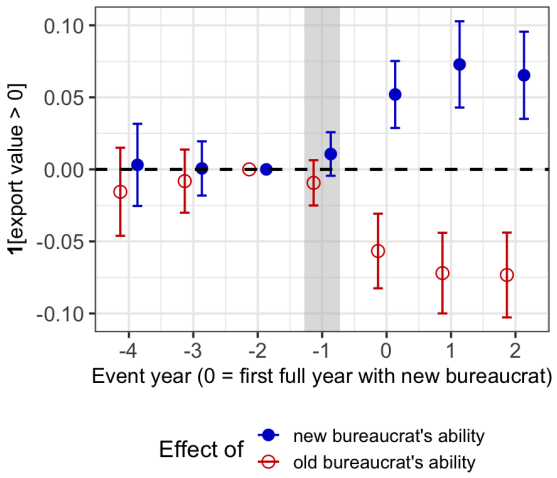

(a) Extensive margin

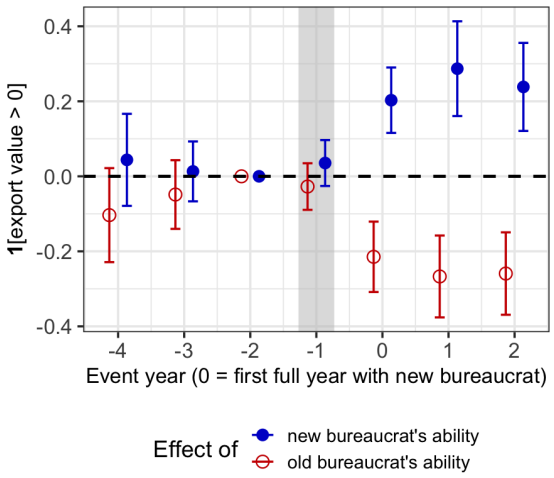

(b) Extensive margin, conditional on a change during the event horizon

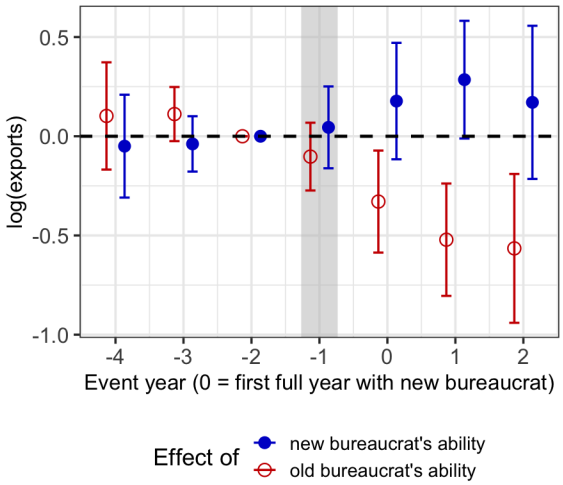

(c) Intensive margin

Notes: The figure shows how the incoming and outgoing bureaucrats' fixed effects affect exports in the extensive and intensive margins, around the time of bureaucrat switches. Specifically, for each plot we plot $\hat{\beta}_k$ and $\hat{\delta}_k$ from equation (5). The outcome variable in Panels (a) and (b) is an indicator for positive export value. For Panel (a), products are included for a given event if South Korea exports this product to any country for all years in the event horizon. For Panel (b), products are included for a given event if South Korea exports this product to this country in some but not all years in the event horizon. For Panel (c), the outcome variable is the inverse hyperbolic sine of export value, and products are included for a given event if South Korea exports this product to this country in all years of the event horizon. 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 (5) includes generated regressors, we report bootstrapped standard errors based on 500 bootstrap iterations.

C.3 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 margin and the intensive margin differently depending on the scale.

First, to focus solely on the extensive margin, Appendix Figure C.4 Panel (a) displays a scatterplot comparing the benchmark bureaucrat fixed effects estimates for $y_{cpt} = \text{exports}_{cpt} = \text{IHS}(\text{export\_value}_{cpt})$, to alternative estimates based on $y_{cpt} = \mathbb{1}\{\text{export\_value}_{cpt} > 0\}$. 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 C.4 Panels (b) and (c) display a scatter plot comparing the benchmark estimates to alternative estimates based on the following transformation:

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

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

dollars to 1,000 nominal 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, clearly evident for $m = 10.00$, 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 valuations of the extensive margin.

Figure C.4: Bureaucrat Fixed Effects Under Alternative Transformations of Exports

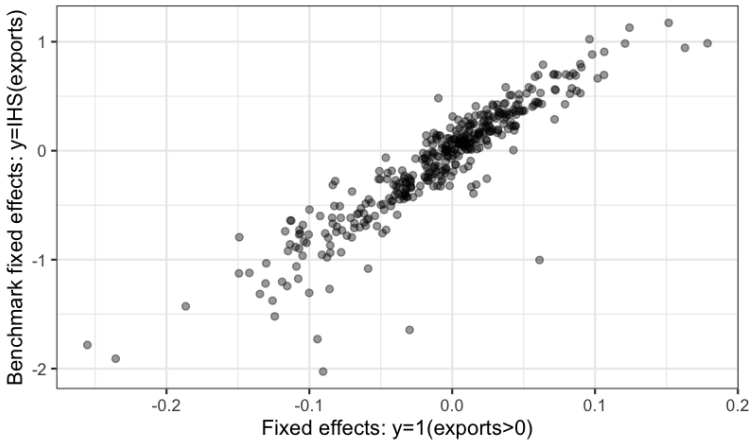

(a) Compared to FE based on the extensive margin

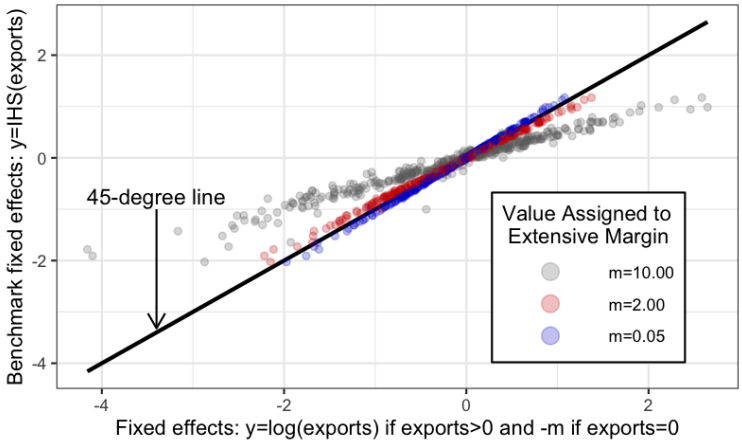

(b) Compared to FE based on explicitly assigning *extreme* values to the extensive margin

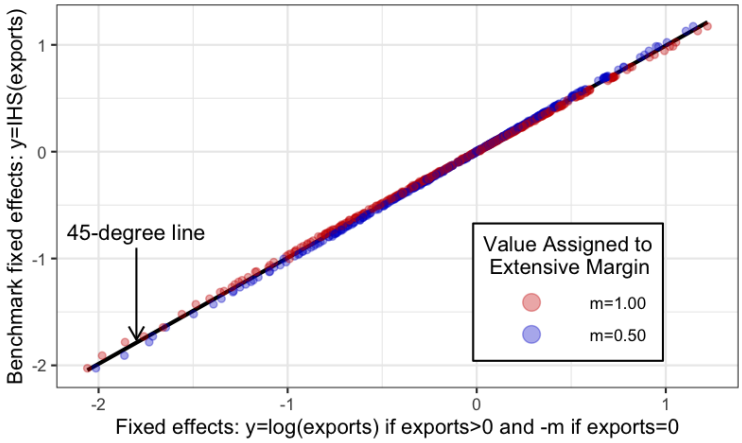

(c) Compared to FE based on explicitly assigning *moderate* values to the extensive margin

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

D Connected Set and Leave-One-Out Connected Set

Appendix Figure D.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, Saggio, and Sølvsten, 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, Saggio, and Sølvsten (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 D.1

illustrate why. They show how many connections the UK has with other countries thanks to just three bureaucrats.

Figure D.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

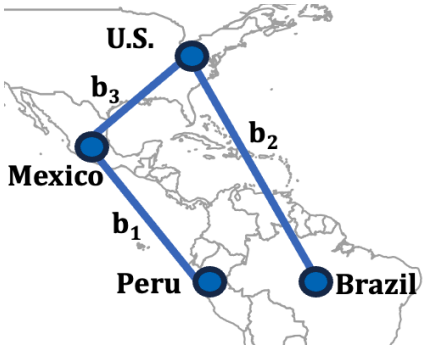

(b) Country-bureaucrat graph composed of a single connected set

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

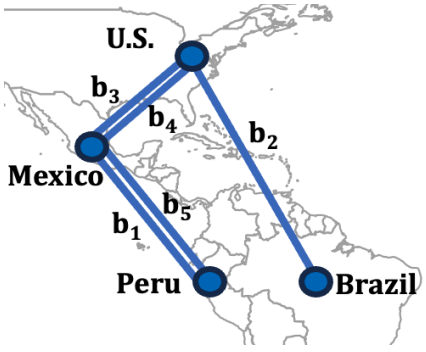

(c) Leave-one-out connected set includes U.S., Mexico, and Peru, but not Brazil

(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.

E 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[83]

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

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, $\mathsf{Var}_b(\hat{\theta}_b) = \mathsf{Var}_b(\theta_b) + \mathsf{Var}_b(v_b) = \sigma_\theta^2 + \sigma_v^2$, will overstate the true dispersion in bureaucrat quality, σ_θ^2.

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

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

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

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[84] equation (2) K times [85] to get the variance of $\hat{\theta}_b$ across bootstrap samples $k = 1, ..., K$:

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

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:

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

$\hat{\sigma}^2_\gamma = 0.682$, 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:

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

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

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

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,

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

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:

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

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

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.

F Robustness of Appointments by Country Importance

This section probes into the finding in Figure 9, that bureaucrats are appointed to the least important countries during their first appointments, the most important countries during their third appointments, with second appointments forming an intermediate case.

Appendix Figure F.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 Figure 9) 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.[86]

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.

[86] 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 F.1: Bureaucrat Flows By Appointment and Country Importance

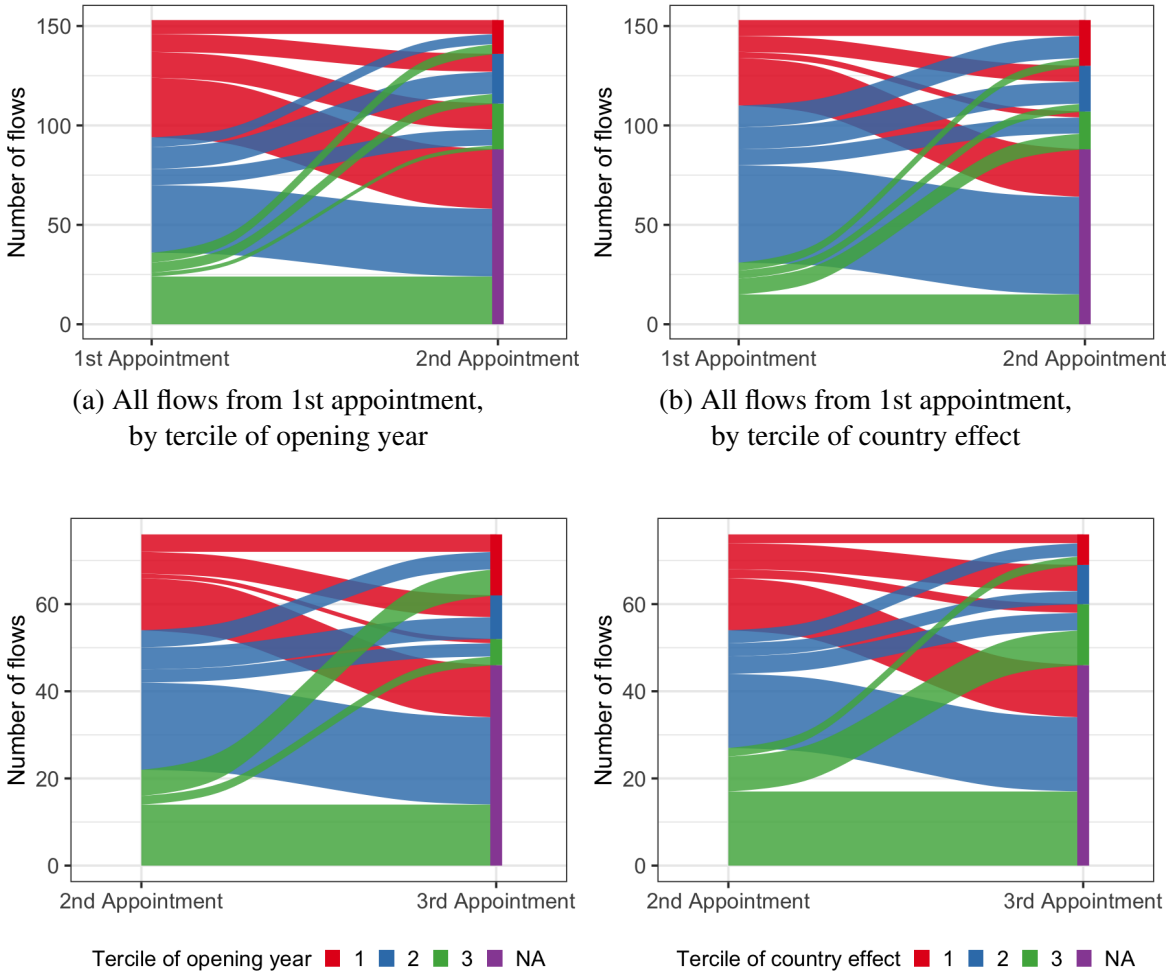

(a) All flows from 1st appointment, by tercile of opening year

(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 roll-out 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 overseas offices 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 roll-out 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.

G Robustness of Office Opening Results

G.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 our main results using never-treated countries as control group, whose lack of differential pre-trends across various sample periods and outcome variables, alleviates concerns that the parallel trends assumption may be violated. 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 G.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 (Figure 10). 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 G.1: Robustness of Office Opening Results to
a Not-Yet-Treated Control Group

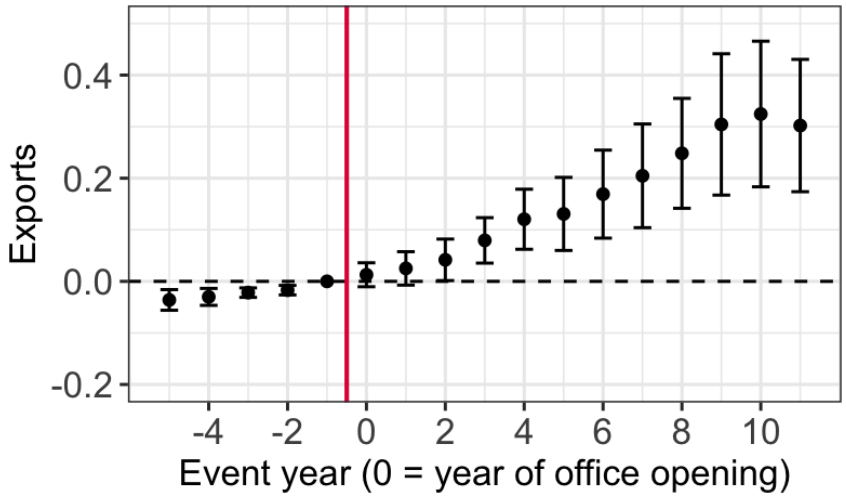

(a) Results with a not-yet-treated control group

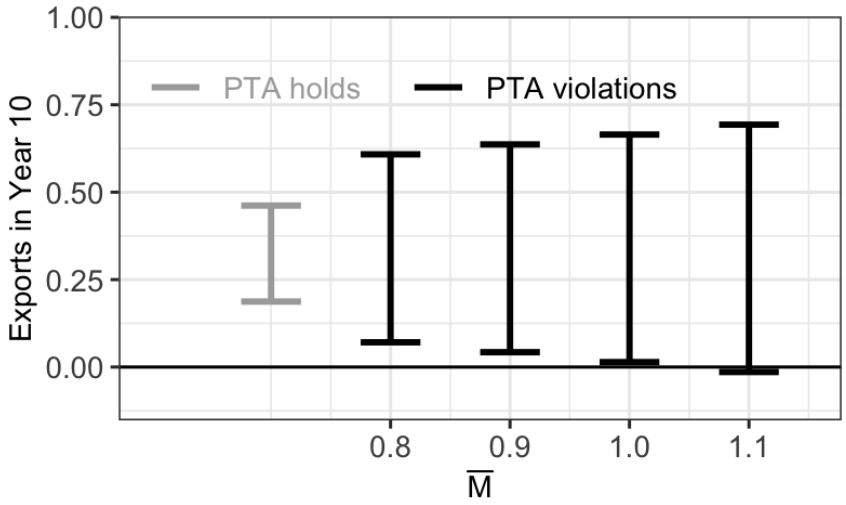

(b) Confidence intervals for the year-10 estimate, allowing for violations of the parallel trends assumption of up to $\bar{M}\times$ the largest pre-treatment violation

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 roll-out. The outcome variable is the inverse hyperbolic sine of Korean export value 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.

G.2 Extensive Margin

This section investigates whether the results in Figure 10 are driven by transforming the raw export values using the IHS, by checking the effect on the product-level extensive margin of exports. Concretely, 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 under investigation becomes: Does a KOTRA office in a country increase the likelihood of Korea exporting a particular 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 G.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 10). Again, the pre-treatment coefficients are very close to zero, corroborating that the parallel pre-trends in Figure 10 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 10, the coefficients stabilize after around ten years.

The results remain similar in Panel (b), where we expand the sample to include openings from 1964.

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

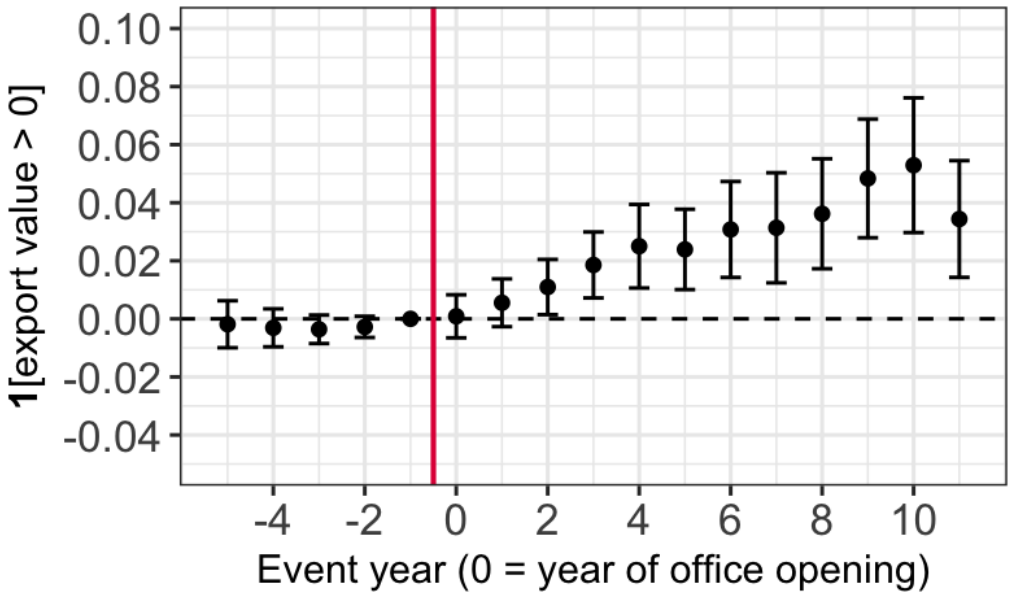

(a) Sample includes openings from 1967

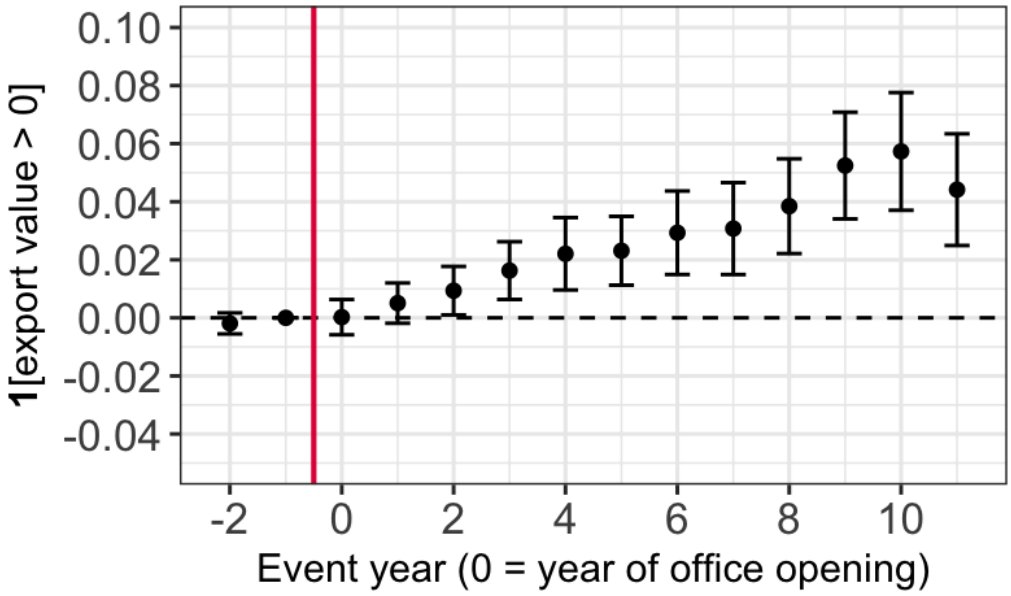

(b) Sample includes openings from 1964

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.