Personnel is Policy (Implementation): Bureaucrats and the Korean Export Miracle*

Philipp Barteska† Jay Euijung Lee‡

March 5, 2025

Please click here for the latest version

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, M50, O1, F13

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

*This paper benefited from comments by Lydia Assouad, 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, 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, Gippeum Lee. We would also like to thank the staff at the KOTRA library, in particular Hyunsun Cho. Financial support from STICERD and the Hayek Programme is gratefully acknowledged.

†Harvard Kennedy School and University of Hong Kong; email: philippbarteska@hks.harvard.edu

‡Department of Economics, Stockholm University, and CEP (LSE); email: jay-euijung.lee@su.se

1 Introduction

Many governments pursue industrial policies, with a marked increase in recent years (Juhász, Lane, Oehlsen, and Pérez, 2022). Proponents often highlight their prevalence during East Asia's unparalleled economic transformation. However, achieving success with industrial policies is challenging because states often lack the necessary information, such as local knowledge needed to tailor policies. The state relies on bureaucrats to collect this information, but their ability and motivation to gather, interpret, and act on it can vary significantly. This variation compounds the challenge for the state 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),[1] little is known about how sensitive these policies are to the bureaucrats implementing them. Investigating how much the effect of an industrial policy depends on implementing capacity has been difficult because it requires the researcher to observe exogenous variation in capacity holding fixed the policy and the economic outcome targeted.

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, i.e., holding fixed time-invariant factors affecting the outcome. The three-yearly rotation further alleviates concerns regarding the potential endogeneity of changes to implementing capacity, both directly and indirectly. Moves of bureaucrats are directly due to the rotation schedule, not other potentially confounding factors. Indirectly, the rotation schedule allows for tests through which we rule out major remaining endogeneity concerns. Third, the principal outcome targeted by each office remained exports to the respective country across the sample period. By studying bureaucrats who targeted exports, a prime economic outcome,[2] this paper stands apart from existing work on

[1] Recent papers that 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; Choi and Levchenko, 2021; Choi and Shim, 2024; Juhász, Lane, Oehlsen, and Pérez, 2022; Kim, Lee, and Shin, 2022).

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

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 strongly depends on the bureaucrat assigned to manage a 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%.[3] Moreover, the variation due to bureaucrats amounts to 1/7 of that due to countries.[4] 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.[5]

The unbiasedness of the bureaucrat fixed effects is supported by numerous diagnostic tests informed by the rigid appointment schedule of bureaucrats. The identifying assumption is that bureaucrat ability is orthogonal to underlying trends in exports. For our main diagnostic check, we perform event studies of exports around switches of office managers, based on the following predictions. First, we should observe differential pre-trends by bureaucrat ability if good bureaucrats were appointed because of underlying trends. Second, the three-yearly 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 less discretion in departures. Empirically, however, we find parallel pre-trends and symmetric effects of gaining and losing a bureaucrat, in support of the identifying assumption. Moreover, in addition to the placebo test described above, we further confirm that the bureaucrat effects are not driven by noise, as they remain stable across different subsamples and retain strong predictive power out-of-sample. Additional diagnostic checks, not discussed here, investigate other potential concerns.

To benchmark the magnitude of the differences in exports due to bureaucrats, we separately estimate the average effect of the policy, using the staggered roll-out of the offices in a difference-in-differences setup.[6] We find that the policy had a substantial effect on average,

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

[4]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).

[5]This alternative bias correction builds on Best, Hjort, and Szakonyi (2023). We make a methodological contribution by fine-tuning the covariance shrinkage estimator and providing insights into its properties in Appendix Section E.

[6]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 roll-out. Recent

increasing exports by 38% ten years after an office opening. Thus, in combination with the main result on bureaucrats, we find that the policy crucially depends on the bureaucrats implementing it; 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.[7]

Appointment patterns suggest that the export promotion organization had the capacity to screen bureaucrats by ability, contributing to the large average effect of the policy. First, bureaucrats are much less likely to be re-appointed if exports underperform during their first appointment as manager of a country office. This suggests the organization learns about a bureaucrat's type during the first appointment. Second, first appointments as country office managers are to less important trade partners, while later appointments tend to be to more important destinations.[8] Hence, the most important offices are rarely run by underperforming bureaucrats.[9] These descriptive patterns align well with archival records on the organization's operations that endorse performance-based pay levels, job assignments, and promotions.[10]

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 in lower tiers of

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).

[7] 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 the exports of the same product from South Korea to other destination countries.

[8] 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.

[9] 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.

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

the state's hierarchy.[11] 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.[12] 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, Creecy, and Kramarz (2002).[13] 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).[14] 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 extensive 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. Similar to industrial policy, most studies on export promotion estimate average treatment effects (Munch and Schaur, 2018; Hayakawa, Lee,

[11] 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 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, 2020; 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 Chetty, Friedman, and Rockoff (2014a,b); Kondylis and Stein (2023); Otero and Muñoz (2025).

[12] 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.

[13] Card, Heining, and Kline (2013); Card, Cardoso, and Kline (2016); Kline, Saggio, and Sølvsten (2020); Bonhomme, Holzheu, Lamadon, Manresa, Mogstad, and Setzler (2023).

[14] as well as more heterogeneous methodologies: 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).

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.

The rest of the paper proceeds as follows. Sections 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, contributing to the large average effect of the policy. 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[15] – to $24,834 (2001).[16] Figure 1 highlights that exports experienced particularly stark growth. Korean exports per capita in 1952 were below 2% of the U.S. level with little convergence between 1952 and 1960. From 1960 on, exports increased rapidly, reaching parity with the U.S. before the end of the century. Over this period, other developing countries saw little growth relative to the U.S.[17]

It is often argued that this growth is partly due to industrial policies conducted by a well-functioning state.[18] On the other hand, the Korean state was described as aid-dependent and corrupt until the 1960s (Kim and Vogel, 2011).[19] 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

[15]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.

[16]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.

[17]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.

[18]Wade (1990) and Cheng, Haggard, and Kang (1998) as cited by Besley et al. (2022); Amsden (1989); Juhász, Lane, and Rodrik (2023). See also the well-known popular book by Studwell (2013). Economists remain broadly skeptical that industrial policies are desirable, including in the case of South Korea's recent history.

[19]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).

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 "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).[20]

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 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).[21] 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 effectiveness of individual bureaucrats in implementing the effect of one important policy pursued by the South Korean government.

[20] 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.

[21] Cheng, Haggard, and Kang (1998) describe the Rhee bureaucracy in unfavourable 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.

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 date until 2000, the end of the sample period.[22]

While the South Korean government had many policies aimed at increasing exports, KOTRA's 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, which were viewed as an efficient way to generate export results in a short time by allowing exporters to engage in direct conversations with local buyers. The overseas offices coordinated logistics and handled strategic tasks such as recruiting, selecting, and briefing exporters. Domestic companies were selected based on newly trending styles and designs that aligned with the market demands at the fair venues. Prior to the fairs, they 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

[22] 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. In 2000, the 88th office opened in Algeria.

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,[23] KOTRA office managers shared similarities with the proverbial "man on the spot" charged with the running of an entire geographic region in the Indian Administrative Service (Bertrand, Burgess, Chawla, and Xu, 2020) or the British colonial administration (Lugard, 1926; Xu, 2018). However, unlike these bureaucrats responsible for a range of policies and outcomes, KOTRA bureaucrats implemented exactly one policy with one 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[24] – 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 Managers to KOTRA's Overseas Offices

Over the entire time period from 1965 to 2000, KOTRA operated 138 offices in 87 countries. The analysis will focus on the main country offices as data on the outcome – exports – is available at the country level.[25]

Official rules do not dictate which bureaucrat gets assigned to which office. The assignment system falls under the discretion of the HR team at the head office. According to interviews we conducted with current and former KOTRA employees, however, there is a general understanding that several factors come into play. The most important factor is language skills; a Spanish speaker is more likely to get sent to a Hispanophone country. Second, 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.

[23] 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).

[24] Archival records and interviews with former bureaucrats offer insight into some broad characteristics on the incentive structure. Salary consisted of a base pay, determined by seniority and position, as well as performance-pay. Performance was primarily assessed at the office level, reflecting the fulfillment of office- and year-specific task allotments and export targets. 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.

[25] The most important or geographically largest countries have multiple offices in different cities. For example, by 1977, KOTRA had 79 overseas offices, of which 64 were the respective country's head office.

Organizational rules do, however, provide substantial rigor regarding the timing of appointments. The regular nature of these directors' appointments is highlighted by the fact that both the modal and median appointment duration is 36 months – three years. Appendix Figure A.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 re-appointment is also largely predetermined: The median duration for the gap between appointments is 29 months, the modal gap is 30 months. Appendix Figure A.3 plots the distribution of gaps between appointments. This rotation limits discretion in appointments as not all bureaucrats are available when a particular country is due to receive a new bureaucrat.

More importantly, the rotation schedule provides exogenous variation in the bureaucrat appointed to manage a country office. In particular, while there is discretion to the decision of appointing bureaucrat b to country c in year t, this decision then largely pre-determines losing bureaucrat b in year $t+3$.

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.[26] 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. In our study, we remove the effect of such policies by netting out product-year fixed effects.

3 Data

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

[26]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.

3.1 Bureaucrat Appointments

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

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

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*.[27] However, this is remedied by a great diversity in first names.[28] 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 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 director'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.

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

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

3.2 Exports

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

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. 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 AKM framework to study how much bureaucrats matter in explaining South Korean exports (Abowd, Kramarz, and Margolis, 1999). 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)[29] 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 ($\theta_{b(c,t)}$),[30] a country component (γ_c), and an error term (ϵ_{cpt}). As in other parts of the paper, exports are at the product-level,

[29] We explore robustness to the IHS transformation. We show in Appendix Section C.2 that bureaucrat fixed effects are predictive of export changes along both the extensive and intensive margins. Furthermore, we show in Appendix Section 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).

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

in line with KOTRA's export targets being typically product-specific. It also avoids a few dominant export products driving the results for a country-year.

The AKM estimation based on equation (1) identifies the bureaucrat and country fixed effects only within the largest connected set.[31] 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 736 appointments of 399 bureaucrats to 87 country offices. The largest connected set among these contains all appointments to 86 out of 87 countries.[32] This connected set contains 735 appointments of 398 lead bureaucrats, of whom 184 had appointments to multiple offices.[33]

The bureaucrat-country graph is interconnected enough such that 75 countries and 93% of appointments form a leave-one-out connected set. The reason is that most country offices remain open for decades, headed by many different office managers. Column (3) indicates that in the leave-one-out connected set, 72 offices have four or more distinct lead bureaucrats, 61 offices have six or more bureaucrats, 49 offices even have eight or more distinct bureaucrats.

Our preferred estimation uses only the appointments in the largest leave-one-out connected set. That way, we can apply Kline, Saggio, and Sølvsten (2020)'s correction method for limited mobility bias in our variance decomposition exercise. Moreover, Kline, Saggio, and Sølvsten (2020) show that restricting attention to the leave-one-out connected set substantially reduces limited mobility bias directly.

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-yearly rotation schedule. As highlighted in Appendix Figure A.2, a new appointment to country c in year

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

[32] Only Cambodia is outside the largest connected set because is only bureaucrat was never appointed to any other country. The data only contains one appointment to Cambodia because its office opening occurred shortly before the end of our sample period.

[33] 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 Cambodia).

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. Section 4.4 tests this hypothesis in an event-study setup for export changes around the event of bureaucrat switches. We find parallel pre-trends.

Second, while KOTRA has some discretion in new appointments, the rigid rotation schedule dictates the year when a country *loses* a 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. Section 4.4 tests this hypothesis 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 result also rules out other alternative hypotheses that 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 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. In qualitative interviews, KOTRA employees mention three additional constraints on the appointment decisions. (1) Bureaucrats are more likely to be appointed to a country when they speak the local language. (2) Bureaucrats prefer being appointed to high-income, English-speaking countries. Because these preferences are commonly shared, KOTRA's Human Resources manages discontent by rotating bureaucrats between low- and high-desirability locations. (3) 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 and importance, both relative to other countries, change little over time. Therefore, while appointments may be correlated with country fixed effects,[34] the scope for appointing bureaucrats based on anticipated export trends is limited.

Despite these constraints, one may still wonder why KOTRA does not strategically ap-

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

point good bureaucrats to countries with growing import demand. There would be no reason to, if there are no match effects between bureaucrats and countries, as specified in equation (1) and verified through various misspecification checks in Appendix Section C.1. Even if KOTRA expected there to be positive match effects, it could be that KOTRA prioritized time-invariant country characteristics over trends in gauging export potential, much like it did for office openings.[35]

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).[36] 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.[37]

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 sampling error in $\hat{\theta}_{b(c,t)}$, which is aggravated when bureaucrats move countries too infrequently (limited mobility bias). Kline, Saggio, and Sølvsten (2020) estimate the degree of bias in $\mathsf{Var}(\hat{\theta}_{b(c,t)})$ under unrestricted heteroskedasticity, building on previous approaches which required homoskedastic error terms (Andrews, Gill, Schank, and Upward, 2008).

Our preferred variance decomposition approach follows the bias correction method of Kline, Saggio, and Sølvsten (2020). We use the computational algorithm of Bonhomme, Holzheu, Lamadon, Manresa, Mogstad, and Setzler (2023) for implementation. Although unreported, the Andrews, Gill, Schank, and Upward (2008) correction method that assumes

[35] The roll-out of export promotion offices mainly followed pre-determined gravity variables, as reported in Figure 11 and discussed in Section 5.3.

[36] 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 $\hat{\lambda}_{\text{cars},1995}$. Table B.2 demonstrate the importance of product-year trends, which explain 35.5% of the total variation in exports.

[37] 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.

homoskedasticity delivers quantitatively very similar results.[38]

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.[39] 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.[40] We bootstrap equation (2) at the appointment level to estimate the variances and covariances of the sampling errors.[41] 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 estimates of the variances and covariances of sampling errors, we can calculate the bias in the plug-in estimators, $\mathsf{Var}(\hat{\theta}_{b(c,t)})$, $\mathsf{Var}(\hat{\gamma}_c)$, and $\mathsf{Cov}(\hat{\theta}_{b(c,t)}, \hat{\gamma}_c)$. We subtract the estimated biases to isolate $\mathsf{Var}(\theta_{b(c,t)})$, $\mathsf{Var}(\gamma_c)$, and $\mathsf{Cov}(\theta_{b(c,t)}, \gamma_c)$. We call this method "bootstrap bias correction."[42] This method builds on Best et al. (2023) by recognizing that the variances or covariances of the shrunk fixed effects, whether from standard shrinkage or covariance shrinkage, do not consistently estimate the true values.[43] In fact, the variance of the covariance-shrunk bureaucrat fixed effects theoretically underestimates the true variance,[44] while in practice, it might overestimate the true variance.[45]

[38] 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, Gill, Schank, and Upward (2008) correction method can also be performed on the largest connected set covering 86 countries and 397 bureaucrats. The two methods delivers extremely similar results for either measure of connectedness.

[39] Although each fixed effect is unbiased, it is noisy. By shrinking them towards the mean value, we introduce bias for reduced variance.

[40] We discovered a minor error in Best et al. (2023)'s derivation of the covariance shrinkage matrix. We use the corrected matrix for our covariance shrinkage. The resulting covariance-shrunk fixed effects are mostly similar to the Best et al. (2023) counterparts but sometimes diverge considerably. See Appendix Section E.3 for our derivation and E.6 for the comparison with Best et al. (2023)'s.

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

[42] Derivations are in Appendix Section E.1.

[43] Best et al. (2023) explain that covariance shrinkage is their preferred variance decomposition approach because the correlation between the covariance-shrunk bureaucrat and organization fixed effects is positive, a more plausible outcome, while other approaches they consider yield negative values. Yet, the positive correlation may partly be mechanically driven by positive off-diagonal elements of the covariance shrinkage matrix. In our setting, we find that $\mathsf{Cov}(\hat{\theta}^{\text{CovSh}}_{b(c,t)}, \hat{\gamma}^{\text{CovSh}}_{c}) > 0$ while $\mathsf{Cov}(\theta_{b(c,t)}, \gamma_c) < 0$. See Appendix Section E.4 for why the off-diagonal elements of the covariance shrinkage matrix may likely be positive.

[44] Shown in Appendix Section E.3.

[45] Appendix E.5 describes the conditions under which a few covariance-shrunk fixed effects might blow up, which would expand the variance.

4.3 Result: Bureaucrats Drive Korean Exports

Figure 3 plots the cumulative distribution function of raw bureaucrat fixed effects. It shows substantial dispersion, although it overstates the true dispersion. Appendix Figure A.7 compares the distributions of the raw fixed effects and the covariance-shrunk fixed effects.

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 column 3-4. The two bias correction methods yield similar results overall. The variance of the bureaucrat component is remarkably similar.

Bureaucrats explain a substantial amount of variation in Korean exports. Columns 2 and 4 show that they explain around 14% of the variation. Furthermore, column 1 shows that one standard deviation of bureaucrat ability is estimated to be 0.318,[46] implying an increase in the dollar value of exports of 37%.[47] 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.[48]

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. The 1/7 number is relatively small compared to effect sizes in recent studies of public-sector managers using similar methodolody. 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.[49] 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 have been found to be very strong predictors of bilateral trade volume. It should be noted, however, that public hospitals, courts, and retail stores also represent more than just the organization, e.g., the demographics of their location.

[46] $0.318 = \sqrt{0.101}$

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

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

[49] 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.

We find a negative covariance between bureaucrat and country fixed effects, suggesting that better bureaucrats work in smaller countries. In the bootstrap bias correction (column 3), the covariance is still negative, but is smaller and statistically insignificant. The negative covariance is almost exclusively driven by single-appointment bureaucrats.[50] 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 they are not driven by noise. Additional diagnostic checks are provided in Appendix Section 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 results 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.

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

First, to test for differential pre-trends, we estimate equation (4), which models how exports change around the event (e) of bureaucrat switches for different terciles of the change in bureaucrat ability induced ($\Delta\hat{\theta}_e$).[51] 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

[50] 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.

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

to country c. 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 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 distinguishing between $\hat{\theta}_e^{\text{new}}$ and $\hat{\theta}_e^{\text{old}}$, the specification is similar to equation (4). As argued in section 4.1, the rotation schedule means we would expect stronger "effects" in absolute terms due to 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 in the direction of the ability of the incoming bureaucrat and symmetrically against the direction of the outgoing bureaucrat's ability. 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, this drop is only relative to South Korean exports to other countries. Given that South Korean exports were growing at more than 35% annually[52] and all the regression equations include product-year fixed effects, losing a good bureaucrat means exports drop *only relative to this trend*.

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

[52]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%.

data, an issue aggravated when the number of observations per bureaucrat is small. The variance decomposition results discussed in Section 4.2 already account for this noise, but we perform a series of checks to reinforce the reliability of the estimates.

First, we perform a placebo check. We apply the variance decomposition on synthetic data where bureaucrats are randomly shuffled across the 735 appointments in the connected set. 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.[53] 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.[54] Remarkably, both the variation in bureaucrat fixed effects, as well as the covariance between bureaucrat and country fixed effects, go to zero, indicating that the Kline et al. (2020) bias-correction method successfully removes the contribution of noise to the variance and covariance values. This finding shows 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 fixed effects' out-of-sample predictiveness. An implication of overfitting is that the fixed effect estimates have low predictive power outside the sample. We directly test and confirm that our estimated bureaucrat effects are, indeed, predictive out of sample.[55] The most natural and conservative way in our setting to obtain fixed effects that are testable out of sample is to only use *other countries* to estimate the fixed effects. For example, to estimate the fixed effects of bureaucrats appointed to the UK, we obtain their fixed effects when excluding the UK from the sample. The estimation of such out-of-sample fixed effects is only possible due to the interconnectedness of our data. When leaving out one country, we always retain one very large connected set, as 75 countries in our data are part of the same leave-one-out 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 0.52.[56] Furthermore, we replicate the event-study regression depicted in Figure 5 using the out-of-sample fixed effects. Appendix Figure

[53] We identify the set of unique bureaucrat-country pairs in the original data, randomly shuffle the matches, and merge it into the original data. For example, if bureaucrat 1 was appointed to country A from 1965 to 1968, a random bureaucrat X replaces bureaucrat 1 in country A for these years.

[54] 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.

[55] Additionally, Table B.2 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.

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

A.8 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. In addition to alleviating overfitting concerns, these findings also strengthen the support for the identifying assumption. The fact that estimated bureaucrat effects are predictive out-of-sample makes it implausible that they are primarily capturing correlations between bureaucrat appointments and underlying export trends.

Third, we conduct a robustness check using the bootstrap bias correction. The main limitation of the out-of-sample predictiveness test is that it cannot be applied to single-appointment bureaucrats, who constitute about half of all bureaucrats and whose fixed effects are the most susceptible to overfitting. The leave-i-out estimation of the bias by Kline et al. (2020) faces the same challenge, so the i-specific variance of the sampling error for a single-appoint bureaucrat is assumed to equal the average among multi-appointment bureaucrats ever appointed to the same country – possibly an underestimate. In contrast, the bootstrap bias correction does not require this homogeneity assumption, as bootstrapping equation (2) yields variance estimates for all bureaucrats.[57] The fact that both correction methods produce extremely similar estimates of the true dispersion in bureaucrat ability (columns 1 and 3, Table 2) offers confidence that we are appropriately accounting for noise also for single-appointment bureaucrats.

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.[58] 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 directors 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 directors (14%, column 2).

Overall, these results demonstrate that bureaucrat fixed effects reliably identify the causal impact of an individual bureaucrat on exports.

[57] As expected, the variance of the sampling error of bureaucrat fixed effects obtained this way steadily decreases in the number of appointments per bureaucrat.

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

4.5 Mechanism: Good Bureaucrats Tap into Import Demand

In this section, we investigate how a high-ability bureaucrat increases exports by more than another bureaucrat of lower ability. 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 to other countries from Korea. 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 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 very 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 timely 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,

similarly to Figure 4, support the orthogonality of bureaucrat appointments and export trends.

4.6 Policy Implications: Screening Bureaucrats and Policy Effect

Bureaucrats are Screened Based on Performance in 1st Appointments

Our findings thus far show that individual bureaucrats' ability significantly impacts the effectiveness of 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) bureaucrat fixed effects (covariance-shrunk), and (b) residualized exports by appointment, splitting the sample by whether a bureaucrat has more than one appointment as director of a country office over their career. The distributions exhibit a substantial mass of lower-ability bureaucrats whose careers only include one appointment.[59]

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 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 the most significant markets (Appendix Figures A.14 and A.15). A country's importance is proxied by how early its KOTRA office opened and how high its fixed effect is.

Together, these two patterns align with a model where the organization learns about bureaucrat ability by initially appointing them to less important countries, and then uses the

[59] We use covariance-shrunk bureaucrat fixed effects to correct for the fact that the raw fixed effects of single-appointment bureaucrats are noisier. Furthermore, panel (b) 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.

information to select out low performers from getting reappointed, particularly to important countries.[60] An appointment as a manager of a small country office is plausibly much more informative about a bureaucrat's ability than a subordinate role in a more important country. Assigning bureaucrats to manage small countries may therefore 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 head more appointments, the overall effect is negative, at -7.1%. 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% 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, following Best, Hjort, and Szakonyi (2023). This counterfactual is based on individual bureaucrat fixed effects, as opposed to the average ability of groups of bureaucrats.

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

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 10%, an even larger gain than counterfactual (2).

These counterfactuals highlight the importance of screening processes in strengthening bureaucratic capacity and suggest that further 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 roll-out of overseas export promotion offices to identify the causal effect of the offices on Korean exports to the respective destination countries.[61] 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,[62] it follows that an office with a bureaucrat one standard deviation below average would have no effect on exports.

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.[63] The comparison is possible because (1) we observe a sufficient number of office openings, and (2) South Korean exports to a country is a well-defined variable even without an office.

The staggered roll-out of offices is depicted in Figure A.1, which plots the cumulative number of offices by year. There were no offices before KOTRA was established in 1962. Over the next two decades, KOTRA opened offices in 75 countries – nearly four new countries every year. After the intensive initial roll-out, KOTRA's expansion came to an abrupt halt, with only three new countries opening their first office over the next seven years.[64] 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 officesto some countries and not others. As this is not feasible, the

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

[62] Recall that mean bureaucrat ability is normalized to 0 in the estimation of equation (2).

[63] 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).

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

analysis here will use the staggered roll-out 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, γ_{cp} indicates country fixed effects that may differ at the product-level. D^k_{ct} are dummies equal to 1 if year t is k years after the first office opened in country c. π_k corresponds to the effect of an office that has been open for k periods. X_{cpt} is a vector of time-varying controls. The main specification uses the IHS of the value of South Korean exports as the outcome variable and does not include any controls.

Two central assumptions 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.[65] 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 roll-out 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 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.[66]

[65] 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).

[66] Positive anticipation may occur if firms decide to export to a market today because KOTRA will open

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, 2023; 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 sample. 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. We have five pre-periods in the baseline specification, and our export data starts in 1962, so the countries whose first office opened between 1962 and 1966 do not contribute to the estimated effect of offices.[67] Appendix Figure A.18 shows that the pre-trends remain small and effect sizes remain comparable when we have four or two pre-periods.

For countries that experience an office opening between 1967 and 1981, we include all observations from five years before to eleven years after the office opening. Hence, the 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.[68]

An obvious 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 all never-treated economies with a population below one million. Furthermore, Appendix Section G.2 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 9 plots $\hat{\pi}_k$, the estimated effects of the first overseas export promortion office in a destination country around the year of the opening. The coefficients in the pre-periods are

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 for any positive anticipation.

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

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

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%[69] higher export value relative to the control group.

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, 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).

5.3 Robustness

This section discusses potential violations of the parallel trends assumption.

Appendix G additionally discusses in detail the robustness of the results to (1) including the countries with office openings between 1964 and 1966, (2) using a not-yet-treated control group following Callaway and Sant'Anna (2021), and (3) considering only the extensive margin of exports – to show that the results are 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 10.[70] The point estimates are very close to zero both before and after an office opening, 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. Panel (a) of Appendix Figure A.18 shows 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.[71] Panel (b) shows that estimates remain unchanged when the effect of non-Korean exports is allowed to vary by year, by controlling for non-Korean exports interacted with year dummies.

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

[70] 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.18).

[71] 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.

Roll-out 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 they are not, the pre-determined order of the roll-out makes it unlikely that office openings were timed to coincide with counterfactual increases of exports.

To predict office openings, we draw on insights from the gravity equation, where distance and market size are the strongest predictors of bilateral trade flows. Apart from the U.S., the first KOTRA offices opened in Taiwan, Thailand, Japan, Singapore, Indonesia, and South Vietnam – among the geographically closest non-communist countries and territories.[72]

Within Europe, the distance from South Korea does not vary much between countries,[73] so the main predictor would be market size.[74] 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 11 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. Appendix Table B.1 further reports predicted opening years for the 17 countries based on their 1962 market size. The actual and predicted years of opening are close and often coincide exactly. The rigidity in the roll-out schedule counters concerns that KOTRA timed openings based on export trends, a violation of the parallel trends assumption.

6 Conclusion

This paper identifies a setting that allows us to show that bureaucratic capacity changes the effect of an economic policy on exports. The setting is of substantial intrinsic interest as it allows us to quantify narratives linking state capacity and the East Asian miracle – an episode with utmost potential to inform development economics as it includes the only transitions from low to high income since the Second World War.

First, 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. This study further broadens the types of bureaucrats represented in the economics literature by studying bureaucrats with substantial autonomy whose responsibility lies at a level between frontline service provision and the high-level drafting of policies.

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

Second, the findings imply that implementation matters substantially in determining the effect of an industrial policy. This may translate to other policies that require a strong discretionary component, tacit knowledge, or frequent exchange of information with firms. This adds nuance to the resurgent debate on industrial policy by qualifying how much we learn from policy success in one setting – even if economic conditions are similar.

Third, we find support that the organization screens bureaucrats based on their performance during the first appointment as country managers, suggesting a path towards managing uncertainty over ability or effort. This possibility provides a rationale for policymakers to choose policies where it is possible to establish such links between individuals and the most relevant policy outcomes.

Bibliography

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

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

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, 09). 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, February). 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. (2020). In Self Interest? Meritocracy in a Bureaucracy. Working paper.

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

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

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

Angrist, J. D. and J.-S. Pischke (2010, June). 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*.

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. Technical report, unpublished Ph.D. thesis, University of Minnesota.

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

Bertrand, M., R. Burgess, A. Chawla, and G. Xu (2020). The Glittering Prizes: Career Incentives and Bureaucrat Performance. *Review of Economic Studies 87*(2), 626–655.

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

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

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

Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2023). How Much Should We Trust Estimates of Firm Effects and Worker Sorting? *Journal of Labor Economics 41*(2), 291–322.

Borusyak, K., X. Jaravel, and J. Spiess (2023). Revisiting event study designs: Robust and efficient estimation. *Review of Economic Studies*.

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, 12). 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, 03). Selecting Top Bureaucrats: Admission Exams and Performance in Brazil. *The Review of Economics and Statistics*, 1–47.

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*.

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.

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). *The Transition to Export-led Growth in South Korea: 1954-1966*. Studies in the modernization of the Republic of Korea, 1945-1975. Harvard East Asian monographs ; v. 91.

Juhász, R. (2018). Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade. *American Economic Review 108*(11), 3339–76.

Juhász, R., N. Lane, E. Oehlsen, and V. Pérez (2022). 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*.

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.

Kondylis, F. and M. Stein (2023). The Speed of Justice. *The Review of Economics and Statistics 105*(3), 596–613.

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, July). 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.

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.

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).

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

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

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

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.

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

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

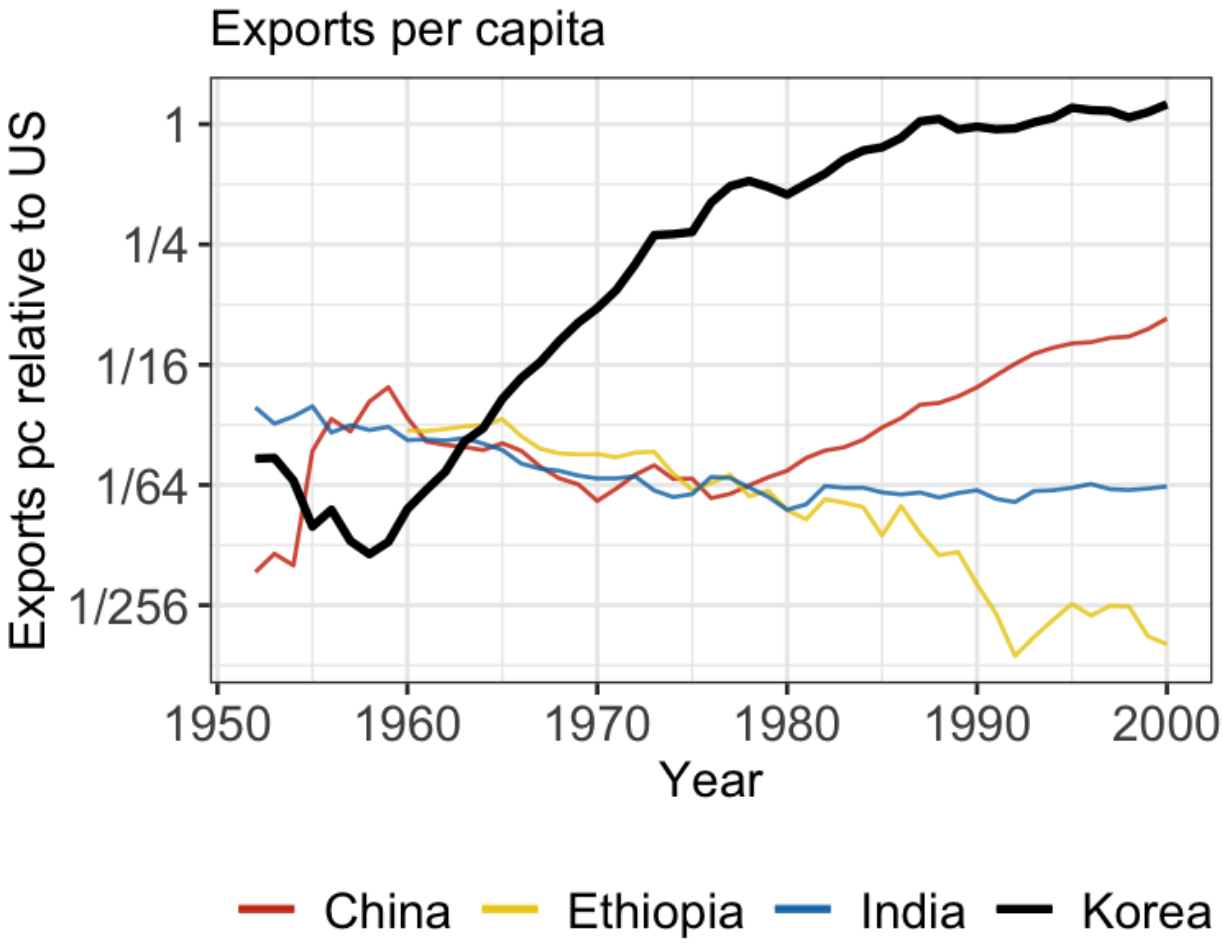
Volpe Martincus, C. and J. Carballo (2012). Export promotion activities in developing countries: What kind of trade do they promote? *The Journal of International Trade & Economic Development 21*(4), 539–578.

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

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

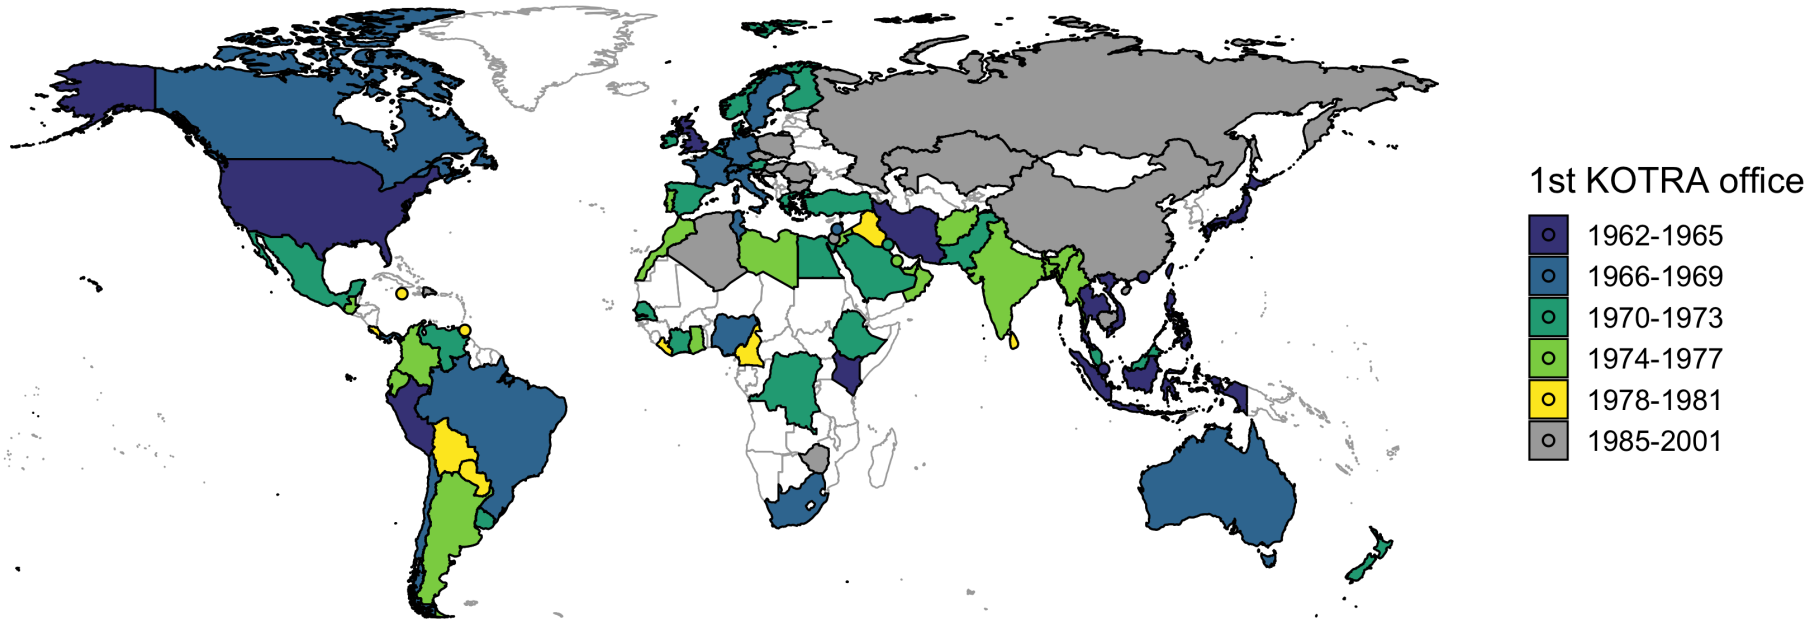
Xu, G. (2018). The Costs of Patronage: Evidence from the British Empire. *American Economic Review 108*(11), 3170–98.

Figure 1: Growth in Korean Exports



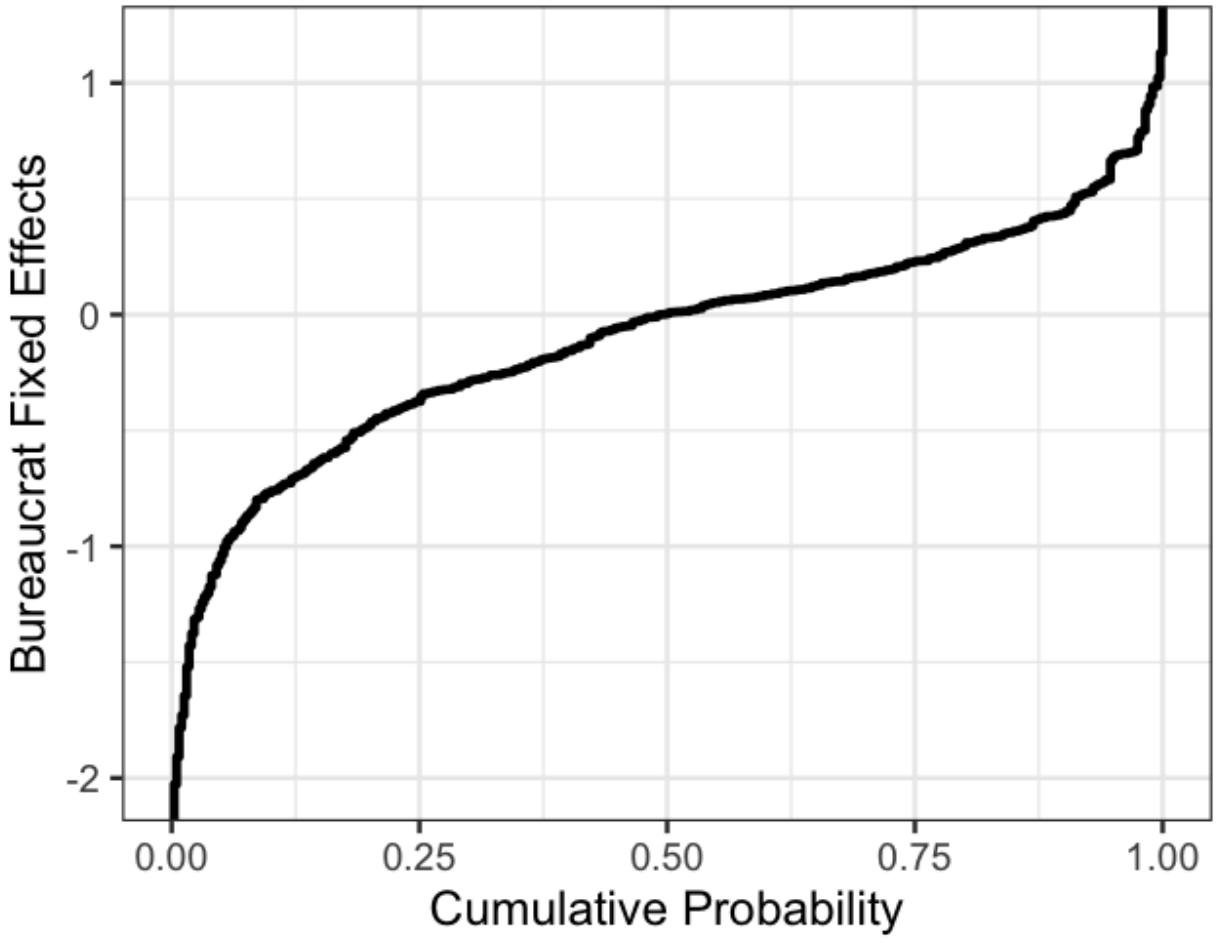
Notes: The figure displays exports per capita relative to the U.S. from 1952 to 2000, for South Korea and a selected group of other countries. Data on exports and population are obtained from International Monetary Fund (2023): Direction of Trade Statistics.

Figure 2: Countries with Korean Export Promotion by KOTRA Office



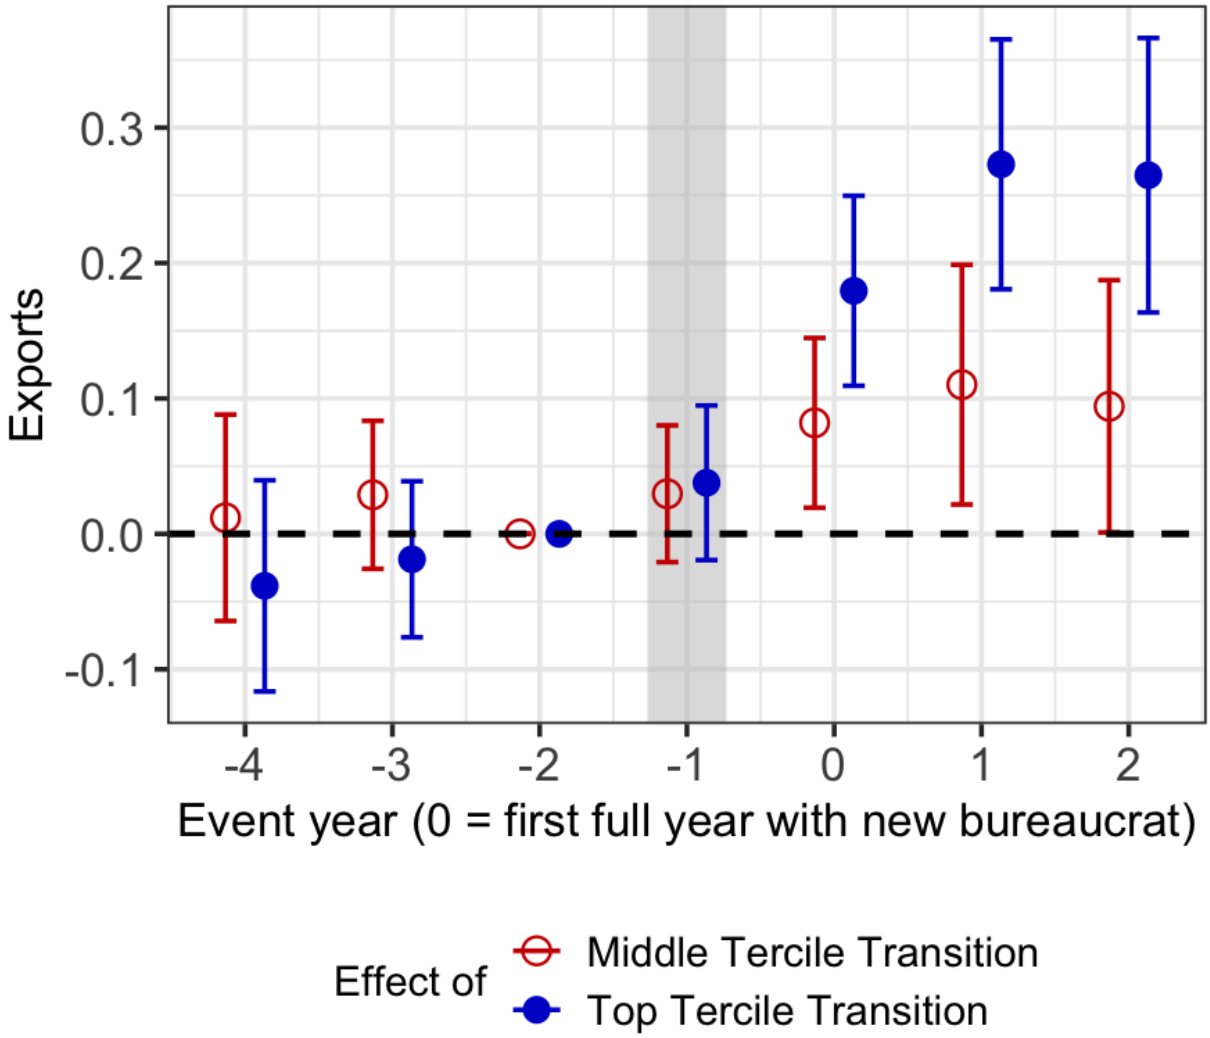
Notes: Colors indicate the opening year of the first Korean export promotion office in the country. Most offices remain open from the opening year for the entire period of study. White indicates countries which have not received an office by 2001.

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



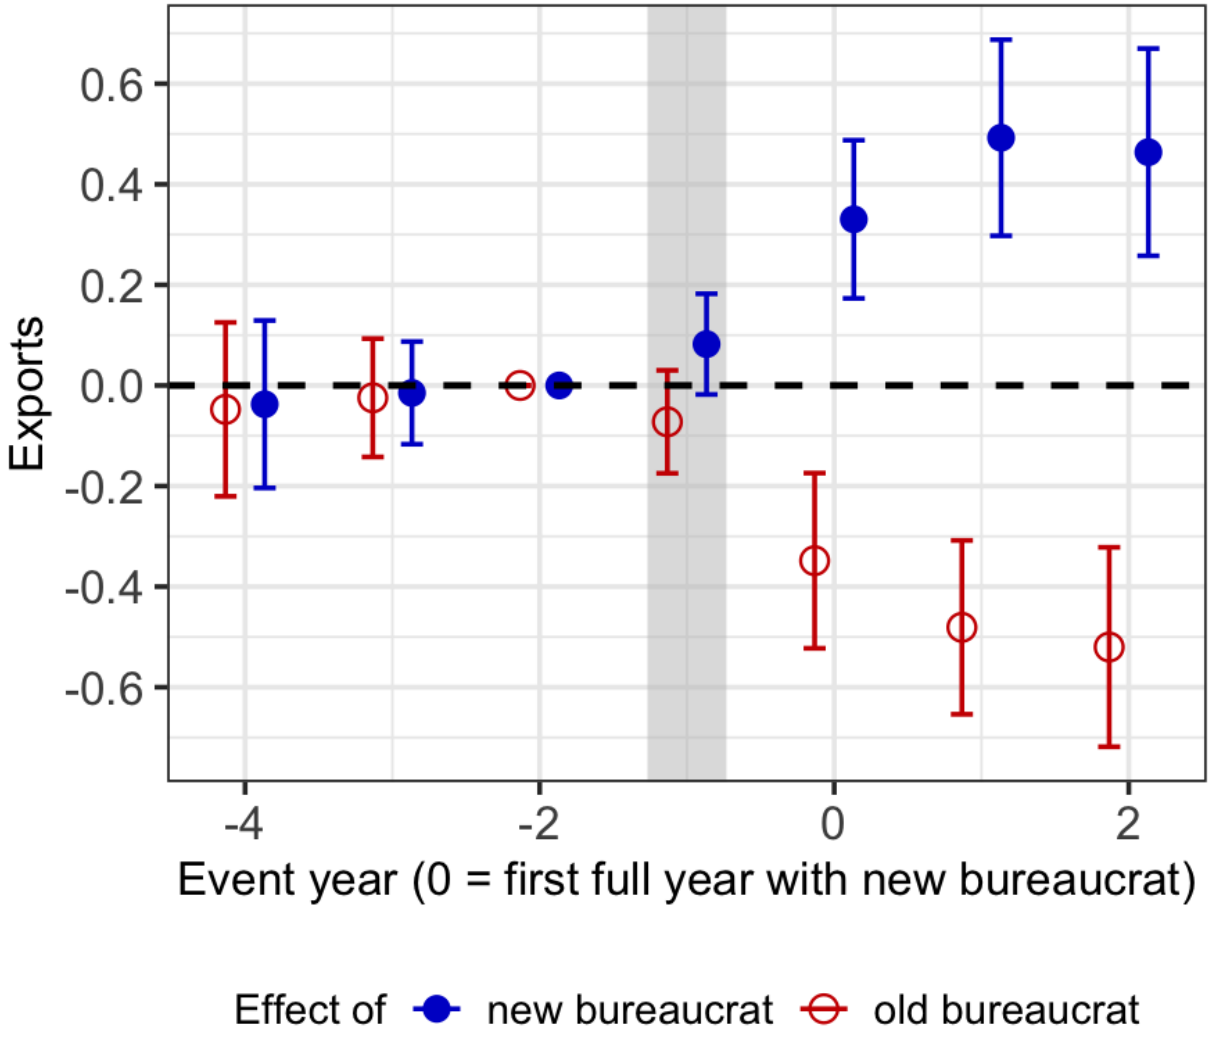
Notes: The figure shows the cumulative distribution function of bureaucrat fixed effects, estimated from equation (2).

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



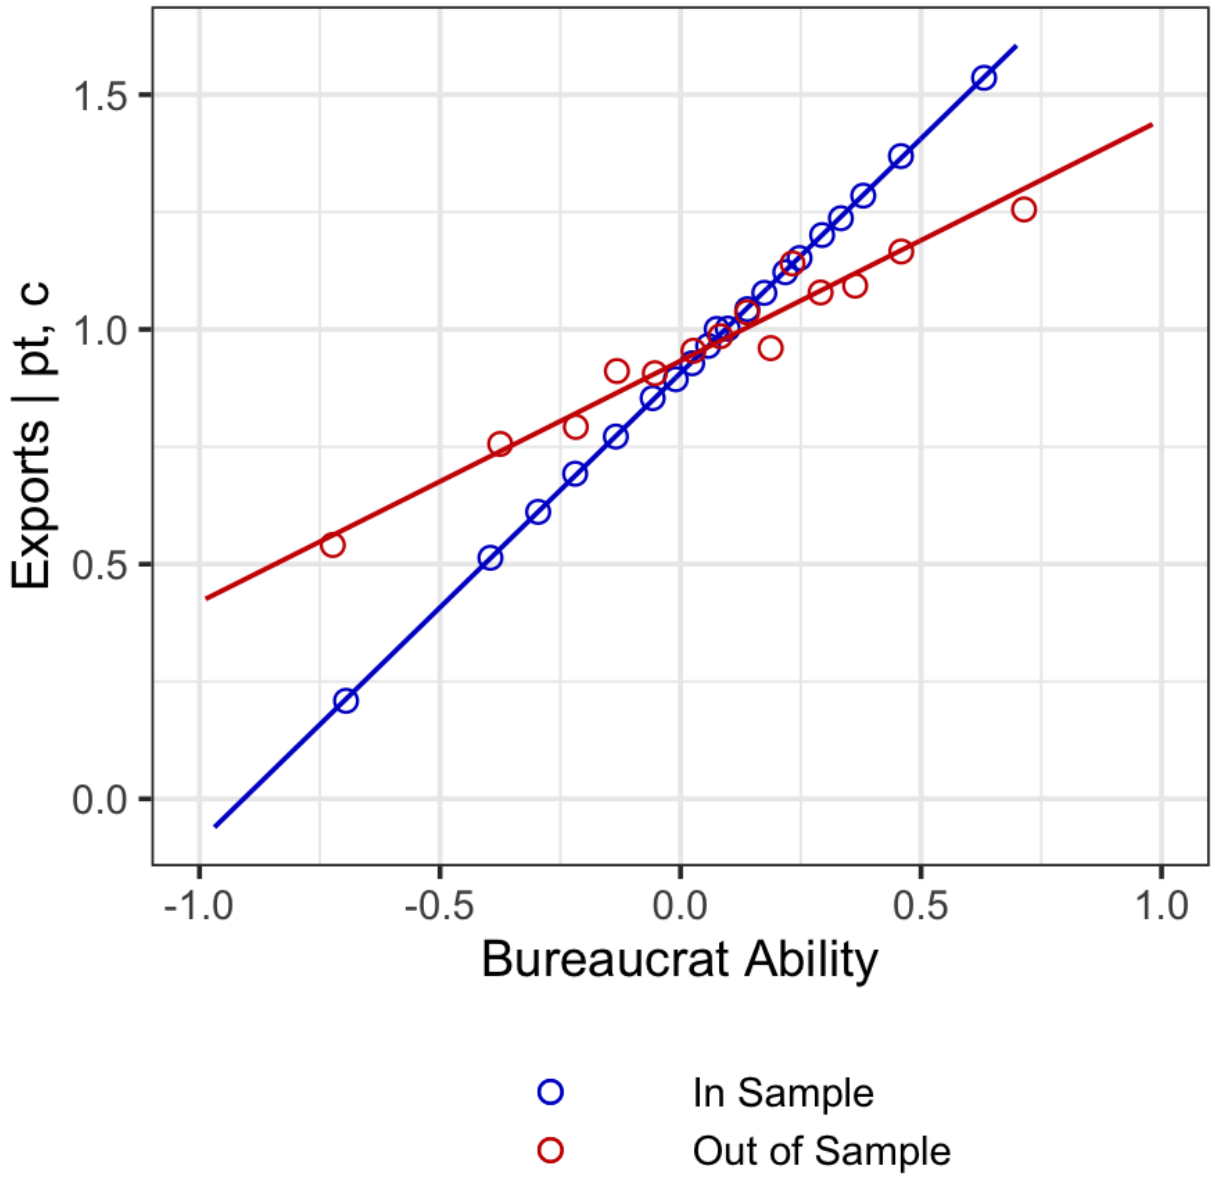
Notes: The figure shows the effect of the change in bureaucrat fixed effects on exports, around the time that the bureaucrat managing a country office changes. Specifically, we plot $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (4). The outcome variable is the inverse hyperbolic sine of exports 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 that 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 the regressors of equation (4) include functions of estimated objects, we report bootstrapped standard errors based on 500 bootstrap iterations, which are found to be more conservative than standard errors clustered by country (unreported).

Figure 5: Symmetric Effects from Gaining and Losing a Bureaucrat.
No Differential Pre-Trends.



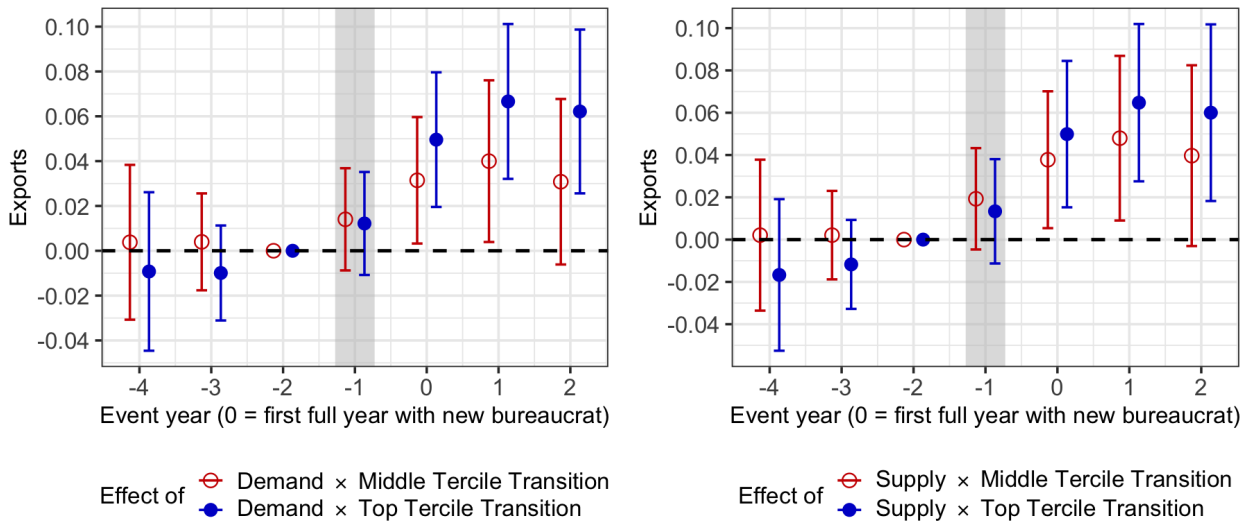
Notes: The figure shows the effect of the incoming and outgoing bureaucrats' fixed effects on exports, around the time that the bureaucrat managing a country office changes. Specifically, we plot $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (5). The outcome variable is the inverse hyperbolic sine of exports 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 that 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. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects. Since the regressors of equation (5) include estimated objects, we report bootstrapped standard errors based on 500 bootstrap iterations, which are found to be more conservative than standard errors clustered by country (unreported).

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 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 1980, we obtain the fixed effect of the bureaucrat in charge of the UK office in 1980 using data from all country-years but the UK. The slope coefficient of the regression of residualized exports on these out-of-sample, i.e. *other country*, fixed effects is 0.52.

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

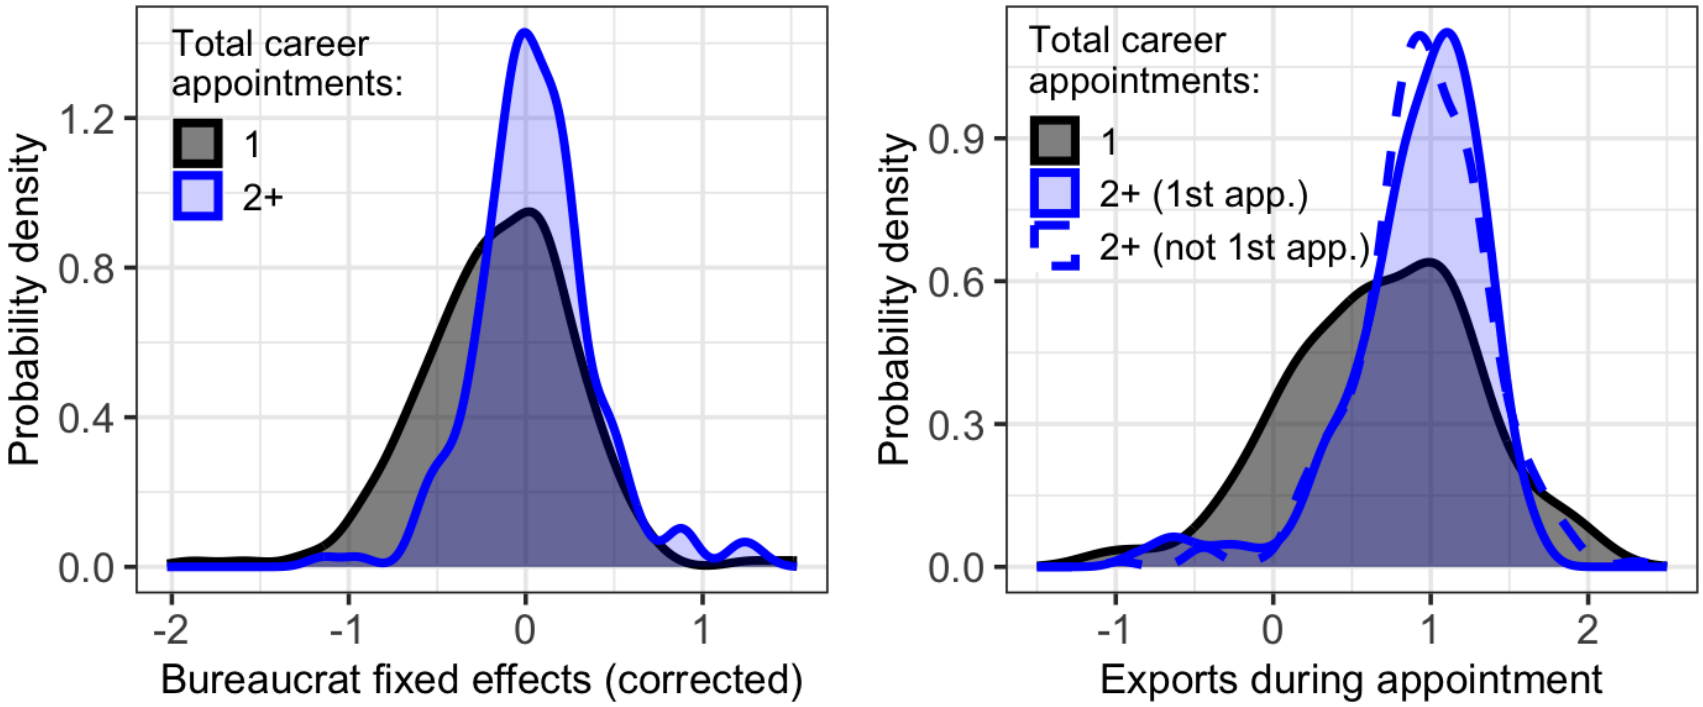


Notes: The figure plots the effect on exports of the change in bureaucrat ability when interacted with two kinds of shocks, estimated from equation (6). The left panel depicts $\hat{\beta}_k^D$, $\hat{\delta}_k^D$, the interaction effects with the destination country's product-specific "demand," proxied by exports of the same product to the same destination by non-Korean countries. 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 that 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 the regressors of equation (6) include estimated objects, we report bootstrapped standard errors based on 500 bootstrap iterations, which are found to be more conservative than standard errors clustered by country (unreported).

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

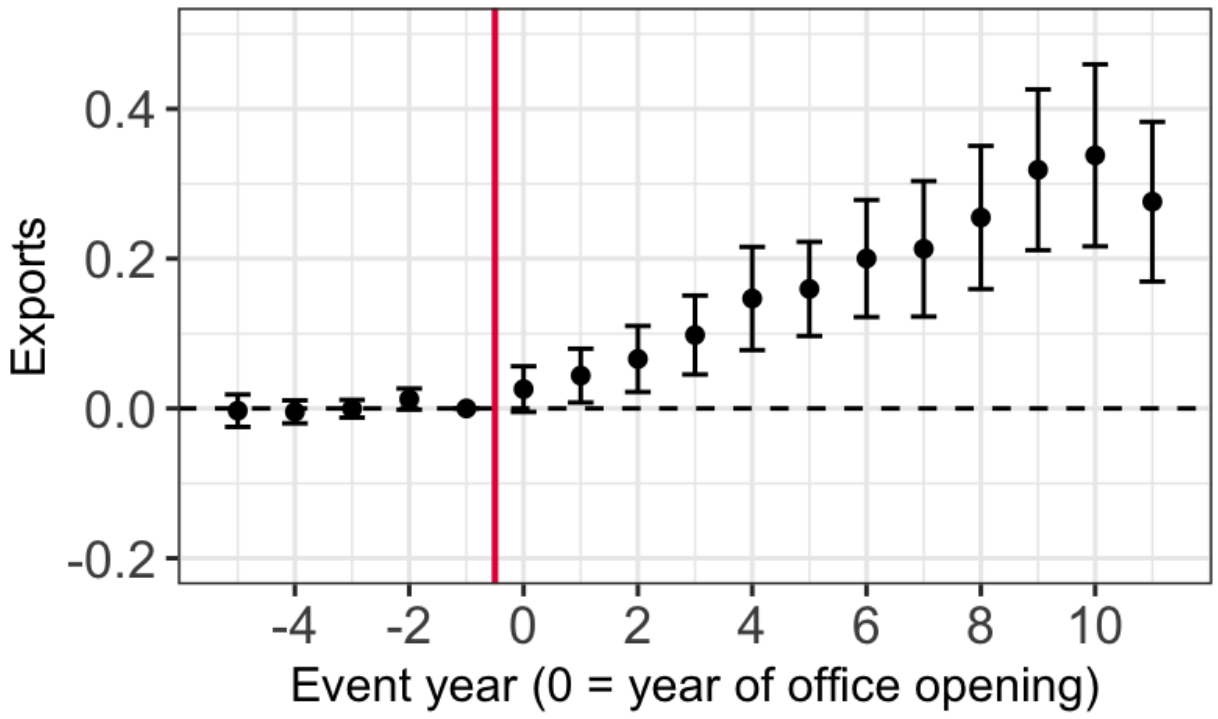
(a) Bureaucrat fixed effects

(b) Residualized exports by appointment



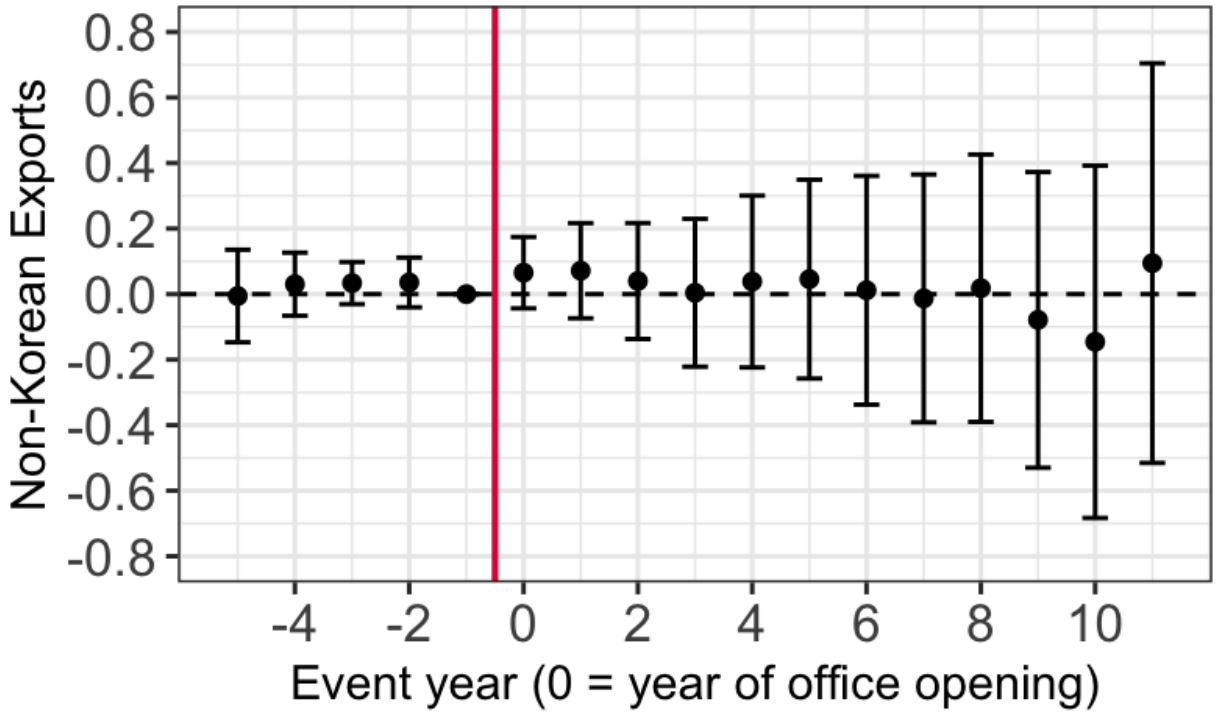
Notes: The figure shows the probability density function of (a) bureaucrat fixed effects (covariance-shrunk) and (b) residualized exports during a bureaucrat's appointment as country office manager. It does so separately for bureaucrats with one appointment (black) and for those with two or more appointments (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 (b) also shows that the distribution of residualized exports does not change much between the first appointment and later appointments for multi-appointment bureaucrats.

Figure 9: Opening an Export Promotion Office Increases Korean Exports to That Country



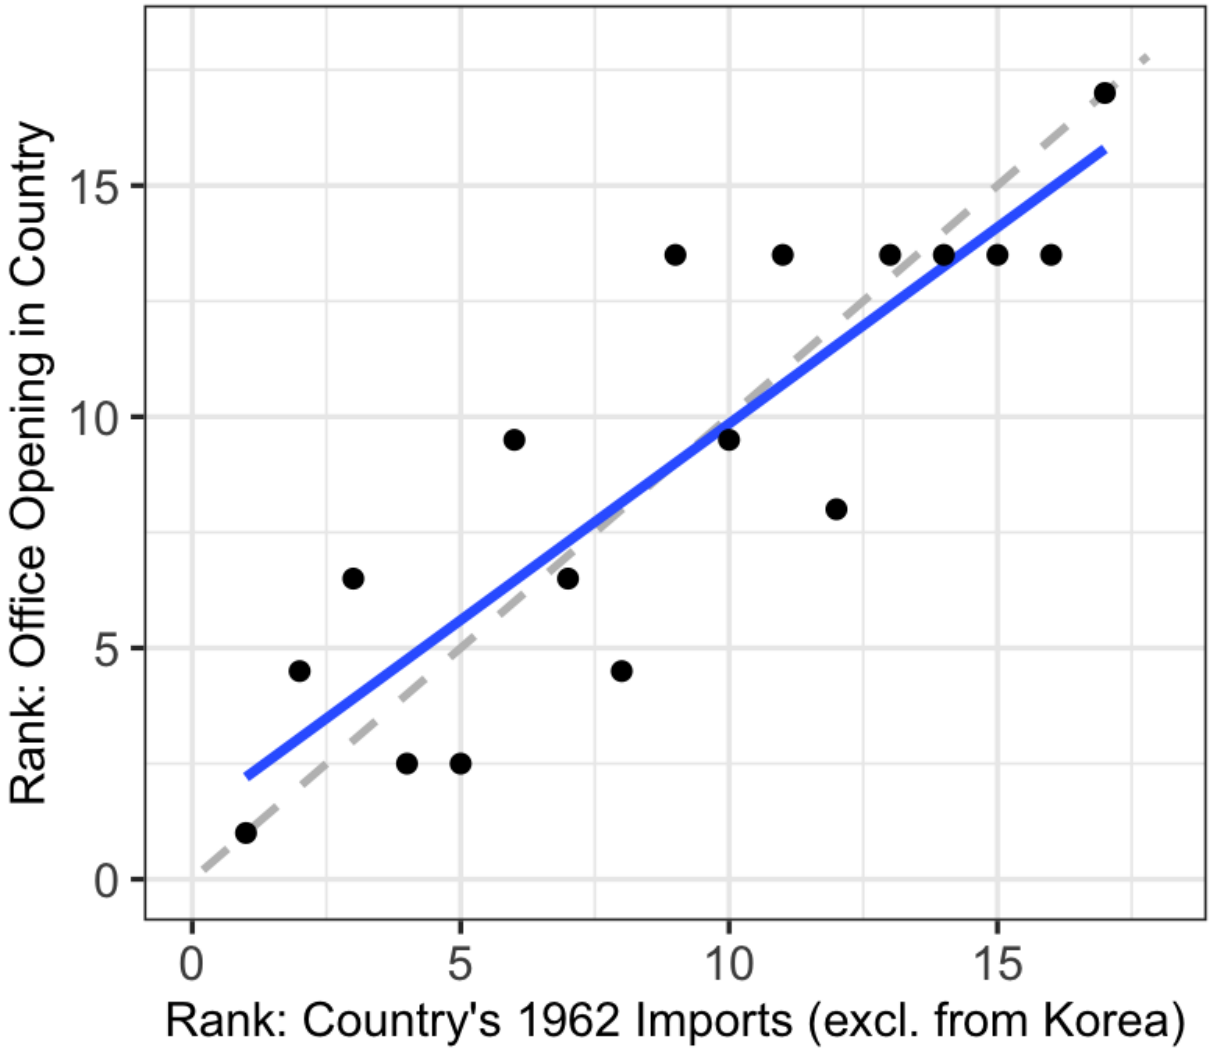
Notes: This figure plots the event-study coefficients ($\hat{\pi}_k$) of the effect of opening export promotion offices on Korean exports, estimated from equation (7) with a never-treated control group. Standard errors are clustered at the country level.

Figure 10: Office Opening Does Not Coincide with
Increases in Import Demand of That Country



Notes: This figure plots the event-study coefficients ($\hat{\pi}_k$) of the effect of opening export promotion offices on *non-Korean* exports, estimated from equation (7) with a never-treated control group. The outcome variable is the inverse hyperbolic sine of non-Korean exports to the country of the office opening, proxying for the country's import demand. Standard errors are clustered at the country level.

Figure 11: 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.

Table 1: Appointment Descriptives

	Full Sample	Connected Set	Leave-One-Out Connected Set
	(1)	(2)	(3)
# Bureaucrats	399	398	380
# Countries/Offices	87	86	75
# Appointments	736	735	676
# Bureaucrats $\geq$ 2 Offices	184	184	180
# Offices $\geq$ 2 Bureaucrats	83	83	75
# Offices $\geq$ 4 Bureaucrats	75	75	72
# Offices $\geq$ 6 Bureaucrats	62	62	61
# Offices $\geq$ 8 Bureaucrats	49	49	49

Notes: The table reports summary statistics 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 this to those countries and managers which form the largest connected set, while column 3 includes only the countries and managers in the largest leave-one-out connected set – i.e. the set of countries and bureaucrats that would remain connected by ommitting connections due to individual appointments. "# Managers" indicates the number of distinct bureaucrats that held a position as office manager. "# Countries/Offices" indicates the number of distinct offices. This is the same as the number of distinct countries. "# Managers $>$ 1 Office" indicates the number of distinct bureaucrats that held a position as manager of at least two offices. "# Offices $>$ x Manager(s") indicates the number of offices with more than x managers over the course of the sample period.

Table 2: Variance decomposition of exports

		Baseline		Robustness check: Bootstrap bias correction		Placebo check: Bureaucrats randomly shuffled 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, Holzheu, Lamadon, Manresa, Mogstad, and Setzler (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 shuffled to countries, preserving the number of appointment spells in the original sample for each bureaucrat. Columns 7-8 replace KOTRA office directors with Korean ambassadors to the same country. The Kline et al. (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 shuffling 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 shuffling produces a different leave-one-out connected set; their similarity to the ones in columns 1-2 imply that the connectivity of the graph is largely maintained in the random shuffling. 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.

Table 3: Counterfactuals

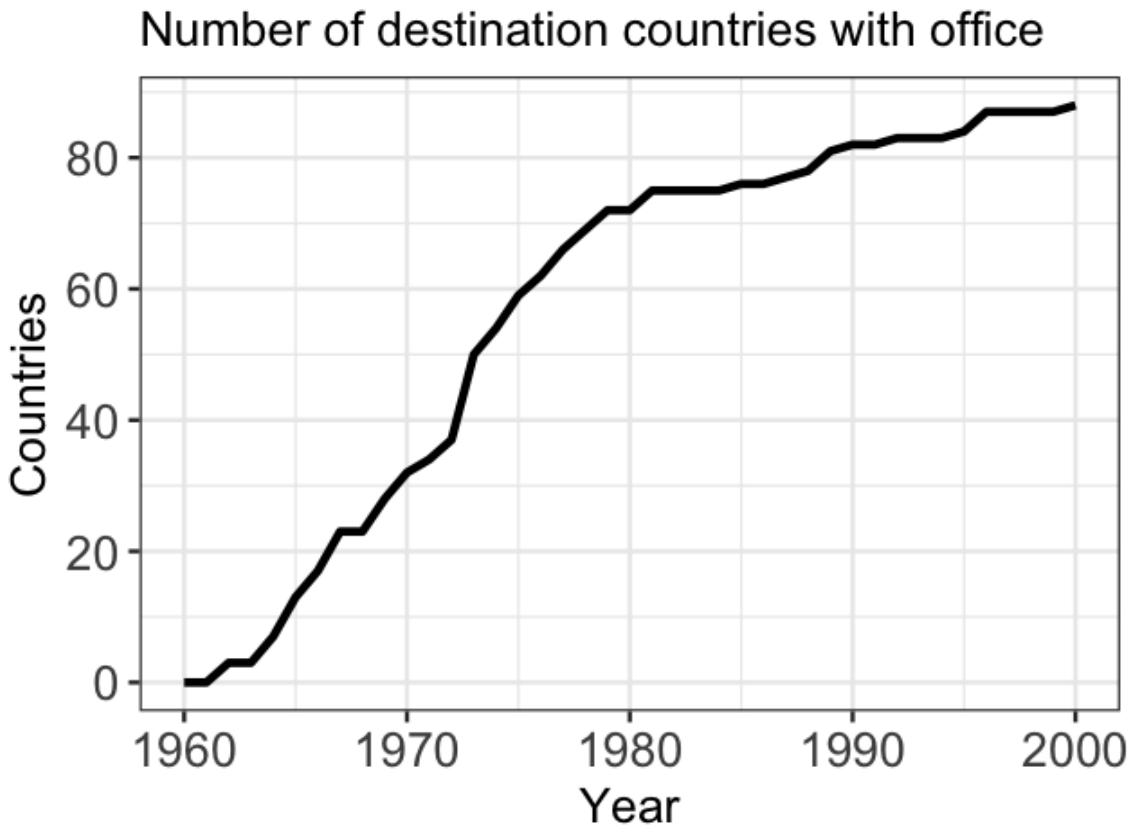
Counterfactual	**Change in Exports** compared to status quo, i.e., screening occurs after 1st appointment			
	Decomposition		Total Effect	Total Effect
	Appointments with single-appointment bureaucrats	Appointments with multi-appointment bureaucrats	All appointments in leave-one-out connected set	All appointments ... using covariance-shrunk bureaucrat effects
(1) No screening occurs after 1st appointment	0.106 (11%)	-0.139 (-13%)	-0.074 (-7.1%)	-0.057 (-5.5%)
(2) Screening occurs before 1st appointment	0.223 (25%)	0	0.066 (6.8%)	0.055 (5.7%)
(3) Replace bottom quartile with median	-	-	-	0.095 (10%)
Number of bureaucrats	200	180	380	380
Number of appointments	200	476	676	676

Notes: The table reports the change in exports (IHS 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).

Online Appendix

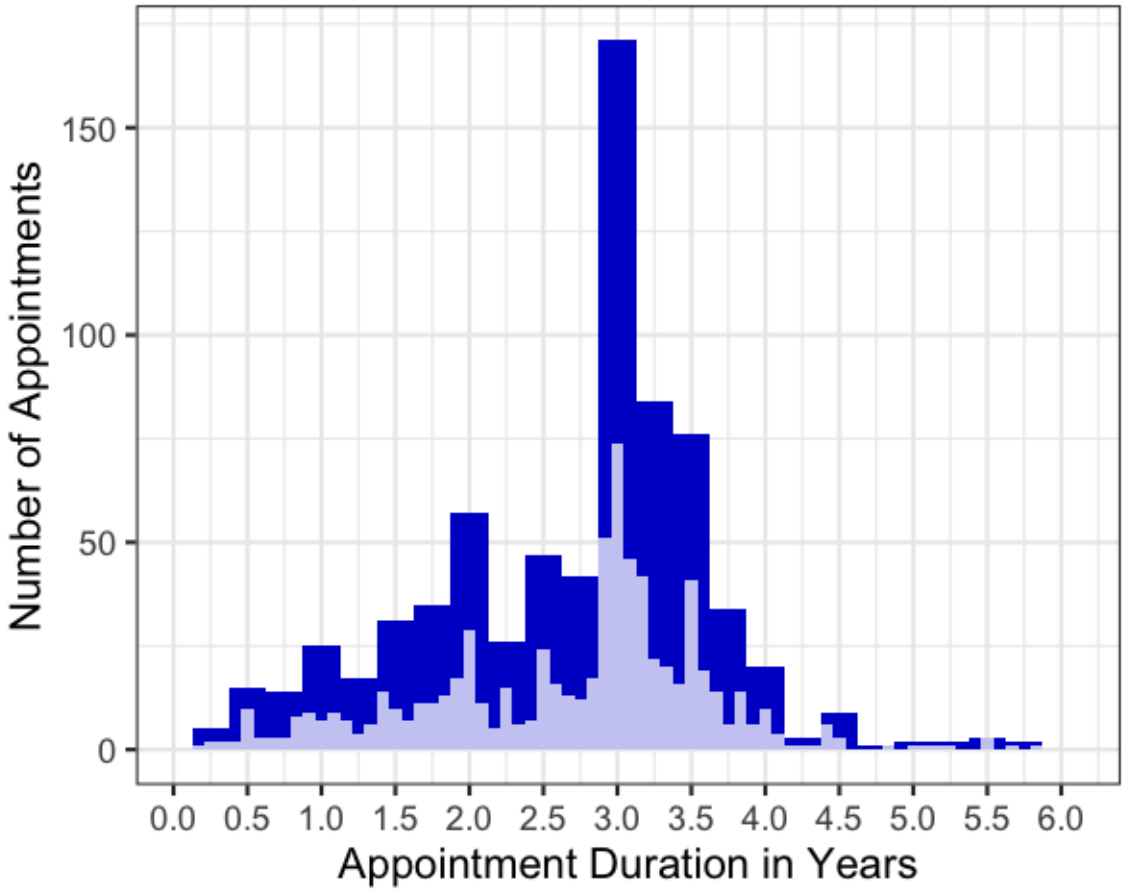
A Appendix Figures

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



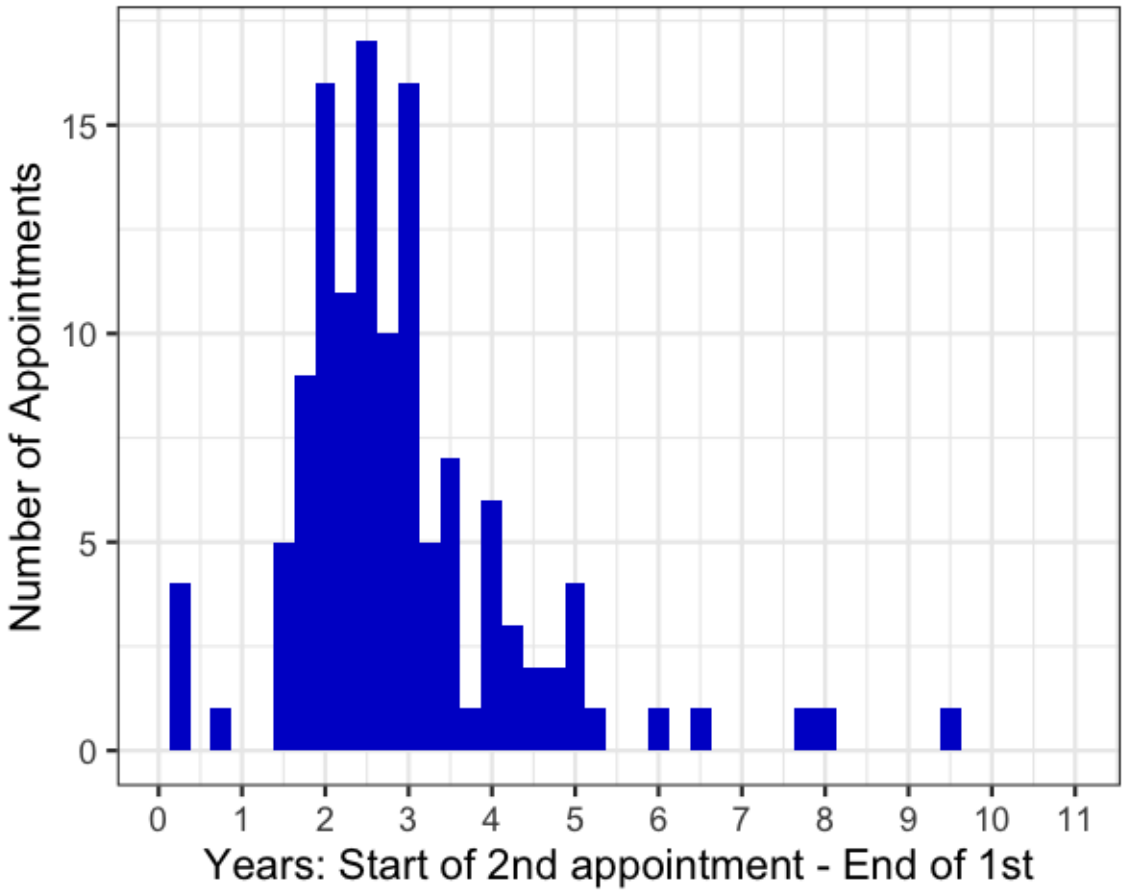
Notes: This figure presents the number of countries in which the Korean export promotion office opened by year t.

Figure A.2: Distribution of appointment durations.
Median and modal duration: 36 months.



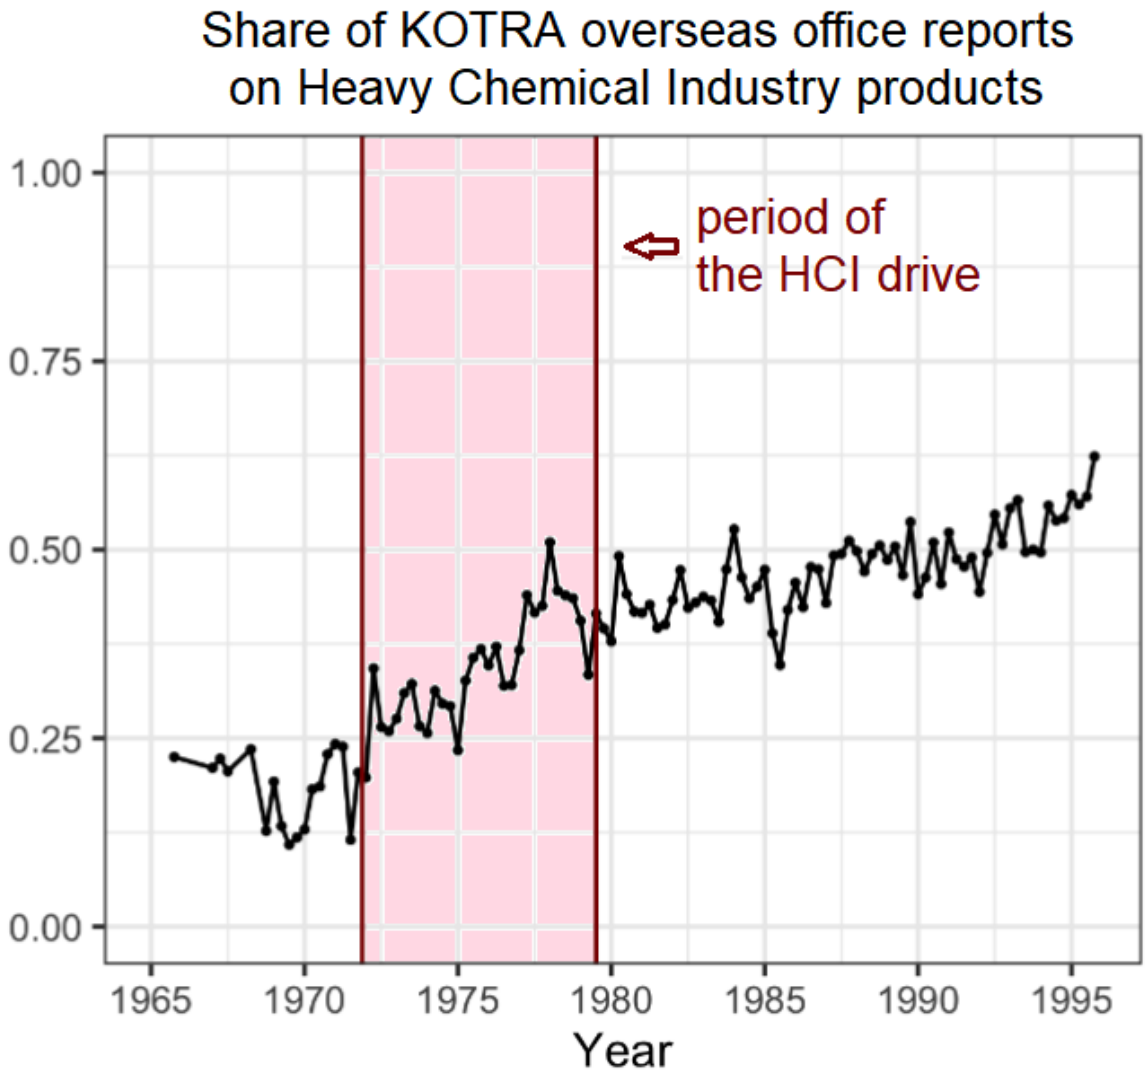
Notes: This figure represents the distribution of appointment durations. The blue bars indicate the number of appointments by quarterly duration whereas the white bars do so for the number of appointments by monthly duration. Hence, as each quarter contains multiple months, the blue bars always (weakly)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

Figure A.3: Distribution of gap lengths.
Median: 29 months. Mode: 30 months.



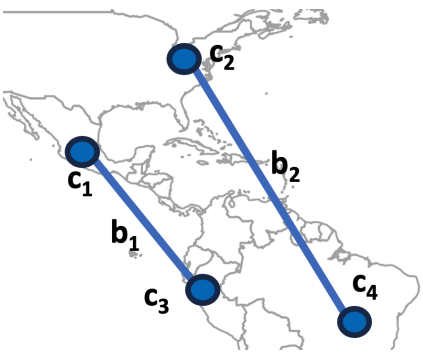
Notes: This figure represents the distribution of the duration of gaps between appointments. The blue bars indicate the number of gaps by quarterly duration.

Figure A.4: Targeting of export promotion activity by product.
Export promotion activity moves in parallel with national industrial policy



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 an HCI product relative to the number of reports that could be linked to any product.

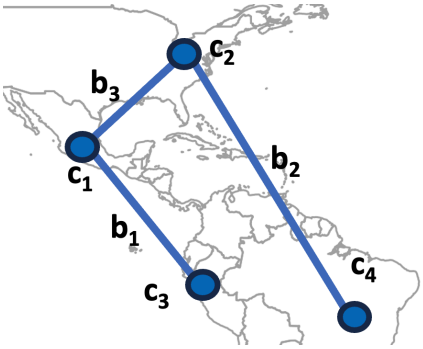
Figure A.5: KOTRA Bureaucrats' Rotation Results in a Single Connected Set



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



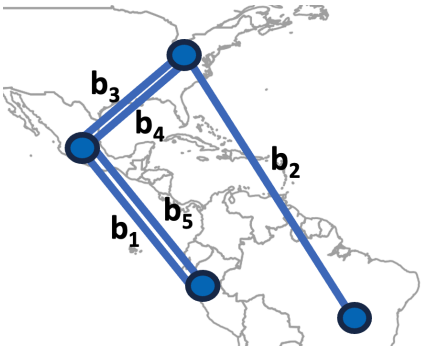
(b) Connections between the UK and other countries due to the bureaucrat appointed to manage the London office in 1981.



(c) Country-bureaucrat graph composed of single connected set.



(d) Connections from UK due to 1981 and 1984 appointments.



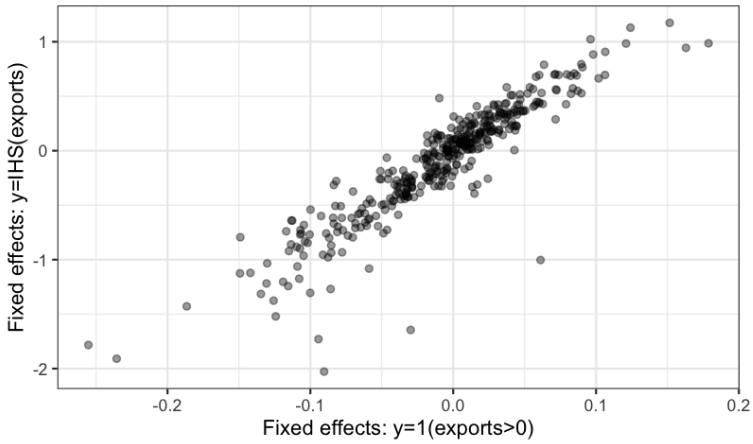
(e) Largest leave-one-out connected set includes c_1, c_2, & c_3, but not c_4.



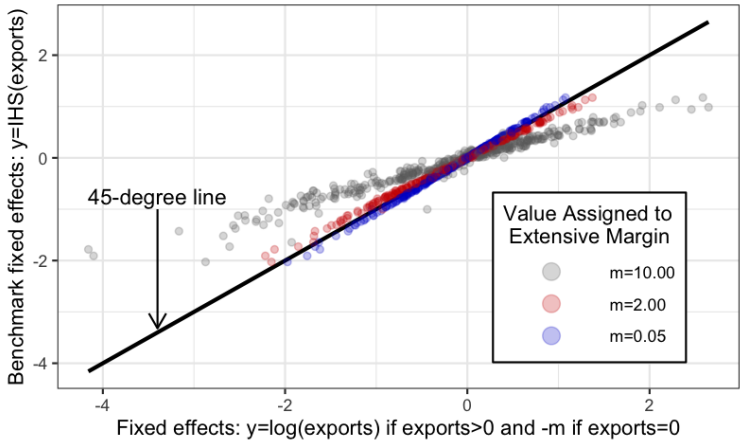
(f) Connections from UK due to 1981, 1984, 1987 appointments.

Notes: This figure highlights how this paper's data fulfills the requirement for the country-bureaucrat graph to form a single connected set. Panels (a), (c), and (e) display a hypothetical country-bureaucrat graphs. The nodes indicate the countries, the edges indicate bureaucrats who are (subsequently) observed as managers of multiple country offices – e.g., b_1 is observed in both Mexico and Peru. b_2 is observed in both Brazil and the U.S. This visualization of the bureaucrat-country graph would be unchanged if there were further appointments of bureaucrats who are only ever appointed to one country. Panels (b), (d), and (f) display the connections between the UK and other countries to the appointment of bureaucrats to manage the London office in 1981, 1984, and 1987.

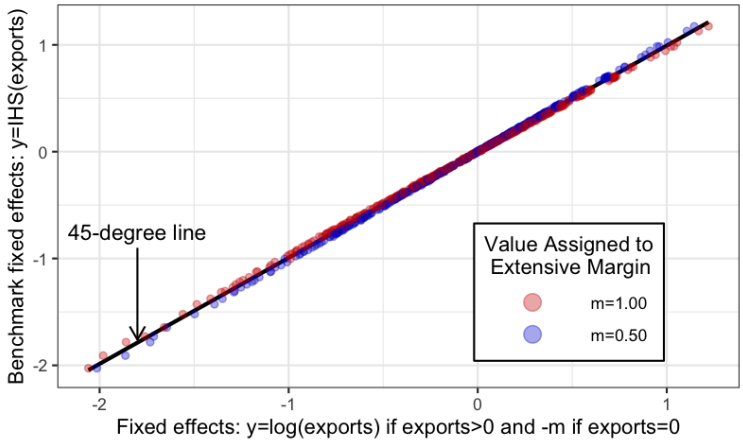
Figure A.6: Bureaucrat fixed effects under alternative transformations of exports.



(a) Compare to fixed effects based on extensive margin.



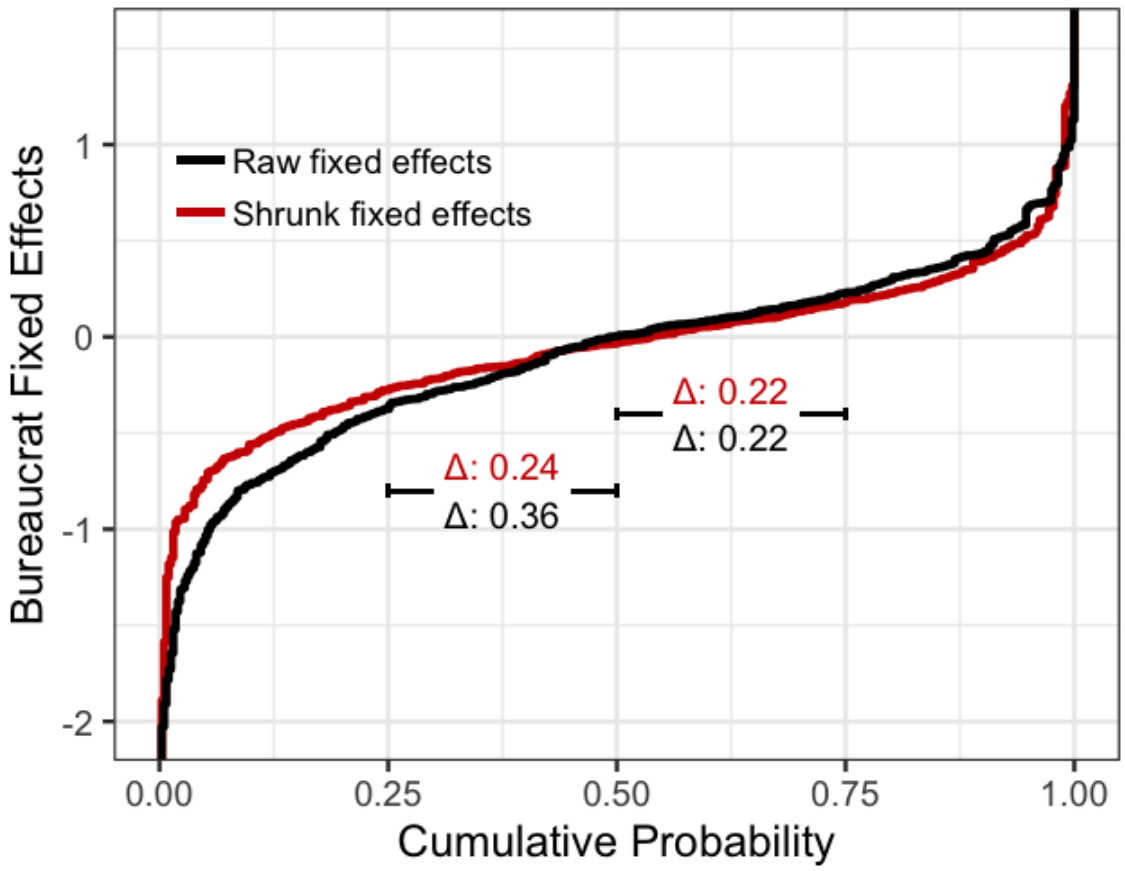
(b) Compare to FE based on explicit (extreme) values assigned to extensive margin.



(c) Compare to FE based on explicit (moderate) values assigned to ext. margin.

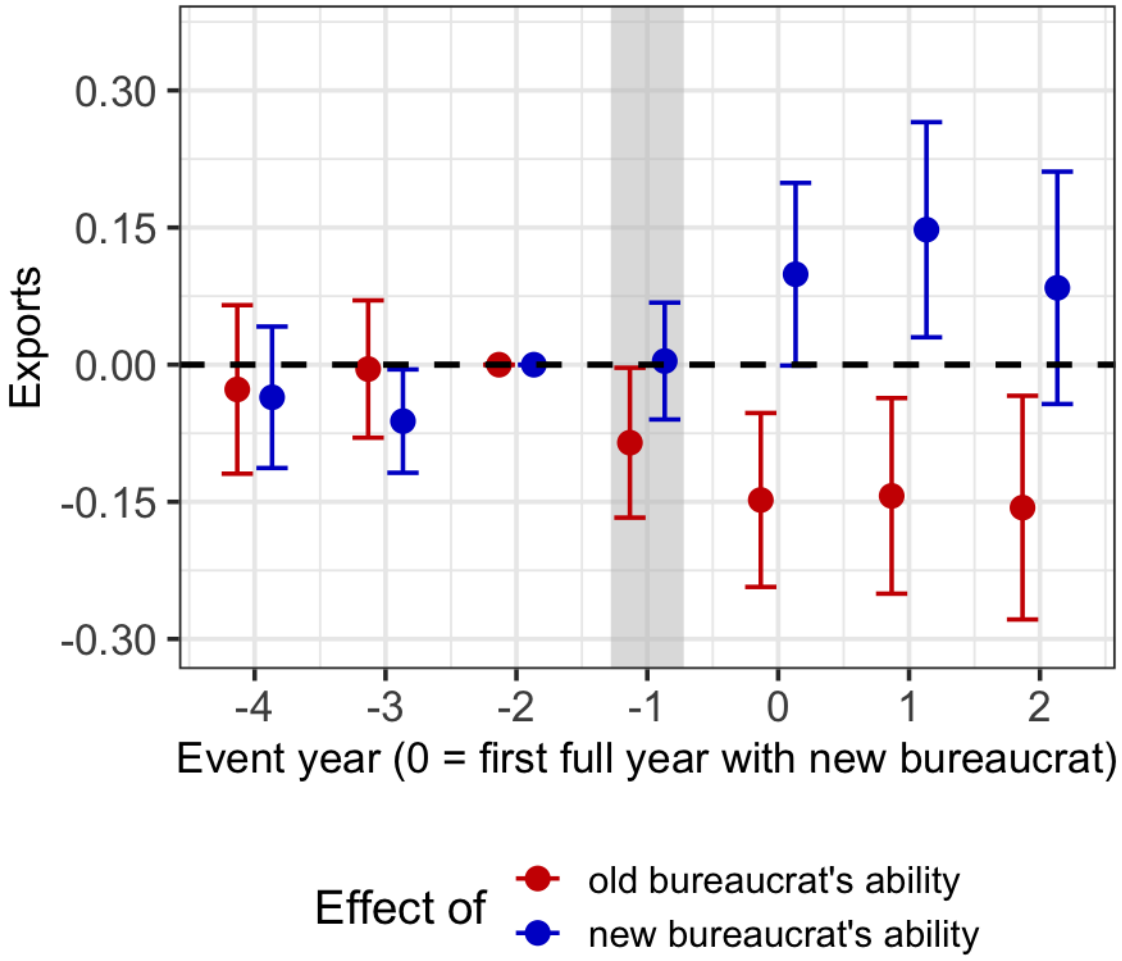
Notes: This figure displays the relationship between the benchmark bureaucrat fixed effects and those estimated under alternative transformations of exports in equation (1).

Figure A.7: CDF of Raw vs Shrunk Fixed Effects of Managers of Country Offices



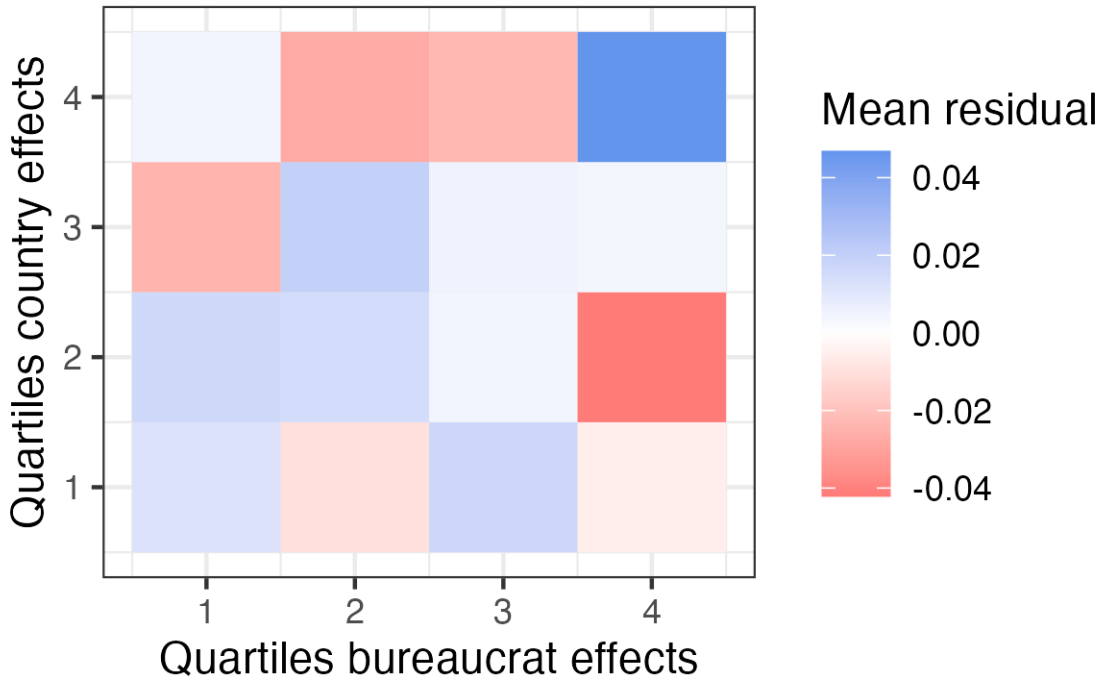
Notes: The figure shows the cumulative distribution function of bureaucrat fixed effects based on equation (3), in addition to covariance-shrunk bureaucrat fixed effects. The first annotation indicates that moving from the 25th to the 50th percentile would imply an increase in the inverse hyperbolic sine of exports to a country of 0.24 based on the corrected fixed effects compared to 0.36 under the raw fixed effects. The second annotation indicates that moving from the 50th to the 75th percentile would imply an increase in the inverse hyperbolic sine of 0.22 based on both the corrected and the raw fixed effects.

Figure A.8: Event study estimates: Out-of-sample bureaucrat fixed effects



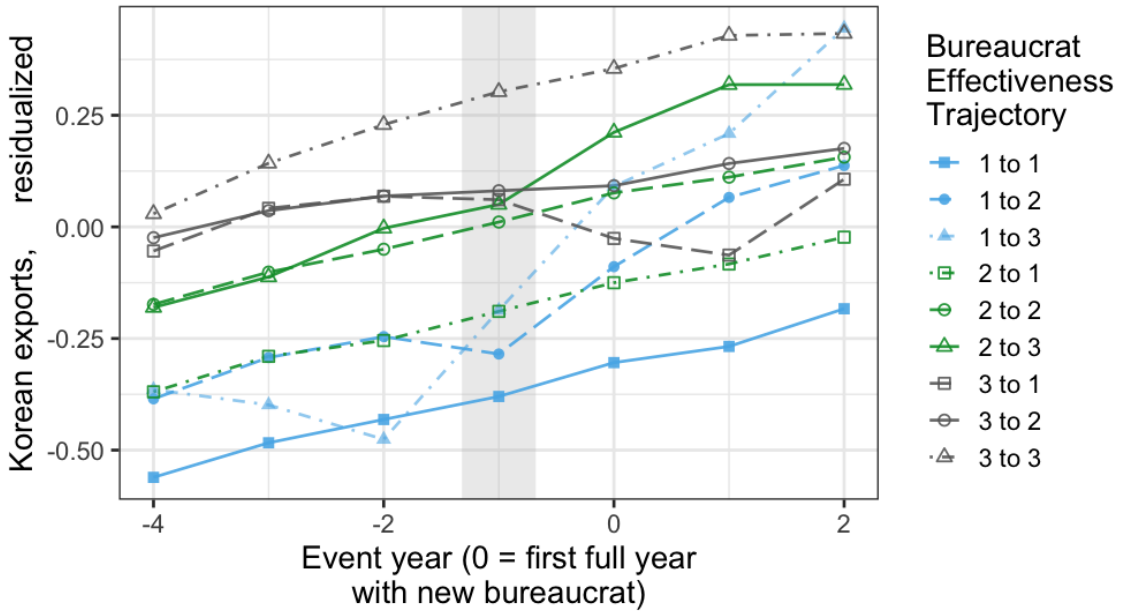
Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects, estimated out of sample, on exports around the time that the manager of a country office changes. These estimates are $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (5). As out-of-sample fixed effects are not available for every bureaucrat, to maximize power, we report coefficients from two different models. First, we estimate equation (5) using *out-of-sample* estimates for the outgoing bureaucrat and *in-sample* estimates for the incoming bureaucrat. Second, we estimate equation (5) using *in-sample* estimates for the outgoing bureaucrat and *out-of-sample* estimates for the incoming bureaucrat. For each model, we only report the out-of-sample coefficients, as these are the ones of interest. For each model, the in-sample coefficients are almost symmetric to the out-of-sample ones. The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat manages the country office. Year -2 is the last full year that the old bureaucrat managed the office. The y axis measures the effect of bureaucrat effectiveness on exports. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure A.9: Residuals by estimated bureaucrat and organization effects. Absence of clear pattern which would point to misspecification.



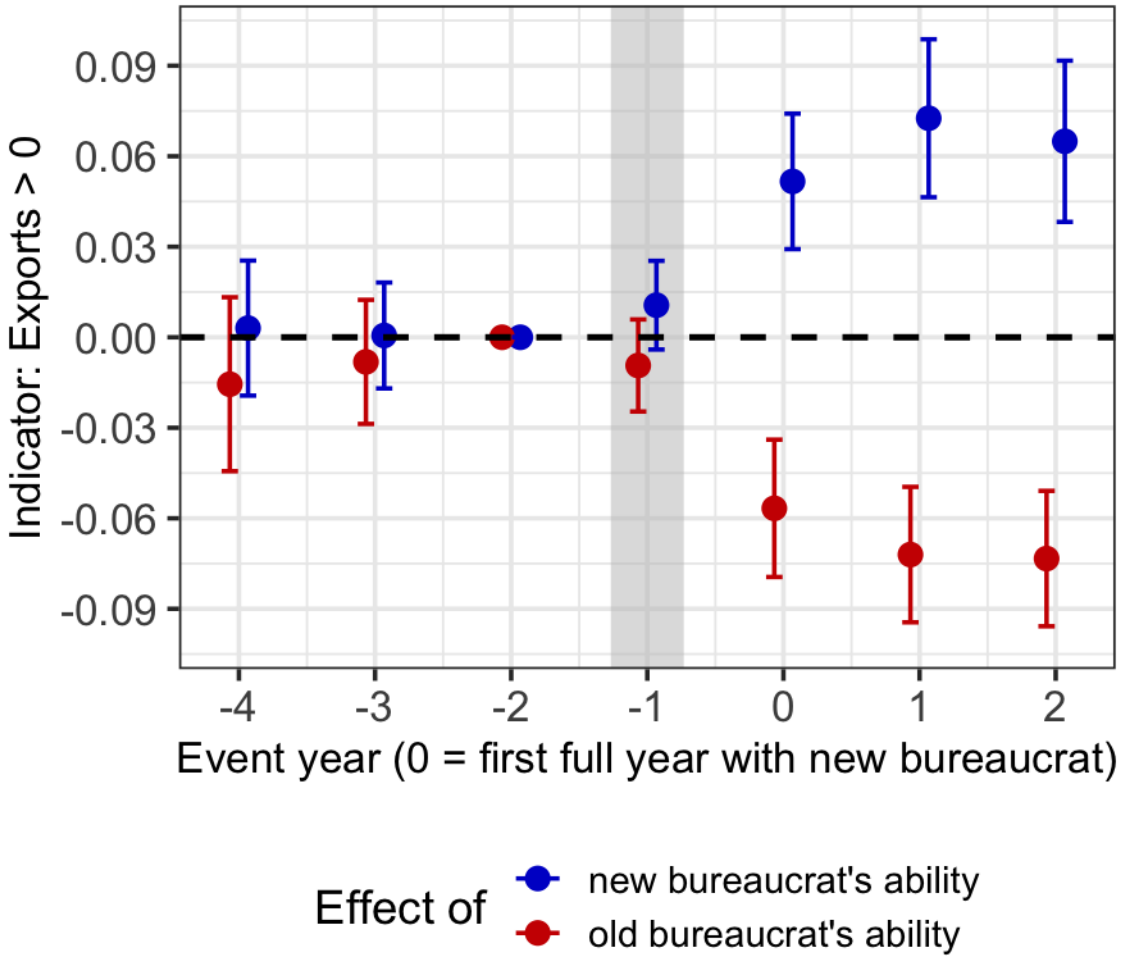
Notes: This figure shows mean residuals from equation 2 with cells defined by quartiles of estimated bureaucrat effect, interacted with quartiles of estimated country effect.

Figure A.10: Mean residualized exports around switches between bureaucrats. Effects consistent across terciles of new and old bureaucrats.



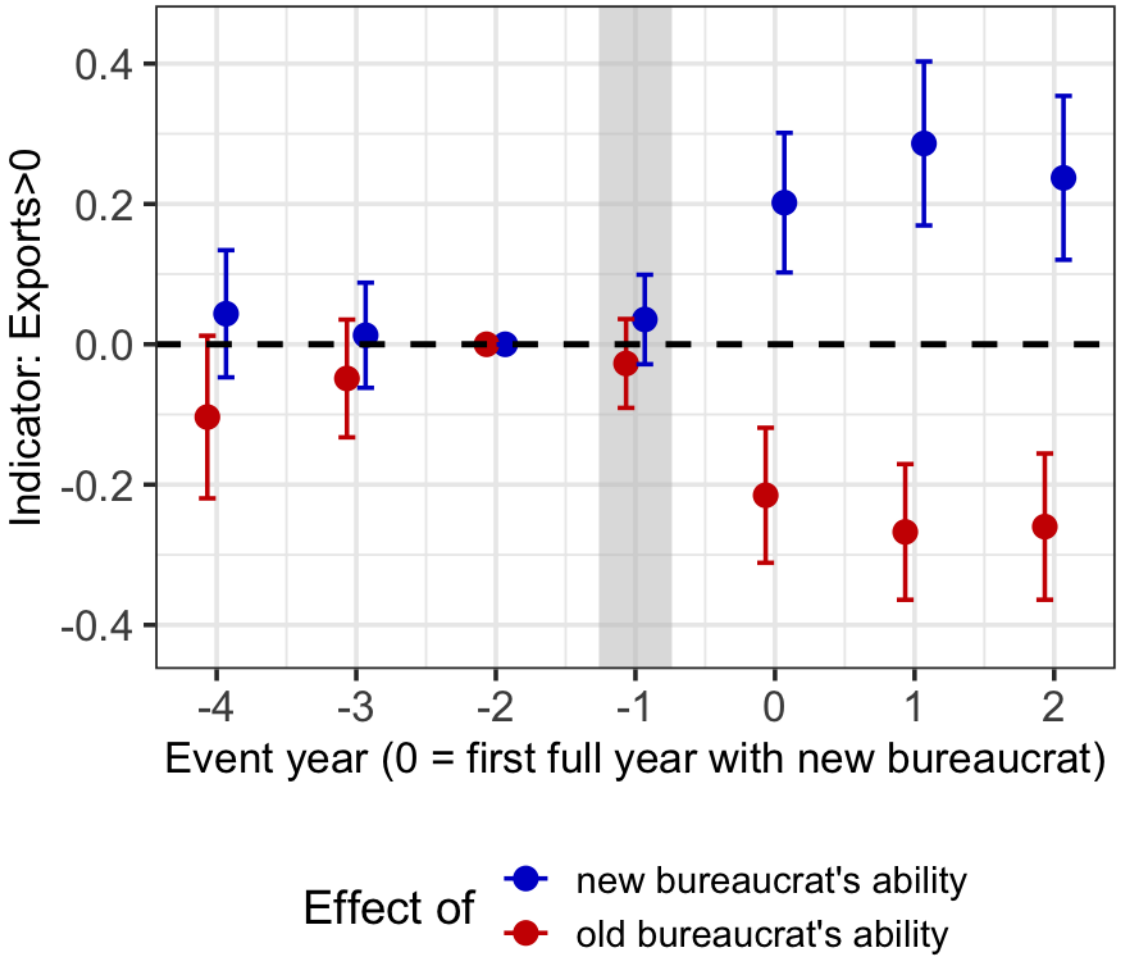
Notes: The figure shows time trends in exports around the time that the manager of a country office changes. The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat manages the country office. Year -2 is the last full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

Figure A.11: Event study – the extensive margin response to switches between bureaucrats



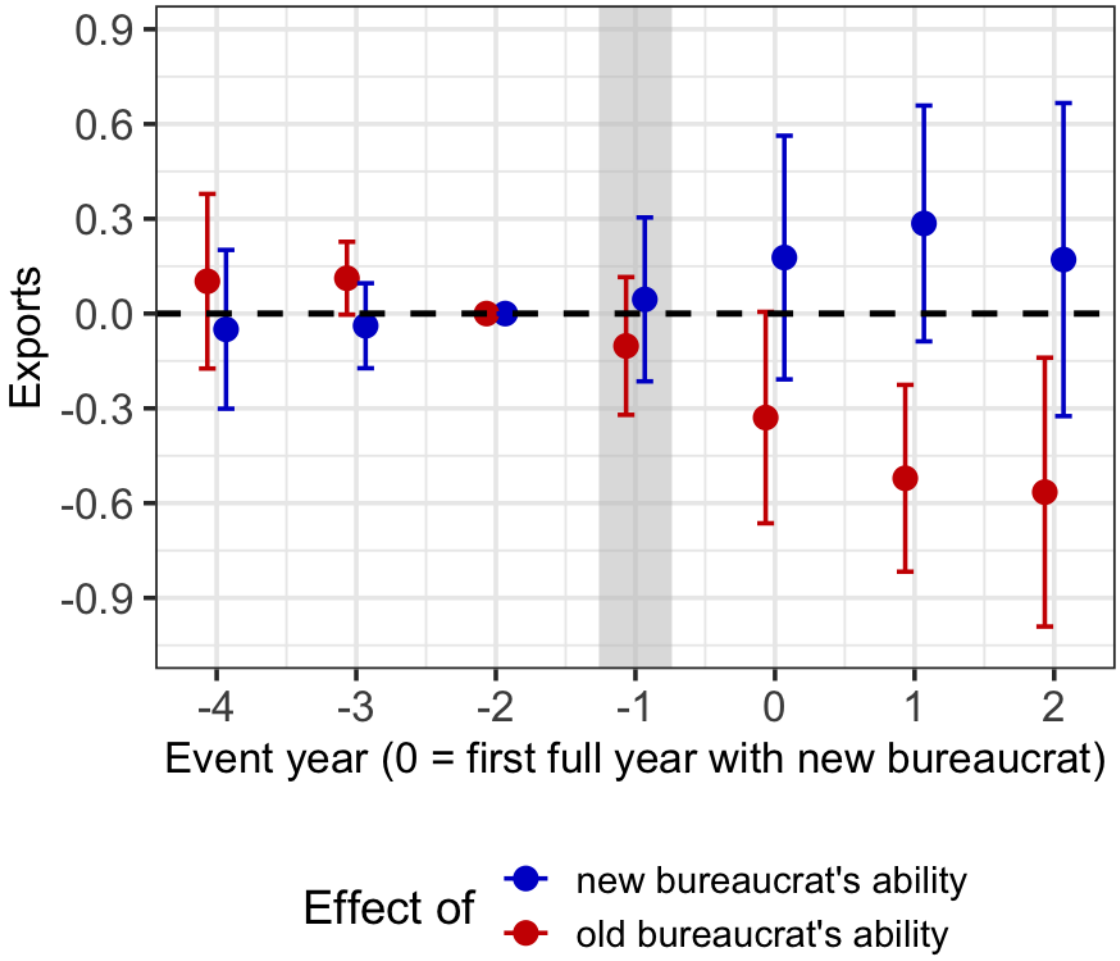
Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects on the likelihood of positive exports in a given product around the time that the manager of a country office changes. These estimates are $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (5). Observations are included for a given event-horizon if South Korea exports this product to any country for all years in the event horizon. The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat manages the country office. Year -2 is the last full year that the old bureaucrat managed the office. The y axis measures the effect of bureaucrat effectiveness on exports. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure A.12: Large extensive margin response to bureaucrat effects for products with any change in extensive margin during event horizon



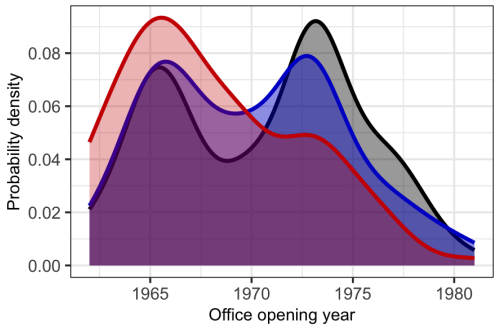
Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects on the likelihood of positive exports in a given product around the time that the manager of a country office changes. These estimates are $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (5). Observations are included for a given event-horizon if South Korea exports this product to this country in one year during the event horizon but not all years in the event horizon. The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat manages the country office. Year -2 is the last full year that the old bureaucrat managed the office. The y axis measures the effect of bureaucrat effectiveness on exports. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

Figure A.13: Event study – the intensive margin response to switches between bureaucrats

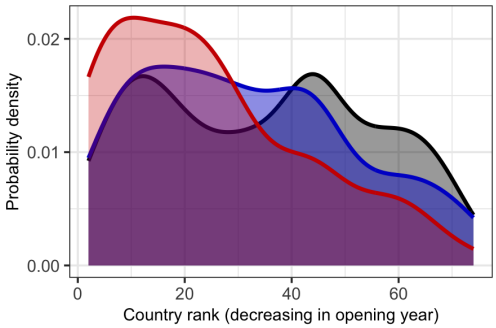


Notes: The figure shows the estimated effect of the change in bureaucrat fixed effects on exports around the time that the manager of a country office changes. This only includes the intensive margin effect as observations are included for a given event-horizon if South Korea exports this product to this country in all years during the event horizon. These estimates are $\hat{\beta}_k$ and $\hat{\delta}_k$ obtained from estimating equation (5). The horizontal axis indexes years in which bureaucrats work in a particular country. Year 0 is the first full year that the new bureaucrat manages the country office. Year -2 is the last full year that the old bureaucrat managed the office. Bureaucrats effectiveness are fixed effects obtained after residualizing exports by product-country and product-year fixed effects.

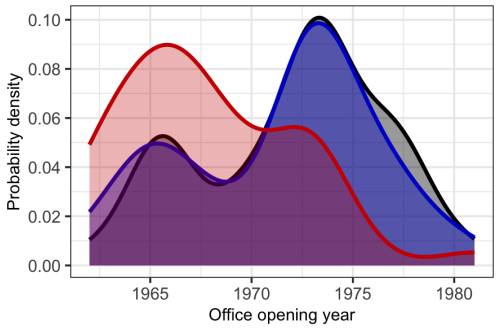
Figure A.14: As bureaucrats' careers progress they are appointed to more offices that opened earlier (proxying for importance).



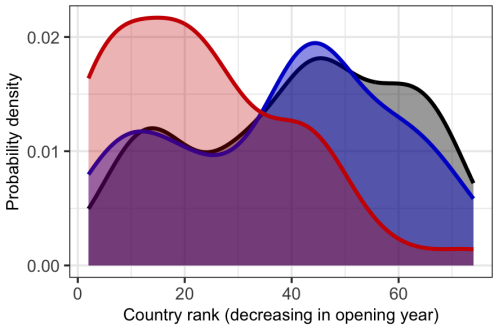
(a) PDF of bureaucrats' 1st, 2nd, 3rd appointment by opening year of the office.



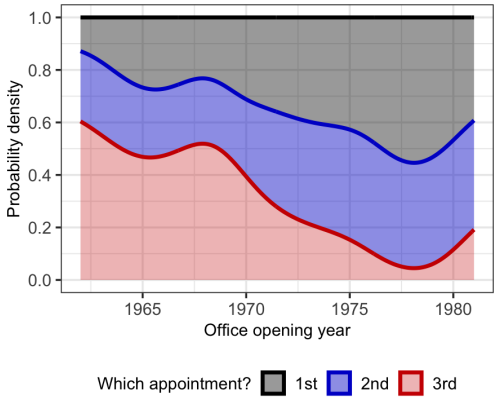
(b) PDF of bureaucrats' 1st, 2nd, 3rd appointment by opening rank of the office.



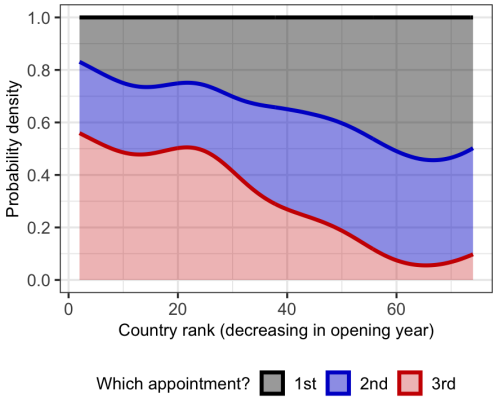
(c) Same as (a), but restricted to bureaucrats whose first appointment started no earlier than 1981.



(d) Same as (b), but restricted to bureaucrats whose first appointment started no earlier than 1981.

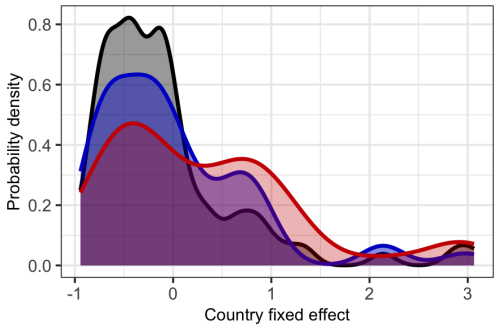


(e) Same as (c), but showing shares of appointments.

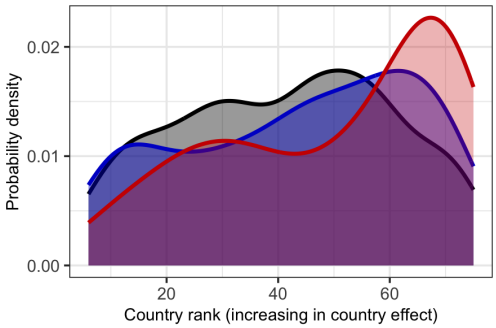


(f) Same as (d), but showing shares of appointments.

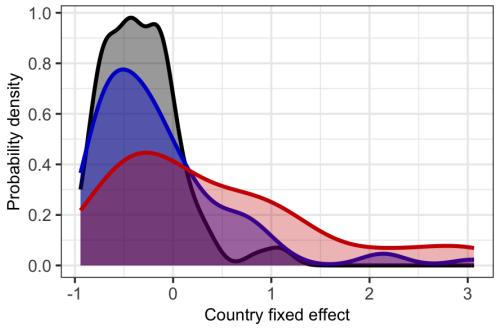
Figure A.15: As bureaucrats' careers progress they are appointed to to countries with higher fixed effects.



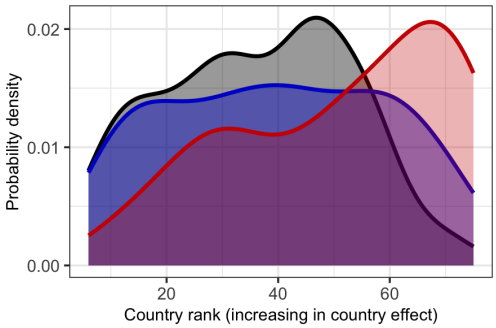
(a) PDF of bureaucrats' 1st, 2nd, 3rd appointment by country effect.



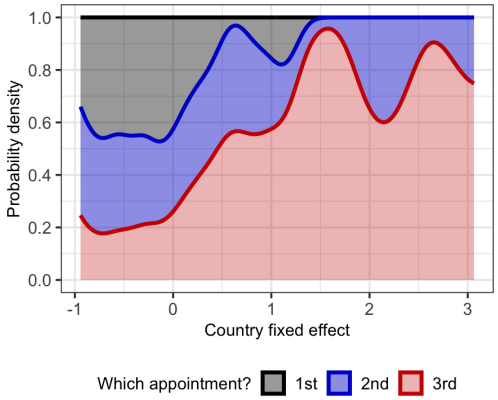
(b) PDF of bureaucrats' 1st, 2nd, 3rd appointment by *rank* of country effect.



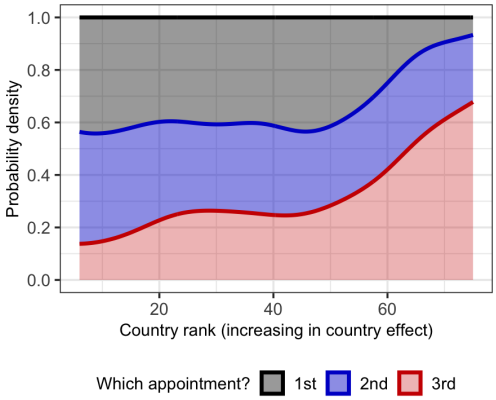
(c) Same as (a), but restricted to bureaucrats whose first appointment started no earlier than 1981.



(d) Same as (b), but restricted to bureaucrats whose first appointment started no earlier than 1981.

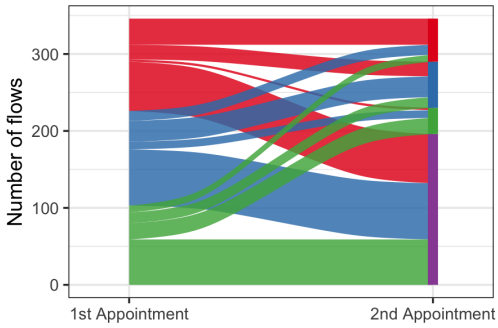


(e) Same as (c), but showing shares of appointments.

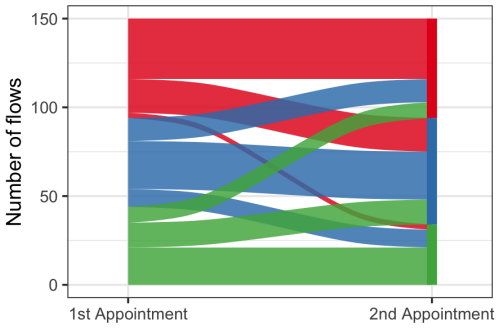


(f) Same as (d), but showing shares of appointments.

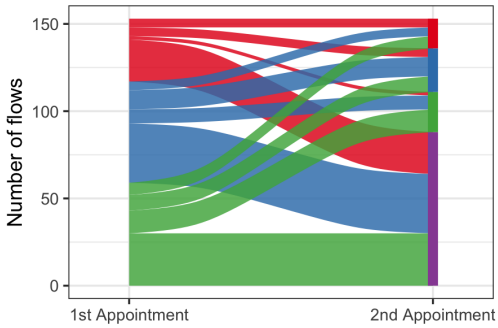
Figure A.16: Bureaucrat flows (by appointment & opening year)



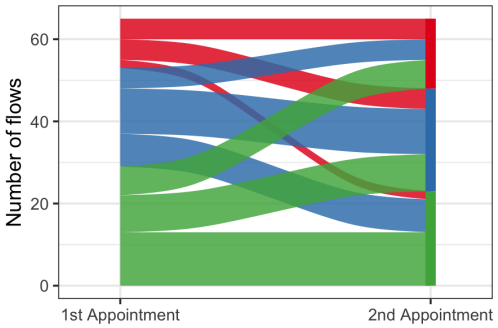
(a) All flows from 1st appointment.



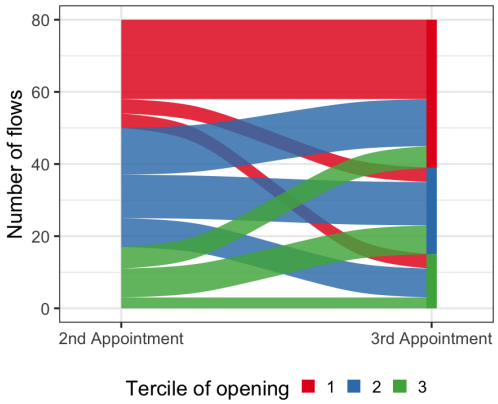
(b) Flows from 1st app. (no exit).



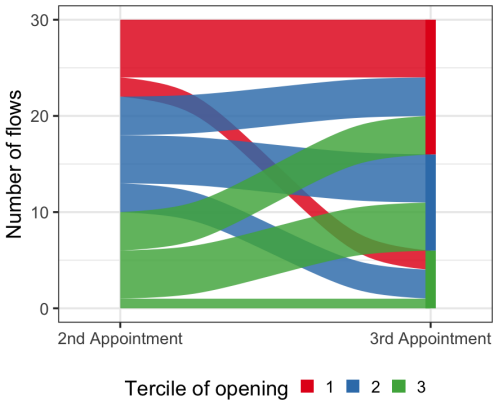
(c) Flows from 1st app. & bureaucrats without appointment before 1981.



(d) Flows (except exits) from 1st app. & bureaucrats without app. until 1981.



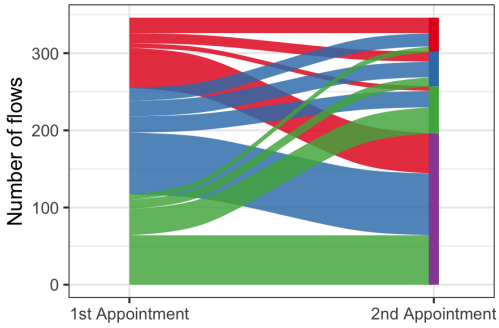
(e) Flows (except exits) from 2nd app.



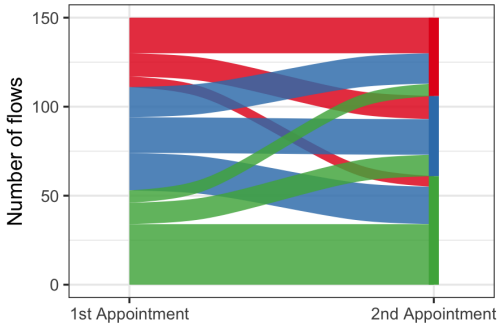
(f) Flows (except exits) from 2nd app. & bureaucrats without app. until 1981.

Notes: This figure shows the flow of bureaucrats from their 1st to 2nd appointment (2nd to 3rd in panels (e) and (f)). We split the offices into groups based on whether the office opening year is in the 1st, 2nd, or 3rd tercile of the original roll-out of offices. We interpret being in an earlier tercile as a revealed preference measure of the importance that KOTRA attributes to an office.

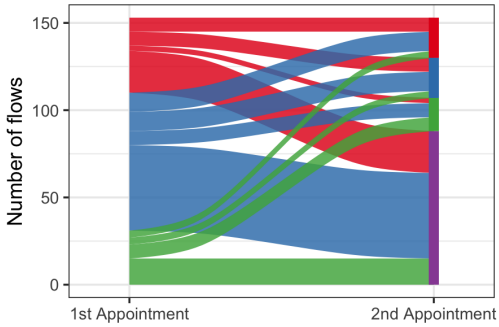
Figure A.17: Bureaucrat flows (by appointment & country effect)



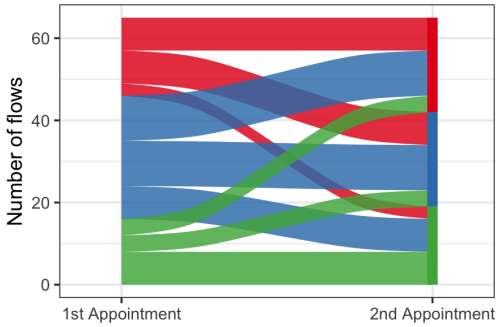
(a) All flows from 1st appointment.



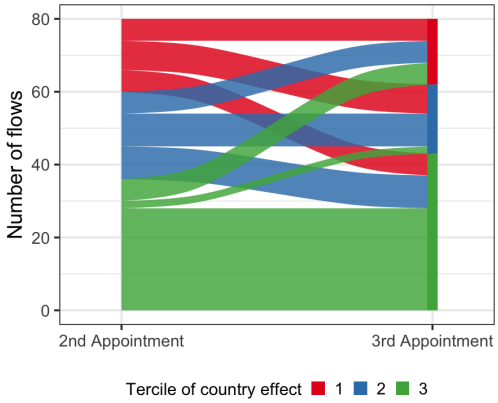
(b) Flows from 1st app. (no exit).



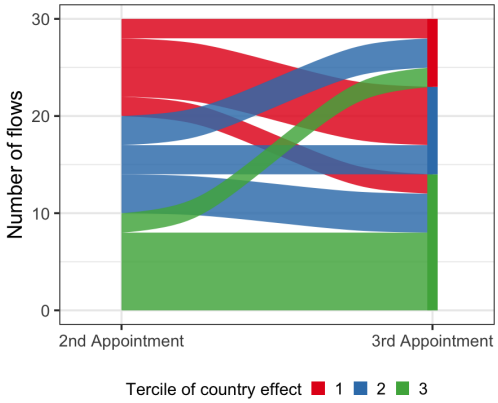
(c) Flows from 1st app. & bureaucrats without appointment before 1981.



(d) Flows (exc. exits) from 1st app. & bureaucrats without app. until 1981.



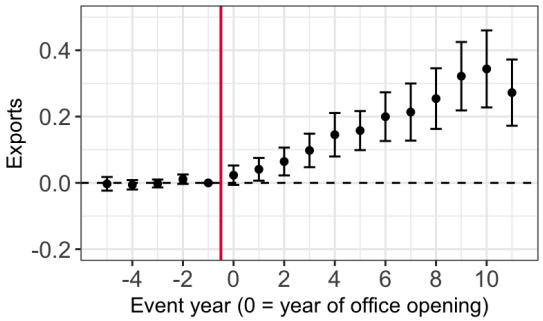
(e) Flows (except exits) from 2nd app.



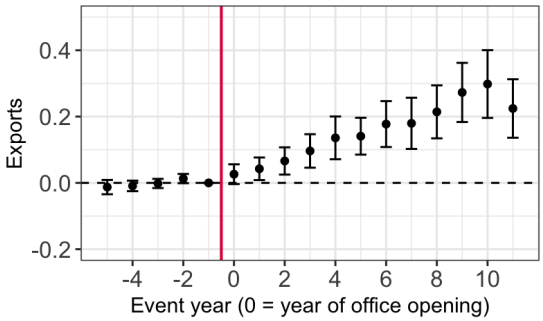
(f) Same as (e) for bureaucrats without app. until 1981.

Notes: This figure shows the flow of bureaucrats from their 1st to 2nd appointment (2nd to 3rd in panels (e) and (f)). We split the offices into terciles based on the country fixed effects. The 3rd tercile consists of the most important countries according to this metric.

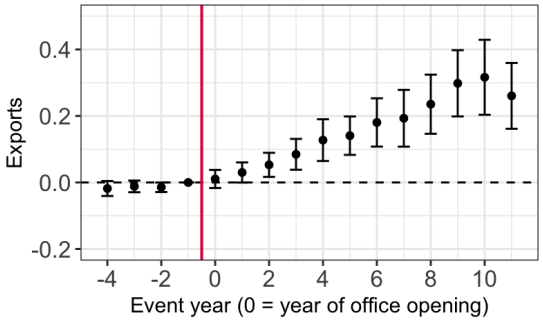
Figure A.18: Robustness: controls, sample, placebo



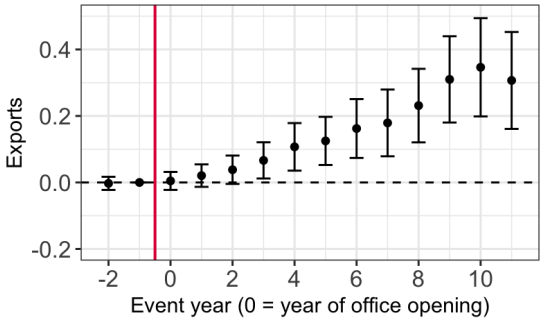
(a) Controlling for non-Korean exports.



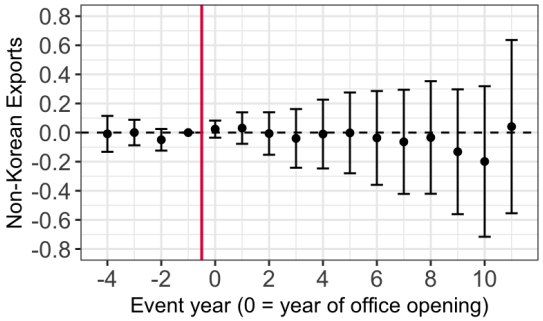
(b) Controlling for non-Korean exports × year dummies.



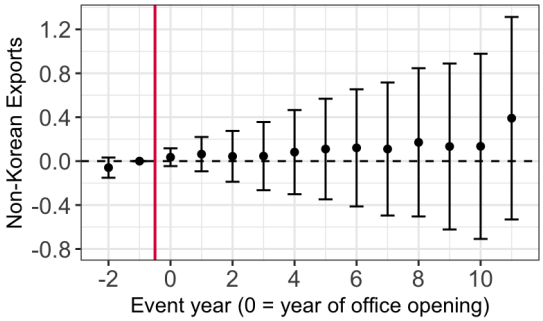
(c) Korean exports as outcome. Openings from 1966. Never-treated as control group.



(d) Korean exports as outcome. Openings from 1964. Never-treated as control group.



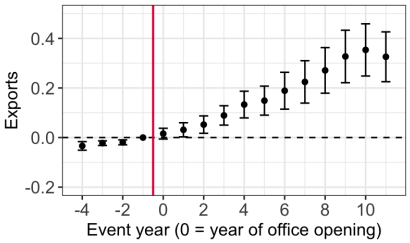
(e) Non-Korean exports as outcome. Openings from 1966. Never-treated as control group.



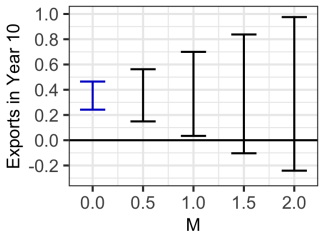
(f) Non-Korean exports as outcome. Openings from 1964. Never-treated as control group.

Notes: For panels (a)-(d), the outcome variable is the inverse hyperbolic sine of Korean exports to the country-year in question. For panels (e) and (f), the outcome is given by the inverse hyperbolic sine of non-Korean exports to the same country-year. An observation is at the product-country-year. Point estimates and standard errors are obtained from estimating equation (7), relying on a never-treated control group. Standard errors clustered at the country-level are reported around each point estimate. A product is included for all the years in which Korea exported it to any country. Product refers to 4-digit SITC Rev. 2 codes.

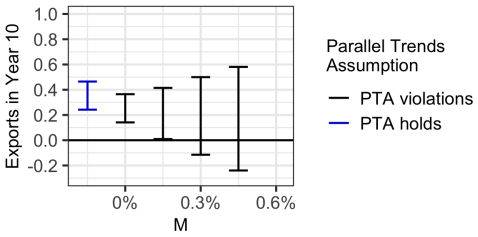
Figure A.19: Robustness: opening with not-yet-treated control



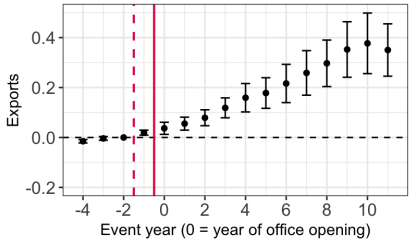
(a.i) CSA estimate, unconditional PTA. 0 periods of anticipation.



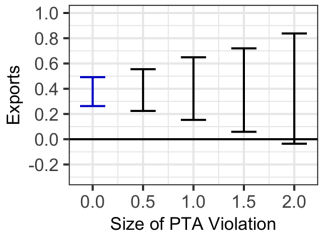
(a.ii) Sensitivity to PTA violation relative to largest pre-treatment violation of PT.



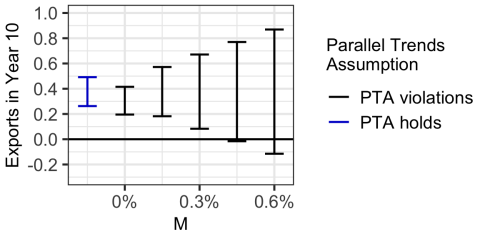
(a.iii) Sensitivity to PTA violations only bounding the extent to which the slope may change across consecutive periods.



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



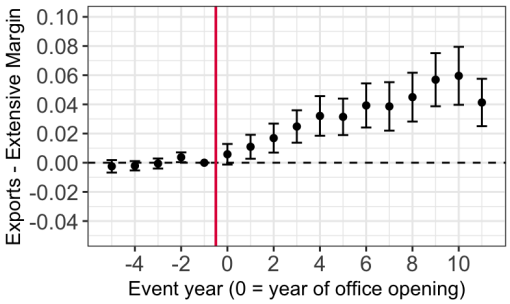
(b.ii) Sensitivity to PTA violation relative to largest pre-treatment violation of PT.



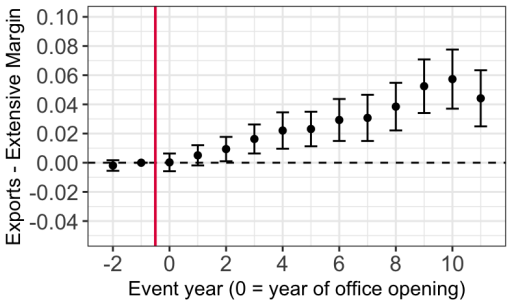
(b.iii) Sensitivity to PTA violations only bounding the extent to which the slope may change across consecutive periods.

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

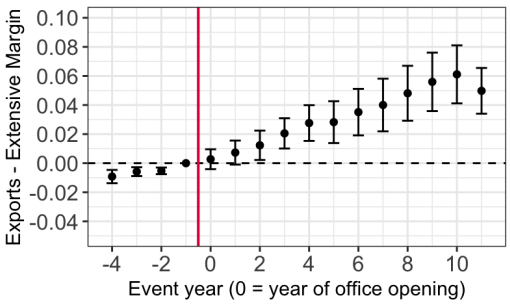
Figure A.20: Extensive margin effect of office opening



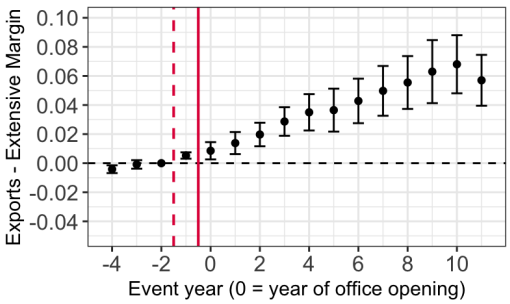
(a) Never-treated as control group. Openings from 1967.



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



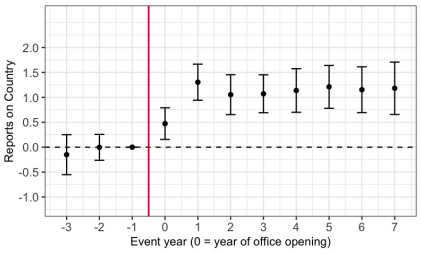
(c) "Not-yet" control. 0 period anticipation.



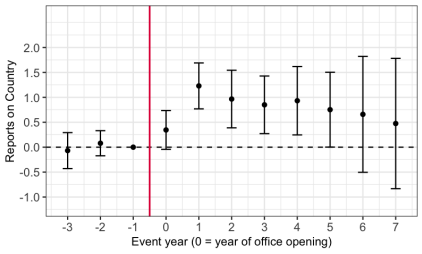
(d) "Not-yet" control. 1 period anticipation.

Notes: In each panel, the outcome variable is a dummy indicating whether South Korea had positive exports in a particular product-country-year – each panel hence corresponds to a linear probability model. An observation is at the product-country-year. For panels (a)-(b), point estimates and standard errors are obtained from estimating equation (7, relying on a never-treated control group. Standard errors clustered at the country-level are reported around each point estimate. In panels (c)-(d), point estimates give the aggregation of treatment-group-specific estimates of the average treatment effect (ATT) using a not-yet-treated control group and Callaway and Sant'Anna (2021) estimator for Difference-in-Differences settings with staggered roll-out using the doubly-robust estimators form Sant'Anna and Zhao (2020). Bootstrapped standard errors are obtained clustering at the level of the destination country. A product is included for all the years in which Korea exported it to any country. Product refers to 4-digit SITC Rev. 2 codes.

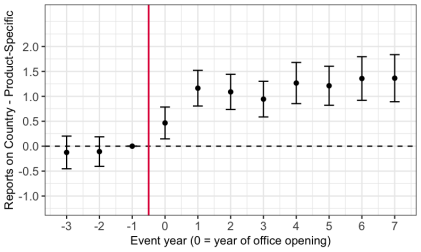
Figure A.21: Event-study estimates of office opening on KOTRA activity



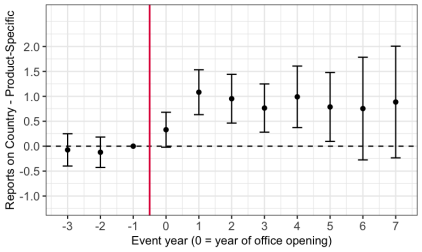
(a) Reports. Control: never-treated.



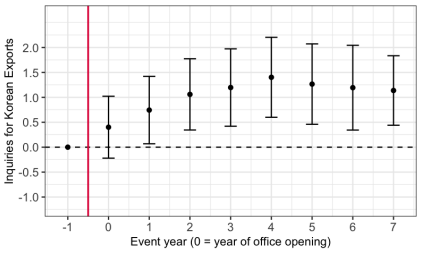
(b) Reports. Control: "not-yet".



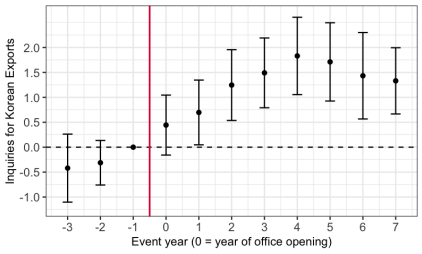
(c) Reports (product). C: never-treated.



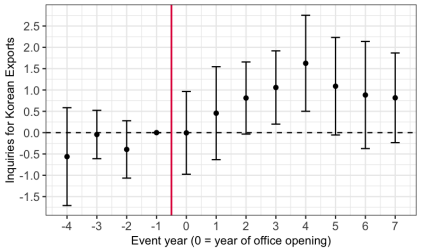
(d) Reports (product). C: "not-yet".



(e) Inquiries. Control: never-treated.



(f) Inquiries. Control: "not-yet".



(g) Inquiries. Control: never-treated. Openings from 1978.

Notes: The left panels reports coefficients $\hat{\pi}_k$ from estimating equation (7) when explaining three different measures of KOTRA activity regarding a specific country: number of reports written, number of product-specific reports, number of inquiries obtained - each transformed by the inverse hyperbolic sine. The right panels do the same following the approach by Callaway and Sant'Anna (2021). Instead of exports, we aim to explain three measures of KOTRA activity, each transformed using the inverse hyperbolic sine. (1) The number of reports about a country, (2) the number of product-specific reports - which may be more specific or informative, (3) the number of inquiries for trade with the country. The data on reports covers the years 1965 to 2001. We thus exclude events before 1968 from the analysis in panels (a)-(d). The data on inquiries covers the years 1974 to 1997. We thus exclude events before 1974 from the analysis in panels (e) and (f). Including events from 1975 comes at the cost of estimating only 1 pre-period in panel (e). Panel (g) takes the alternative approach of including multiple pre-periods, at the cost that the sample of treated countries is restricted to those with an event between 1978 and 1981.

B Appendix Tables

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

	Actual year of opening	Rank by 1962 market size	Predicted year of opening
UK	1965	1	1965
Italy	1966	4	1967
Netherlands	1966	5	1967
W Germany	1967	2	1966
Switzerland	1967	8	1970
France	1969	3	1966
Sweden	1969	7	1969
Austria	1970	12	1973
Belgium	1972	6	1969
Spain	1972	10	1972
Denmark	1973	9	1972
Norway	1973	11	1973
Finland	1973	13	1973
Greece	1973	15	1973
Turkey	1973	16	1973
Ireland	1973	14	1973
Portugal	1974	17	1974

Notes: The table compares the actual and predicted year of office openings for 17 European countries that received offices in the main office roll-out period (1962-1981). The predicted year reassigns the actual year based on the rank by pre-determined market size, measured by 1962 non-Korean exports to that country. For example, Italy ranks 4, so its predicted year is the 4^{th} among the actual opening years, 1967.

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

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

Results from estimating equation (1) reported. An observation is a product-country-year. The outcome variable is exports after residualizing by product-year and country fixed effects. A country is included for all the years that it has an office and is linked to a bureaucrat. A product is included for all the years in which Korea exported it to any country. Product refers to 4-digit SITC Rev. 2 codes. S.D. of IHS exports : 2.45, s.d. of IHS exports | tp, c: 1.83. The increase in R^2 due to bureaucrat FE is most meaningfully compared to the increase due to country FE – 0.018 compared to 0.097. These levels are lower than reported in the variance decomposition as the latter bundles all observations within an appointment while this table retains separate observations for each product, thus including variation that cannot be explained by product-invariant explanatory variables such as country FE and bureaucrat FE.

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

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

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

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 bureaucrats and countries 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, we observe that violations of additive separability would result in residuals with high absolute values for certain kinds of bureaucrat-country pairs (Fenizia, 2022; Otero and Muñoz, 2025). Following the literature, we divide our observations based on the quartile of the estimated manager fixed effect and the quartile of the estimated country fixed effect. For example, if – contrary to the linearity assumption – bureaucrats mattered more in small countries, we would expect large positive (negative) residuals for observations with top (bottom) quartile bureaucrats in bottom quartile countries. Figure A.9 shows that mean residuals do not exhibit any clear pattern such as the ones described above. 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 assumption of additive separability.

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.[75] Table B.2 reports the variation explained by the different levels of fixed-effects when estimating equation (1), which identifies the causal effect of the two sets of fixed effects under the assumptions discussed above.[76] Adding bureaucrat fixed effects increases R^2 by 0.018, about 18.6% of the increase in explanatory power from adding country fixed effects.[77] To understand whether bureaucrat effects differ between appointments, we

[75] 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).

[76] While informative, these are subject to some of the criticisms addressed by the Kline, Saggio, and Sølvsten (2020) bias correction reported in Table 2.

[77] 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 is smaller than other recent papers, studying managers of organizations that process insurance claims (increase in R^2 of 0.11, Fenizia), or hospital CEOs (0.09, Otero

compare the explanatory power when including appointment fixed effects (column 4) compared to column 3 which reports results from our main specification that assumes that bureaucrat and country effects are additively separable. The increase in explanatory power from this is negligible, suggesting that bureaucrat effects are relatively stable across appointments, which assuages concerns that the productivity of a bureaucrat-country match is not approximated well by the linear combination of the bureaucrat fixed effect. Further, if bureaucrat fixed effects mainly picked up statistical noise, allowing for appointment-specific effects would likely increase explanatory power more than 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 A.10 shows time trends in residualized exports around the year when an office experiences a change in the manager. 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). These are obtained from average de-trended exports of a product during a bureaucrat's appointments, i.e. bureaucrat fixed effects after residualizing exports by product-country and product-year fixed effects.

Appendix Figure A.10 finds that the main takeaways from Figure 5 are present for transitions across all terciles of incoming and outgoing bureaucrats bureaucrats. First, in the pre-periods, exports are highest when the outgoing bureaucrat is in the top tercile and lowest when the outgoing bureaucrat is in the bottom tercile. Second, in the post-period, the effect of the outgoing bureaucrat's tercile becomes less important, the effect of the incoming bureaucrat's tercile becomes dominant. In year one – the second full year of the incoming bureaucrat – exports are lowest when the incoming bureaucrat is in the bottom tercile. They are highest when the incoming bureaucrat is in the top tercile. Third, exports change sharply, and in the expected direction, precisely when a destination switches to a less or more effective bureaucrat. Exports increase the most upon a switch to the highest tercile and (relatively) decrease the most upon a switch to the lowest tercile. The figure shows little sign that exports are differentially increasing in countries that subsequently switch to a better bureaucrat, and vice versa. This suggests that drift in effectiveness and switches are uncorrelated.

Overall, this section assuages concerns regarding the additive separability between bureaucrat and country effects implicit in equation (3).

and Muñoz). Relative to the explanatory power increase from adding country or organization fixed effects, the increase in R^2 due to bureaucrats is slightly smaller than other recent papers studying bureaucrats who run organizations that process insurance claims (23.4%, Fenizia) and public hospitals (28.0%, Otero and Muñoz).

C.2 How Estimated Ability Affect Extensive and Intensive Margin

This section unpacks the effect on the inverse hyperbolic sine of exports into the extensive and intensive margin. We find that bureaucrat effects cause increases both along the extensive and the intensive margin. Hence, both margins together explain the increase in the inverse hyperbolic sine of exports implied by the fixed effects.

Appendix Figure A.11 reports the event study estimates of bureaucrat effects estimated from equation (5) with the outcome variable replaced by a dummy indicating whether South Korean exports of a particular product to this country-year exceeded 0. There is no indication of differential pre-trends. In event years 0 and 1, the new bureaucrat's ability increases the likelihood of positive exports of a given product by 5-7 percentage points, a sizable effect. The old bureaucrat's ability decreases it by the same amount, suggesting that losing bureaucrat ability has symmetric effects to gaining such ability.

Appendix Figure A.12 reports the estimates using only the sample of products with extensive margin changes. 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.

Appendix Table B.3 shows that the number of products with extensive margin changes remains roughly constant across decades. So the extensive margin response remains similarly important over time. However, Appendix Table B.3 shows an increase over time in the number of products for which only the intensive margin matters.[78] Appendix Figure A.13 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 inverse hyperbolic sine of exports – increases both the intensive and extensive margin of exports.

[78] The omitted – shrinking – category contains products without any exports throughout the event horizon.

C.3 Sensitivity of Fixed Effects to Alternative Transformations of Exports

This section investigates how much the bureaucrat fixed effect estimates would change under alternative transformations of export values.

Appendix Figure A.6 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 the extensive margin.

Appendix Figure A.6 panels (b) and (c) display a scatter plot comparing the benchmark estimates to alternative estimates based on the following transformation, following Chen and Roth (2023):

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

Under this transformation, the weight of the extensive margin is explicitly assigned via the choice of m. Panel (b) displays alternative vs. benchmark estimates based on rather extreme values of m: 0.05, 2.00, 10.00. While the estimates differ based on the value of m, they remain close to the 45-degree line. Thus, benchmark fixed effects are very predictive of fixed effects under these alternative transformations. There is slightly less dispersion when valuing the extensive margin less ($m = 0.05$) and somewhat more when valuing the extensive margin greatly ($m \in \{2, 10\}$).

Panel (c) shows that the benchmark estimates' dispersion of fixed effects lies between estimates based on $m = 0.50$ and $m = 1.00$, providing bounds for weight implicitly assigned to the extensive margin by the inverse hyperbolic sine transformation. It's further noteworthy that 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.

D Connected Set and Leave-One-Out Connected Set

Appendix Figure A.5 illustrates the connected set – connected via moves of bureaucrats connecting all countries. Panel (c) also highlights that if the movement of a single bureaucrat (b_3) connects two separate connected sets – e.g., the bottom-left to the top-right – then the estimated difference in fixed effect between countries, and hence bureaucrats, in the two sets of countries is strongly affected by any shock that occurs during b_3's appointments. For instance, if there is a positive shock of size $\hat{\epsilon}$ in country c_2 during b_3's appointment to c_2, estimating equation (1) would yield a positive bias in the fixed effect estimate for $\hat{\gamma}_{c_2}$ (relative to $\hat{\gamma}_{c_1}$ and $\hat{\gamma}_{c_3}$). This would spill over into a negative bias in $\hat{\beta}_{b_2}$, and hence a positive bias in $\hat{\beta}_b$ for all bureaucrats b that are ever only appointed to countries c_2 or c_4. These biases result in the limited mobility bias: the variance of estimated bureaucrat fixed effects overstates the variance in exports due to individual bureaucrats.

Panel (e) 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, under weak assumptions the leave-one-out connected set allows us to correct the variance in bureaucrat fixed effects and obtain consistent estimates of the variance in ability (Kline, Saggio, and Sølvsten, 2020).

Appendix Figure A.5, panels (b), (d), and (f), highlight how many connections offices or countries can have with only three different bureaucrats. The panels display all the connections between the UK and other countries because of only three consecutive appointments of managers to London (1981, 1984, and 1987).

E Covariance Shrinkage and 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[79]

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

E.1 Bootstrap bias correction

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: $\hat{\theta}_b = \theta_b + \nu_b$, where ν_b is sampling error. The variance of the estimated bureaucrat fixed effects, $\mathsf{Var}_b(\hat{\theta}_b) = \mathsf{Var}_b(\theta_b) + \mathsf{Var}_b(\nu_b) = \sigma^2_\theta + \sigma^2_\nu$, will overstate the true dispersion in bureaucrat quality, σ^2_θ.

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

$$\sigma^2_\nu = \mathsf{E}_b(\sigma^2_{\nu_b}) \Rightarrow \hat{\sigma}^2_\nu = \frac{\sum_{b=1}^n \hat{\sigma}^2_{\nu_b} w_b}{\sum_{b=1}^n w_b}$$

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, $\hat{\sigma}^2_{\nu_b}$, is estimated by bootstrapping[80] equation (2) K[81] times to get the variance of $\hat{\theta}_b$ across bootstrap samples

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

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

[81] We choose $K = 500$, which we confirm to be a sufficiently large number so that the sampling-

$k = 1, ..., K$:

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

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:

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

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

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

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

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

$$\sigma_{vu} = \mathsf{E}_{bc}(\sigma_{v_b u_c}) \Rightarrow \hat{\sigma}_{vu} = \frac{\sum_{b=1}^n \sum_{c=1}^C \hat{\sigma}_{v_b u_c} w_{bc}}{\sum_{b=1}^n \sum_{c=1}^C w_{bc}} \tag{8}$$

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[82] is again estimated from the bootstrap results:

$$\sigma_{v_b u_c} = \mathsf{Cov}(v_b, u_c) \Rightarrow \hat{\sigma}_{v_b u_c} = \widehat{\mathsf{Cov}}_k(\hat{\theta}^k_b, \hat{\gamma}^k_c) = \frac{\sum_{k=1}^K (\hat{\theta}^k_b - \bar{\bar{\hat{\theta}}}_b)(\hat{\gamma}^k_c - \bar{\bar{\hat{\gamma}}}_c) w^k_{bc}}{\sum_{k=1}^K w^k_{bc}} \tag{9}$$

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

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

error variance and covariance estimates stabilize.

[82] If b and c are not matched, i.e. bureaucrat b is never appointed to country c, then $\sigma_{v_b u_c} = 0$.

E.2 Shrinkage

Although $\hat{\theta}_b$ are unbiased estimates of θ_b, they are inconsistent in small samples. Shrinking $\hat{\theta}_b$ makes them biased but reduces their variance. The shrinkage factor for each bureaucrat is chosen to minimize their mean-squared-error:

$$\lambda_b^* = \arg\min_{\lambda_b} \underbrace{\mathsf{E}\left[\left(\theta_b - \lambda_b \hat{\theta}_b\right)^2\right]}_{\text{MSE}_b^{\text{Sh}}} = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2}$$

The shrinkage estimator of bureaucrat b's fixed effect is therefore

$$\hat{\theta}_b^{\text{Sh}} = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2} \cdot \hat{\theta}_b$$

The variance of the shrunk bureaucrat fixed effects theoretically understates the true variance:

$$\begin{aligned}
\mathsf{Var}_b(\hat{\theta}_b^{\text{Sh}}) &= \mathsf{E}_b\left[\mathsf{Var}(\hat{\theta}_b^{\text{Sh}})\right] + \mathsf{Var}_b\left(\mathsf{E}(\hat{\theta}_b^{\text{Sh}})\right) \quad \text{(Law of total variance)} \\
&= \mathsf{E}_b\left[\mathsf{Var}(\hat{\theta}_b^{\text{Sh}})\right] + \mathsf{Var}_b\left(\mathsf{E}\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2}\hat{\theta}_b\right)\right) \\
&= \mathsf{E}_b\left[\mathsf{Var}(\hat{\theta}_b^{\text{Sh}})\right] + \mathsf{Var}_b\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2}\mathsf{E}(\hat{\theta}_b)\right) \\
&= \mathsf{E}_b\left[\mathsf{Var}(\hat{\theta}_b^{\text{Sh}})\right] + \mathsf{Var}_b\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2} \cdot 0\right) \\
&= \mathsf{E}_b\left[\mathsf{Var}(\hat{\theta}_b^{\text{Sh}})\right] \\
&= \mathsf{E}_b\left[\mathsf{Var}\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2} \cdot \hat{\theta}_b\right)\right] \\
&= \mathsf{E}_b\left[\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2}\right)^2 \mathsf{Var}(\hat{\theta}_b)\right] \\
&= \mathsf{E}_b\left[\left(\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2}\right)^2 (\sigma_\theta^2 + \sigma_{v_b}^2)\right] \\
&= \mathsf{E}_b\left[\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2} \cdot \sigma_\theta^2\right] \\
&= \mathsf{E}_b\left[\frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_{v_b}^2}\right] \cdot \sigma_\theta^2 < \sigma_\theta^2
\end{aligned} \tag{10}$$

To compute $\hat{\theta}_b^{\text{Sh}}$ in practice, σ_θ^2 and $\sigma_{v_b}^2$ are not known so we replace them with $\hat{\sigma}_\theta^2$ and $\hat{\sigma}_{v_b}^2$, calculated according to the formula in Appendix E.1.

E.3 Covariance-shrinkage

Covariance-shrinkage, developed by Best et al. (2023), is an extension of the common shrinkage method applicable in the two-way fixed effects setup, which explicitly accounts for the correlation between the two dimensions of fixed effects.

Covariance shrinkage weights the full vector of bureaucrat and country fixed effects to minimize the mean-squared-error for each bureaucrat or country fixed effect. The covariance shrinkage matrix is found by stacking the minimization problem for all fixed effects:

$$\mathbf{\Lambda}^* = \arg\min_{\mathbf{\Lambda}} \underbrace{\mathsf{E}[(\boldsymbol{\rho} - \mathbf{\Lambda}\hat{\boldsymbol{\rho}})'(\boldsymbol{\rho} - \mathbf{\Lambda}\hat{\boldsymbol{\rho}})]}_{\text{MSE}^{\text{CovSh}}} \tag{11}$$

where $\boldsymbol{\rho} = (\theta_1, \theta_2, ..., \theta_n, \gamma_1, ..., \gamma_C)'$.

$$\begin{aligned} \text{MSE}^{\text{CovSh}} &= \mathsf{E}(\boldsymbol{\rho}'\boldsymbol{\rho}) - 2\mathsf{E}(\boldsymbol{\rho}'\mathbf{\Lambda}\hat{\boldsymbol{\rho}}) + \mathsf{E}\left((\mathbf{\Lambda}\hat{\boldsymbol{\rho}})'\mathbf{\Lambda}\hat{\boldsymbol{\rho}}\right) && (12) \\ &= \text{Tr}(\mathbf{\Sigma}_{\boldsymbol{\rho}\boldsymbol{\rho}}) - 2\text{Tr}(\mathbf{\Lambda}\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\boldsymbol{\rho}}) + \text{Tr}(\mathbf{\Lambda}\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}\mathbf{\Lambda}') && (13) \end{aligned}$$

$$\frac{\partial \text{MSE}^{\text{CovSh}}}{\partial \mathbf{\Lambda}}\Big|_{\mathbf{\Lambda}=\mathbf{\Lambda}^*} = -2\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\boldsymbol{\rho}} + 2\mathbf{\Lambda}^*\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}} = 0 \tag{14}$$

$$\Rightarrow \mathbf{\Lambda}^* = \mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\boldsymbol{\rho}}\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}^{-1} \tag{15}$$

where $\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\boldsymbol{\rho}}$ is the covariance matrix between $\hat{\boldsymbol{\rho}}$ and $\boldsymbol{\rho}$, and $\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}$ is the variance-covariance matrix of $\hat{\boldsymbol{\rho}}$.

Specifically,

$$\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\boldsymbol{\rho}} = \begin{pmatrix} \sigma_\theta^2 \mathbf{I}_n & \sigma_{\theta\gamma}\mathbf{1}_{n\times C} \\ \sigma_{\theta\gamma}\mathbf{1}_{C\times n} & \sigma_\gamma^2 \mathbf{I}_C \end{pmatrix} \tag{16}$$

because

$$(\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\boldsymbol{\rho}})_{ij} = \mathsf{Cov}(\hat{\rho}_i, \rho_j) = \mathsf{Cov}(\rho_i, \rho_j) = \begin{cases} \mathsf{Var}(\theta_i) = \sigma_\theta^2 & \text{if } i = j, i \leq n \\ \mathsf{Cov}(\theta_i, \theta_j) = 0 & \text{if } i \neq j, i \leq n, j \leq n \\ \mathsf{Cov}(\theta_i, \gamma_{j-n}) = \sigma_{\theta\gamma} & \text{if } i \neq j, i \leq n, j > n \\ \mathsf{Var}(\gamma_{j-n}) = \sigma_\gamma^2 & \text{if } i = j, j > n \\ \mathsf{Cov}(\gamma_{i-n}, \gamma_{j-n}) = 0 & \text{if } i \neq j, i > n, j > n \\ \mathsf{Cov}(\gamma_{i-n}, \theta_j) = \sigma_{\theta\gamma} & \text{if } i \neq j, i > n, j \leq n \end{cases}$$

Moreover,

$$\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}} = \mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\boldsymbol{\rho}} + \begin{pmatrix} \text{diag}(\{\sigma_{v_b}^2\}_{b=1}^n) & \mathbf{\Gamma} \\ \mathbf{\Gamma}' & \text{diag}(\{\sigma_{u_c}^2\}_{c=1}^C) \end{pmatrix} \tag{17}$$

where $\mathbf{\Gamma} = [\sigma_{v_b u_c}]$ is a $n \times C$ matrix of the covariances of sampling errors, because

$$(\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}})_{ij} = \mathsf{Cov}(\hat{\rho}_i, \hat{\rho}_j) = \mathsf{Cov}(\rho_i, \rho_j) + \begin{cases} \mathsf{Var}(v_i) = \sigma^2_{v_i} & \text{if } i = j, i \leq n \\ \mathsf{Cov}(v_i, v_j) = 0 & \text{if } i \neq j, i \leq n, j \leq n \\ \mathsf{Cov}(v_i, u_{j-n}) = \sigma_{v_i u_{j-n}} & \text{if } i \neq j, i \leq n, j > n \\ \mathsf{Var}(u_{j-n}) = \sigma^2_{u_{j-n}} & \text{if } i = j, j > n \\ \mathsf{Cov}(u_{i-n}, u_{j-n}) = 0 & \text{if } i \neq j, i > n, j > n \\ \mathsf{Cov}(u_{i-n}, v_j) = \sigma_{v_j u_{i-n}} & \text{if } i \neq j, i > n, j \leq n \end{cases}$$

and $\sigma_{v_b u_c} = 0$ if (b, c) is not a matched bureaucrat-county pair.

The covariance-shrunk fixed effect of bureaucrat b is

$$\hat{\theta}_b^{\text{CovSh}} = \mathbf{\Lambda}_b^{*\prime} \hat{\boldsymbol{\rho}}$$

where $\mathbf{\Lambda}_b^*$ is the vector of the b^{th} row of $\mathbf{\Lambda}^*$, i.e.

$$\mathbf{\Lambda}^* = \begin{bmatrix} \mathbf{\Lambda}_1^{*\prime} \\ \vdots \\ \mathbf{\Lambda}_{n+C}^{*\prime} \end{bmatrix} \tag{18}$$

Thus, $\mathbf{\Lambda}_b^* = \mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}^{-1} \boldsymbol{\sigma}_{\theta_b \hat{\boldsymbol{\rho}}}$, where $\boldsymbol{\sigma}_{\theta_b \hat{\boldsymbol{\rho}}}$ is the vector of the covariance between θ_b and $\hat{\boldsymbol{\rho}}$.

As with standard shrinkage, the variance of the covariance-shrunk bureaucrat fixed effects also theoretically understates the true variance, though by less:

$$\begin{aligned} \mathsf{Var}_b(\hat{\theta}_b^{\text{CovSh}}) &= \mathsf{E}_b\left(\mathsf{Var}(\hat{\theta}_b^{\text{CovSh}})\right) + \mathsf{Var}_b\left(\mathsf{E}(\hat{\theta}_b^{\text{CovSh}})\right) \\ &= \mathsf{E}_b\left(\mathsf{Var}(\hat{\theta}_b^{\text{CovSh}})\right) + 0 \\ &= \mathsf{E}_b\left(\mathsf{Var}(\mathbf{\Lambda}_b^{*\prime}\hat{\boldsymbol{\rho}})\right) \\ &= \mathsf{E}_b\left(\mathsf{E}\left[(\mathbf{\Lambda}_b^{*\prime}\hat{\boldsymbol{\rho}} - \mathsf{E}(\mathbf{\Lambda}_b^{*\prime}\hat{\boldsymbol{\rho}}))^2\right]\right) \\ &= \mathsf{E}_b\left(\mathsf{E}\left[(\mathbf{\Lambda}_b^{*\prime}\hat{\boldsymbol{\rho}})^2\right]\right) \\ &= \mathsf{E}_b\left(\mathbf{\Lambda}_b^{*\prime}\mathsf{E}(\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}^\prime)\mathbf{\Lambda}_b^*\right) \\ &= \mathsf{E}_b\left(\mathbf{\Lambda}_b^{*\prime}\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}\mathbf{\Lambda}_b^*\right) \\ &= \mathsf{E}_b\left(\left(\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}^{-1}\boldsymbol{\sigma}_{\theta_b\hat{\boldsymbol{\rho}}}\right)^\prime \mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}\mathbf{\Lambda}_b^*\right) \\ &= \mathsf{E}_b\left(\boldsymbol{\sigma}_{\theta_b\hat{\boldsymbol{\rho}}}^\prime\mathbf{\Lambda}_b^*\right) \\ &= \mathsf{E}_b\left(\mathbf{\Lambda}_b^{*\prime}\boldsymbol{\sigma}_{\theta_b\hat{\boldsymbol{\rho}}}\right) \end{aligned} \tag{19}$$

Denote the minimized mean-squared-error of bureaucrat b as $\text{MSE}_b^{\text{CovSh}*}$. Since

mean-squared-error is always positive,

$$\begin{aligned} \text{MSE}_b^{\text{CovSh}^*} &= \mathsf{E}\left[(\theta_b - \boldsymbol{\Lambda}_b^{*\prime}\hat{\boldsymbol{\rho}})^2\right] \\ &= \mathsf{E}(\theta_b^2) - 2\boldsymbol{\Lambda}_b^{*\prime}\mathsf{E}(\theta_b\hat{\boldsymbol{\rho}}) + \mathsf{E}\left[(\boldsymbol{\Lambda}_b^{*\prime}\hat{\boldsymbol{\rho}})^2\right] \\ &= \sigma_\theta^2 - 2\boldsymbol{\Lambda}_b^{*\prime}\boldsymbol{\sigma}_{\theta_b\hat{\boldsymbol{\rho}}} + \boldsymbol{\Lambda}_b^{*\prime}\boldsymbol{\sigma}_{\theta_b\hat{\boldsymbol{\rho}}} \\ &= \sigma_\theta^2 - \boldsymbol{\Lambda}_b^{*\prime}\boldsymbol{\sigma}_{\theta_b\hat{\boldsymbol{\rho}}} > 0 \end{aligned} \tag{20}$$

$$\Rightarrow \boldsymbol{\Lambda}_b^{*\prime}\boldsymbol{\sigma}_{\theta_b\hat{\boldsymbol{\rho}}} < \sigma_\theta^2 \quad \forall\, b = 1, ..., n$$

Thus, $\mathsf{Var}_b(\hat{\theta}_b^{\text{CovSh}}) = \mathsf{E}_b(\boldsymbol{\Lambda}_b^{*\prime}\boldsymbol{\sigma}_{\theta_b\hat{\boldsymbol{\rho}}}) < \sigma_\theta^2$.

We also know that $\text{MSE}_b^{\text{CovSh}^*} < \text{MSE}_b^{\text{Sh}^*}$, because $\lambda_b\hat{\theta}_b \subseteq \boldsymbol{\Lambda}_b\hat{\boldsymbol{\rho}}$, i.e. minimizing bureaucrat b's mean-squared-error by weighting $\hat{\boldsymbol{\rho}}$ would do better than minimizing it by weighting $\hat{\theta}_b$ alone. Then,

$$\begin{aligned} & 0 < \text{MSE}_b^{\text{CovSh}^*} = \sigma_\theta^2 - \mathsf{Var}(\hat{\theta}_b^{\text{CovSh}}) < \text{MSE}_b^{\text{Sh}^*} = \sigma_\theta^2 - \mathsf{Var}(\hat{\theta}_b^{\text{Sh}}) \\ \Rightarrow\ & \mathsf{Var}(\hat{\theta}_b^{\text{Sh}}) < \mathsf{Var}(\hat{\theta}_b^{\text{CovSh}}) < \sigma_\theta^2 \\ \Rightarrow\ & \mathsf{E}_b\left[\mathsf{Var}(\hat{\theta}_b^{\text{Sh}})\right] < \mathsf{E}_b\left[\mathsf{Var}(\hat{\theta}_b^{\text{CovSh}})\right] < \sigma_\theta^2 \\ \Rightarrow\ & \mathsf{Var}_b(\hat{\theta}_b^{\text{Sh}}) < \mathsf{Var}_b(\hat{\theta}_b^{\text{CovSh}}) < \sigma_\theta^2 \end{aligned} \tag{21}$$

Therefore, the variance of covariance-shrunk bureaucrat fixed effects underestimates the true dispersion of bureaucrat ability, but by less than the variance of the (standard) shrunk ones.

Again, to compute $\hat{\theta}_b^{\text{CovSh}}$ in practice, we do not know the true variances and covariances σ_θ^2, σ_γ^2, $\sigma_{\theta\gamma}$, $\sigma_{v_b}^2$, $\sigma_{u_c}^2$, and $\sigma_{v_b u_c}$, so we replace them with the sample analogs described in Appendix E.1.

E.4 A 2×2 Example of the Covariance Shrinkage Weights

For illustrative purposes, we derive the covariance-shrinkage factor for the simple case with 1 bureaucrat and 1 country: $\boldsymbol{\rho} = (\theta_1, \gamma_1)'$

$$\boldsymbol{\Sigma}_{\hat{\boldsymbol{\rho}}\boldsymbol{\rho}} = \begin{bmatrix} \sigma_\theta^2 & \sigma_{\theta\gamma} \\ \sigma_{\theta\gamma} & \sigma_\gamma^2 \end{bmatrix} \tag{22}$$

$$\boldsymbol{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}} = \begin{bmatrix} \sigma_\theta^2 + \sigma_{v_1}^2 & \sigma_{\theta\gamma} + \sigma_{v_1 u_1} \\ \sigma_{\theta\gamma} + \sigma_{v_1 u_1} & \sigma_\gamma^2 + \sigma_{u_1}^2 \end{bmatrix} \tag{23}$$

$$\begin{aligned} \boldsymbol{\Lambda}^* &= \boldsymbol{\Sigma}_{\hat{\boldsymbol{\rho}}\boldsymbol{\rho}}\boldsymbol{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}^{-1} \\ &= \frac{1}{(\sigma_\theta^2 + \sigma_{v_1}^2)(\sigma_\gamma^2 + \sigma_{u_1}^2) - (\sigma_{\theta\gamma} + \sigma_{v_1 u_1})^2} \begin{bmatrix} \sigma_\theta^2 & \sigma_{\theta\gamma} \\ \sigma_{\theta\gamma} & \sigma_\gamma^2 \end{bmatrix} \begin{bmatrix} \sigma_\gamma^2 + \sigma_{u_1}^2 & -(\sigma_{\theta\gamma} + \sigma_{v_1 u_1}) \\ -(\sigma_{\theta\gamma} + \sigma_{v_1 u_1}) & \sigma_\theta^2 + \sigma_{v_1}^2 \end{bmatrix} \end{aligned}$$

$$\begin{pmatrix} \hat{\theta}_1^{\text{CovSh}} \\ \hat{\gamma}_1^{\text{CovSh}} \end{pmatrix} = \mathbf{\Lambda}^* \begin{pmatrix} \hat{\theta}_1 \\ \hat{\gamma}_1 \end{pmatrix} \tag{24}$$

Focusing on bureaucrat 1's fixed effect,

$$\hat{\theta}_1^{\text{CovSh}} = \mathbf{\Lambda}_{11}^* \hat{\theta}_1 + \mathbf{\Lambda}_{12}^* \hat{\gamma}_1 \tag{25}$$

where

$$\mathbf{\Lambda}_{11}^* = \frac{\sigma_\theta^2(\sigma_\gamma^2 + \sigma_{u_1}^2) - \sigma_{\theta\gamma}(\sigma_{\theta\gamma} + \sigma_{v_1 u_1})}{(\sigma_\theta^2 + \sigma_{v_1}^2)(\sigma_\gamma^2 + \sigma_{u_1}^2) - (\sigma_{\theta\gamma} + \sigma_{v_1 u_1})^2}, \quad \mathbf{\Lambda}_{12}^* = \frac{\sigma_{\theta\gamma}\,\sigma_{v_1}^2 - \sigma_\theta^2\,\sigma_{v_1 u_1}}{(\sigma_\theta^2 + \sigma_{v_1}^2)(\sigma_\gamma^2 + \sigma_{u_1}^2) - (\sigma_{\theta\gamma} + \sigma_{v_1 u_1})^2}$$

Equipped with the precise formula of the covariance shrinkage weights ($\mathbf{\Lambda}_{11}^*$ and $\mathbf{\Lambda}_{12}^*$), we can derive some comparative statics or compare them to the shrinkage factor (λ_1^*) to understand the covariance shrinkage method better. In general, covariance shrinkage improves on shrinkage by additionally utilizing the information about θ_1 contained in $\hat{\gamma}_1$, while taking into account the noisiness of $\hat{\gamma}_1$ and the strength of the correlation.

- If $\sigma_{\theta\gamma} = \sigma_{v_1 u_1} = 0$ ($\hat{\gamma}_1$ contains no information about θ_1), then $\hat{\theta}_1^{\text{CovSh}} = \hat{\theta}_1^{\text{Sh}}$.
- $\frac{\partial \mathbf{\Lambda}_{11}^*}{\partial \sigma_{v_1}^2} < 0$: Similarly to shrinkage, $\hat{\theta}_1$ is shrunk more if it is estimated with greater noise (higher $\sigma_{v_1}^2$).
- $\frac{\partial \mathbf{\Lambda}_{12}^*}{\partial \sigma_{u_1}^2} < 0$: Greater noise in estimating γ_1 (higher $\sigma_{u_1}^2$) reduces the covariance shrinkage weight on $\hat{\gamma}_1$.
- It is likely that $\mathbf{\Lambda}_{12}^* > 0$. Because the denominator is positive, the sign of $\mathbf{\Lambda}_{12}^*$ follows the sign of $(\sigma_{\theta\gamma}\,\sigma_{v_1}^2 - \sigma_\theta^2\,\sigma_{v_1 u_1})$. Moreover, limited mobility bias implies $\sigma_{v_1 u_1} < 0$ (Andrews et al., 2008).
  - If $\sigma_{\theta\gamma} > 0$, $\mathbf{\Lambda}_{12}^*$ is clearly positive.
  - Even if $\sigma_{\theta\gamma} < 0$, $\mathbf{\Lambda}_{12}^*$ is still positive under the likely scenario that $\hat{\theta}_1$ is more informative about θ_1 than $\hat{\gamma}_1$ is. The specific condition is that the signal-to-noise ratio (SNR) of $\hat{\theta}_1$ is larger than the SNR of $\mathsf{Cov}(\hat{\theta}_1, \hat{\gamma}_1)$, i.e. $\frac{\sigma_\theta^2}{\sigma_{v_1}^2} > \frac{|\sigma_{\theta\gamma}|}{|\sigma_{v_1 u_1}|} = \frac{\sigma_{\theta\gamma}}{\sigma_{v_1 u_1}}$. In that case, $\mathbf{\Lambda}_{12}^* > 0$ and $\mathbf{\Lambda}_{11}^* > \lambda_1^*$.[83] Intuition: Standard shrinkage reduces mean-squared-error by shrinking the noisy estimator $\hat{\theta}_1$ towards its mean value, $\mathsf{E}(\theta_1) = 0$. Less shrinkage is needed when the SNR of $\hat{\theta}_1$ is high, which means that $\hat{\theta}_1$ carries reliable information about θ_1. With covariance shrinkage, the SNR of the covariance also comes into play. $\mathbf{\Lambda}_{11}^* > \lambda_1^*$ reflects that $\hat{\theta}_1$ has higher SNR than the covariance, i.e. $\hat{\theta}_1$ is more informative about the value of θ_1 than $\hat{\gamma}_1$ is. At the same time,

[83] $(\mathbf{\Lambda}_{11}^* - \lambda_1^*) \propto (\sigma_{\theta\gamma} + \sigma_{v_1 u_1})(\sigma_{\theta\gamma}\,\sigma_{v_1}^2 - \sigma_\theta^2\,\sigma_{v_1 u_1})$.

a positive weight on $\hat{\gamma}_1$ ($\mathbf{\Lambda}^*_{12} > 0$) keeps $\hat{\theta}_1^{\text{CovSh}} = \mathbf{\Lambda}^*_{11}\hat{\theta}_1 + \mathbf{\Lambda}^*_{12}\hat{\gamma}_1$ from diverging away from 0 too much, because $\sigma_{\theta\gamma} < 0$.

– If $\sigma_{\theta\gamma} < 0$ and the SNR of $\hat{\theta}_1$ is *smaller* than the SNR of $\mathsf{Cov}(\hat{\theta}_1, \hat{\gamma}_1)$, i.e. $\frac{\sigma^2_{\theta}}{\sigma^2_{v_1}} < \frac{|\sigma_{\theta\gamma}|}{|\sigma_{v_1 u_1}|} = \frac{\sigma_{\theta\gamma}}{\sigma_{v_1 u_1}}$, then $\mathbf{\Lambda}^*_{12} < 0$ and $\mathbf{\Lambda}^*_{11} < \lambda^*_1$.

- If $\mathbf{\Lambda}^*_{12} > 0$, then $\frac{\partial \mathbf{\Lambda}^*_{12}}{\partial \sigma_{v_1 u_1}} < 0$: The positive weight on $\hat{\gamma}_1$ is higher the more closely connected the bureaucrat and country are (the stronger the negative covariance between the sampling errors of $\hat{\theta}_1$ and $\hat{\gamma}_1$ is).

E.5 Potential Pitfalls of Covariance Shrinkage

In general, the covariance shrinkage estimator improves upon the noisy raw fixed effect estimators by minimizing mean-squared-error. Yet, the covariance shrinkage method comes with two caveats.

1. The covariance-shrunk fixed effects are weighted averages of all raw fixed effects. Hence, it is possible for some of them to blow up – contrary to the idea of "shrinkage" toward the mean value 0 – if they non-trivially weigh a particularly large (positive or negative) raw fixed effect. The issue is compounded if they happen to non-trivially weigh *multiple* raw fixed effects that are large and of the same sign. Given that covariance shrinkage weights tend to be larger on closely connected nodes in the bureaucrat-country graph, this may occur, for instance, for a bureaucrat who got assigned to multiple large export destination countries by chance.

2. If $\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}$ is near-singular, often the case with large covariance matrices, then $\mathbf{\Sigma}^{-1}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}$ can contain some extreme values. Then, some covariance-shrunk fixed effects might explode due to these extremely large covariance shrinkage weights, even if the raw fixed effects they are weighting are small in size.

Both caveats imply that covariance shrinkage may produce a few extreme values. In that case, the variance of the covariance-shrunk fixed effects may not reliably measure the true dispersion of bureaucrat or country fixed effects. However, other characteristics of the distribution that are less sensitive to outliers, such as the median or quartiles, would still remain informative.

E.6 Comparison with Best et al. (2023)

We discovered a minor error in Best et al. (2023)'s derivation of the covariance shrinkage matrix. Specifically, the error is in the covariance matrix between $\hat{\boldsymbol{\rho}}$ and $\boldsymbol{\rho}$. In

their version,

$$\boldsymbol{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}}^{\text{BHS}} = \mathsf{diag}(\sigma_\theta^2 \mathbf{1}_n, \sigma_\gamma^2 \mathbf{1}_C),$$

so the covariance terms ($\sigma_{\theta\gamma}$) were omitted.

To gauge the impact of the error, we investigate the difference in the covariance shrinkage weights in the Best et al. (2023) version and the correctly derived version in Appendix Table E.4.

In Panel A, we list the top-10 covariance shrinkage weights in each version for $\hat{\theta}_1^{\text{CovSh}}$, the covariance shrunk fixed effect of bureaucrat 1 (`B1`; arbitrarily selected). Bureaucrat and country identifiers are assigned based on alphabetical order (e.g., `C1`: Afghanistan). In the correct version, the largest weight is put on the fixed effect of bureaucrat 1 himself (`B1`), followed by countries he was appointed to (`C36` and `C55` – the shortest-path distance between `B1` and `C36` or `C55` is 1 in the country-bureaucrat graph). The fourth largest weight goes to `B4`, who was also appointed to `C36` – the shortest-path distance between `B4` and `B1` is 2. More generally, more closely connected nodes in the graph are assigned larger weights, reflecting their stronger correlations with $\hat{\theta}_1$. In the Best et al. (2023) version, however, the ordering of the weights and the shortest-path distances are not as well-aligned.

In Panel B, we report the summary statistics on the covariance shrinkage weights for all bureaucrat and country fixed effects. The summary statistics are more sensible in the correct version than in the Best et al. (2023) version, again highlighting that more closely connected nodes are assigned larger weights: 1) the share of bureaucrats and countries whose top covariance shrinkage weight is on themselves is higher at 0.970; 2) the average share of matched pairs that get the next-largest weights[84] is higher at 0.782; 3) the average weight on non-matched pairs[85] is lower and closer to zero, at 0.0004 on average, implying fewer instances of non-matched pairs getting sizable weights.

Appendix Figure E.22 visually compares the two versions of covariance shrinkage to each other, as well as to standard shrinkage. The two covariance shrinkage versions are typically very similar to one another.

Shrinkage always compresses the bureaucrat fixed effects toward zero, such that all points lie between the 45-degree line and the horizontal axis. Both covariance

[84]The calculation of the share of matched pairs getting the next-largest weights is illustrated through the following example. For `B1` with two matched countries (`C36`, `C55`), the next-2 largest weights after the top weight go to both `C36` and `C55` in the correct version (share=1), but they go to only `C36` in the Best et al. (2023) version (share=0.5). The statistic in the table is the average of these shares for all bureaucrats and countries.

[85]Non-matched pairs refer to all bureaucrats and countries besides oneself and the countries or bureaucrats that they are directly matched to. E.g. for `B1`, they are all bureaucrats and countries besides himself, `C36`, and `C55`.

shrinkage versions mostly compress the bureaucrat fixed effects, too. Even for the fixed effects that get expanded away from zero (lying between the 45-degree line and the vertical axis) or switch in sign (lying in the second and fourth quadrants), most remain clustered around zero.

However, covariance shrinkage produces a small number of extreme values, related to caveats under Appendix E.5. Panel (b) shows that the covariance-shrunk (either version) fixed effects of bureaucrats appointed to Japan, one of Korea's largest trade partners, blow up because they absorb Japan's large fixed effect. Panels (c) and (d) show that the two versions of covariance shrinkage occasionally diverge. The Best et al. (2023) version produces extreme values for bureaucrats appointed to Turkey and Taiwan, the two countries that increase the condition number of $\mathbf{\Sigma}_{\hat{\rho}\hat{\rho}}$ the most.[86]

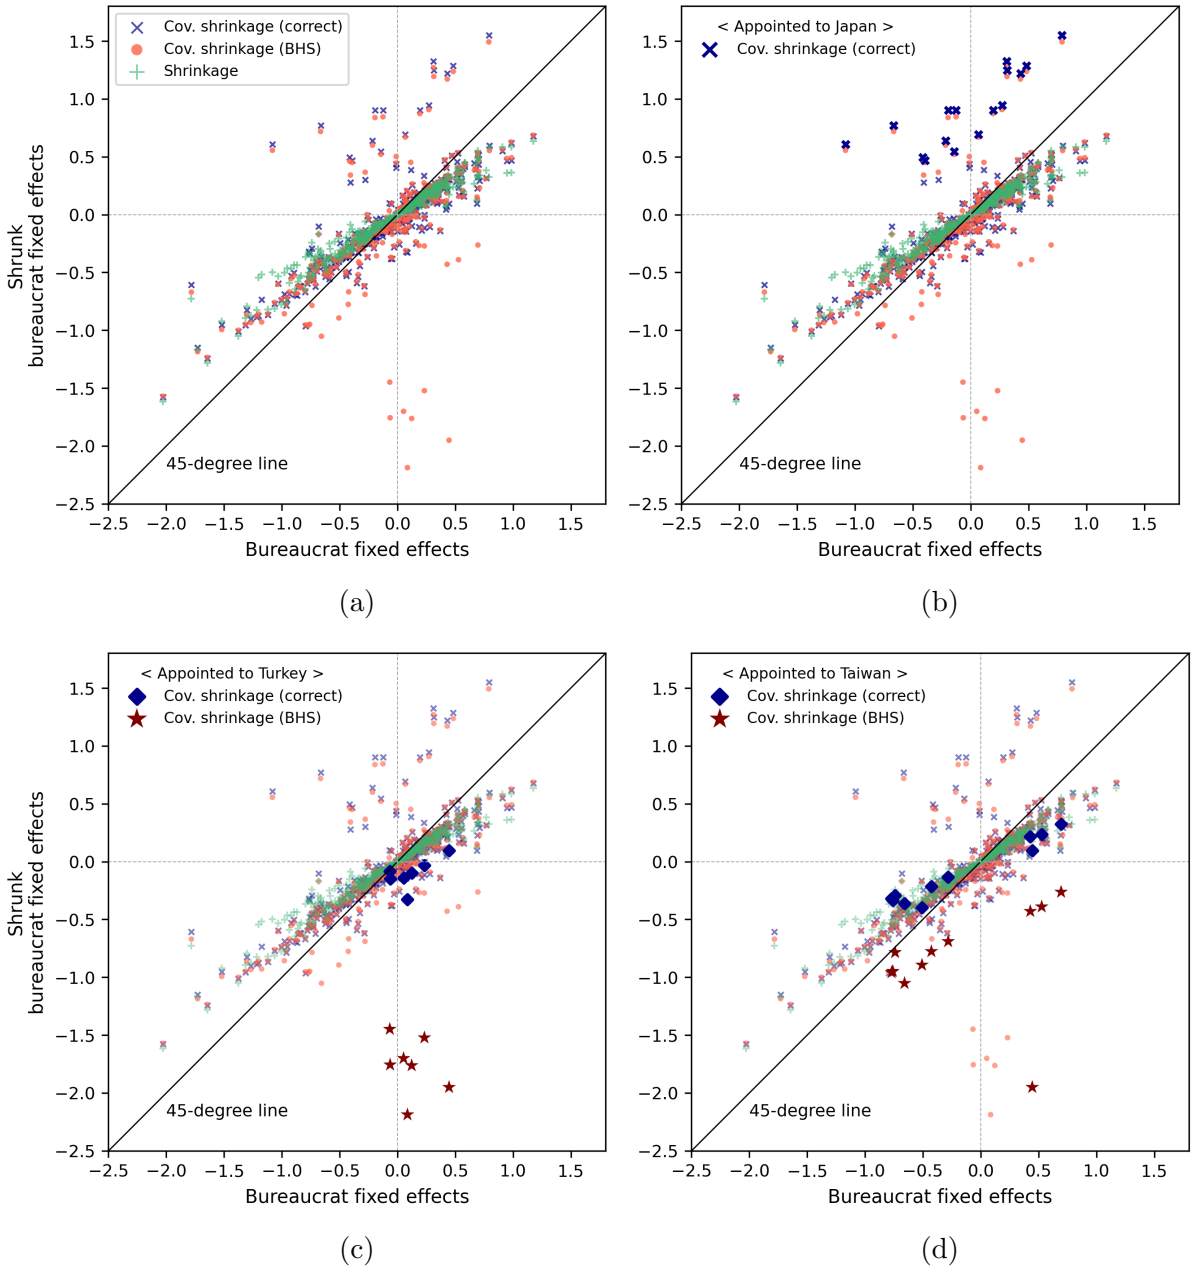
F Robustness of Appointments by Country Importance

This section conducts robustness checks on the finding in Appendix Figure A.14, 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. This main message is equally striking when reporting these distributions differently. Panels (c)-(f) restrict attention to bureaucrats whose first appointment started no earlier than 1981. This is to avoid that some results are due to a mechanical association between an office's total number of appointments and the duration of its existence. Panels (e) and (f) report the share of appointments to a given opening year (rank). This clearly shows the much higher likelihood of *late* offices to be managed by bureaucrats in their first appointment – around 50% of appointments to these offices – compared to *early* offices – around 25% of appointments to these offices. For bureaucrats in their third appointment, the numbers are similarly stark. Only around 15% of appointments to *late* offices are of bureaucrats in their third appointment, while the share for *early* offices is around 50%.

Appendix Figure A.15 displays similar results when measuring a country's importance as a market for South Korean exports by its fixed effect based on equation (2).

[86] For covariance matrices, a high condition number often indicates near-singularity, which can lead to extreme or unreliable results in matrix inversion. We compare the condition number of $\mathbf{\Sigma}_{\hat{\rho}\hat{\rho}}$ and $\mathbf{\Sigma}^{\text{BHS}}_{\hat{\rho}\hat{\rho}}$ as we remove one country from the sample at a time. The condition number falls the most dramatically when Turkey is removed, followed by Taiwan, for both matrices, but by more for $\mathbf{\Sigma}^{\text{BHS}}_{\hat{\rho}\hat{\rho}}$. Attesting to the importance of connection closeness for the covariance shrinkage weights, these two countries are connected through one common bureaucrat.

Figure E.22: Comparison of Different Shrinkage Methods



Notes: This figure shows the scatter plots of the raw bureaucrat fixed effects (horizontal axis) against the fixed effects shrunk using correctly-derived covariance shrinkage, the Best, Hjort, and Szakonyi (2023) version of covariance shrinkage, or standard shrinkage (vertical axis). Panel (a) is the benchmark figure, while panels (b), (c), and (d) each highlight the bureaucrats appointed to Japan, Turkey, and Taiwan, respectively. Panel (b) shows that covariance shrinkage (both versions) sometimes produces extreme values. Panels (c) and (d) show that the two versions of covariance shrinkage occasionally diverge.

Table E.4: Comparison of covariance shrinkage formulae

Panel A: Top-10 covariance shrinkage weights for shrinking $\hat{\theta}_1$

Correct version			**Best et al. (2023) version**		
i	Cov. shrinkage weight on i	Shortest-path distance to B1	i	Cov. shrinkage weight on i	Shortest-path distance to B1
B1	0.754	0	B1	0.755	0
C36	0.052	1	C36	0.055	1
C55	0.043	1	C80	0.051	5
B4	0.019	2	C55	0.045	1
B49	0.019	2	B203	0.044	4
B187	0.019	2	B83	0.042	4
B397	0.017	2	B251	0.036	6
B145	0.017	2	B179	0.034	6
B38	0.017	2	B91	0.033	6
B87	0.017	2	B16	0.032	6

Panel B: Summary statistics on covariance shrinkage weights

	Average across all bureaucrats and countries	
	Correct version	**Best et al. (2023) version**
$\mathbb{1}$(Top weight is on oneself)	0.970	0.945
Share of matched pairs getting largest weights	0.782	0.683
Weight on oneself	0.686	0.674
Average weight on matched pairs	0.074	0.085
Average weight on others	0.0004	0.0027

Notes: The table compares the covariance shrinkage weights in two versions: 1) when they are derived correctly ($\mathbf{\Lambda}^* = \mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}} \left(\mathbf{\Sigma}_{\hat{\boldsymbol{\rho}}\hat{\boldsymbol{\rho}}} + \mathbf{\Sigma}\right)^{-1}$, where $\mathbf{\Sigma} = \left[\mathsf{Cov}_k(\hat{\rho}_b^k, \hat{\rho}_c^k)\right]_{bc}$ is the covariance matrix of the bootstrap estimates), and 2) when they follow Best et al. (2023) ($\mathbf{\Lambda}^* = \mathsf{diag}(\sigma_\theta^2 \mathbf{1}_n, \sigma_\gamma^2 \mathbf{1}_C) \left(\mathsf{diag}(\sigma_\theta^2 \mathbf{1}_n, \sigma_\gamma^2 \mathbf{1}_C) + \mathbf{\Sigma}\right)^{-1}$). In Panel A, we list the top-10 covariance shrinkage weights in each version for $\hat{\theta}_1^{\text{CovSh}}$, the covariance shrunk fixed effect of bureaucrat 1 (B1; arbitrarily selected). Bureaucrat and country identifiers are assigned based on alphabetical order (e.g., C1: Afghanistan). In Panel B, we report the summary statistics on the covariance shrinkage weights for all bureaucrat and country fixed effects. Both panels show that closely connected nodes are more consistently assigned larger weights in the correct version than in the Best et al. (2023) version. See the text description for details.

The density of bureaucrats' first appointments is higher in low fixed effect countries than for third appointments. The opposite holds for bureaucrats' third appointments, with 2nd appointments forming an intermediate case. Panels (e) and (f) again provide the starkest contrast. Among countries with a negative fixed effect, 40% of appointments are bureaucrats' first appointments, 20% of appointments are bureaucrats' third appointments. This reduces to 10-20% of first appointments for the 15 countries with the highest fixed effects. For third appointments, the share among the top 15 countries is much higher: 40-70%.

Appendix figures A.16 and A.17 highlight how bureaucrats move from their first appointment to second and third appointments, or to exiting KOTRA. First, around half of bureaucrats exit after their first appointment. The exit rate appears unrelated to the tercile of the office of the bureaucrat's first appointment. This holds when classifying into terciles both by opening year as well as country effects. Second, between appointments one and two, bureaucrats move between all terciles of countries – despite some persistence. Third, between appointments two and three, few bureaucrats stay in the least important countries: almost all bureaucrats leave the third tercile of openings (latest openings) and the first tercile of country effects (lowest fixed effect). The opposite is not true: bureaucrats largely stay in the most important countries – the first tercile of openings and the third tercile of country effects.[87]

It should be noted that more important countries may also be more desirable for bureaucrats. So such patterns of appointments are 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.

G Robustness of Office Opening Results

G.1 Assessing Longer Pre-Trends by Restricting the Sample

This section investigates concerns that the results in 9 may be influenced by the choice of excluding events occurring before 1967. Appendix Figure A.18 reports estimates when expanding attention to events that happen from 1966 (from 1964) on. Panel (c) points to parallel trends between periods -4 and -2 with a moderate uptick in period -1. Panel (d) more convincingly finds the same convincing parallel trends reported in Figure 9. The uptick in period -1 is discussed further below when we check the sensitivity of our estimates to parallel trend violations following Rambachan and Roth (2023).

G.2 Not-Yet-Treated Control Group

This section allays concerns that the estimated effect of opening an office is driven by the choice of the never-treated as the control group. We check robustness to using a not-yet-treated control group instead, following Callaway and Sant'Anna (2021)'s approach, which yields consistent estimators in cases where the two-way fixed effects

[87] From appointments one to two, there is already somewhat limited mobility from the first to third tercile of openings. This is not as striking between terciles of country effects.

regression with a simple treatment indicator could fail. The estimates, plotted in Appendix Figure A.19 Panel (a.i), are of very similar magnitude and precision to the ones from the main estimation strategy. 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 (Rambachan and Roth, 2023).

Panel (a.ii) shows that the estimate of the effect on exports 10 years after an office opening remains statistically significant at the 95% level when allowing for parallel trend violations up to one times the largest pre-treatment violation. Panel (a.iii) shows that it remains significant when allowing for the slope of the violation of parallel trends to change by up to 0.15.

One period of anticipation would suggest that KOTRA has an effect on exports in the year before opening an office. A statistically significant negative effect in period -2 may point to violations of the parallel trends assumption. On the one hand, the measured opening year is meant to capture the year when the office becomes operational, i.e. the first year in which it can affect exports. If the office can have no effect in year -1, the jump in exports in this year could be explained by KOTRA choosing to open offices in the year when a country becomes a more important destination market. This would constitute a violation of the parallel trends assumption that would upwardly bias the estimated treatment effect. Panels (a.ii) and (a.iii) report the sensitivity of the estimated parallel trends assumption when concluding that the jumps in year -1 cannot be a causal effect of the office openings. On the other hand, setting up an office already requires resources dedicated to a country that may have a direct or indirect effect on exports. A direct effect would occur if in year -1 the KOTRA bureaucrats setting up the office already engage in KOTRA's usual activities. This appears plausible if testing different strategies to promote exports is an important component of the activities of setting up an office. Appendix Figure A.21 shows that this is not the case regarding reports – which do not go up prior to an office opening — though it may be the case regarding obtaining inquiries from potential importers. An indirect effect may occur in year -1 if the (planned) presence of a KOTRA office partly functions as a signal that is interpreted as indicating export potential by potential Korean exporters. An effect due to coordination would be interesting, as an important role of industrial policy is to coordinate industrial activity, often justified by making reference to potential positive externalities.

G.3 Extensive Margin

This section investigates whether the results in Figure 9 are artifacts of transforming the raw export values using the inverse hyperbolic sine. To do so, it investigates

the product-level extensive margin of exports. Concretely,were-estimate equation (7) with y_{cpt} changed to a dummy indicating whether there are positive exports from South Korea of product p to country c in year t. The question under investigation becomes: Does a KOTRA office in a country increase the likelihood of positive exports of a particular product from South Korea to the country? Or alternatively: Does a KOTRA office in a country raise the share of products that Korea exports to that country?

Appendix Figure A.20 reports the estimated effects of office openings 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. While the magnitude is not directly comparable, the trajectory of the point estimates in panel (a) is very similar to the main results reported in Figure 9. Again, the pre-treatment coefficient is very close to zero, corroborating that the parallel pre-trends in Figure 9 are not due to the inverse hyperbolic sine somehow obscuring differential pre-trends. On the other hand, there appears to be a treatment effect due to the office opening: the effect on the outcome variable slopes upwards starting with the office opening. The estimated coefficients become economically sizable as early as one year after the opening and statistically significant at the 5%-level two years after the opening. As before, the coefficients stabilize after around ten years.

The results remain qualitatively similar when restricting attention to openings between 1967 and 1981, which allows for estimating coefficients in the five years prior to the opening (panel (b)).

Panels (c) and (d) replicate these results using the not-yet-treated control group. Panel (c) does so while assuming 0 periods of anticipation. While the estimated coefficients in the post-period are very similar to the ones in panels (a) and (b), panel (c) casts some doubts on the parallel trends assumption: the estimated coefficients for 2, 3, and 4 years prior to the opening are all negative and statistically significant. Panel (d) shows that this anticipation effect seems to mainly occur in the year prior to the office opening. Hence, the discussion about the timing of KOTRA's effect from Section G.2 applies here.