Export-Oriented Industrial Policy and the East Asian Miracle

Nathan Lane* Max Marczinek† Kristen Dilan Yang‡

Preliminary conference draft
This version: November 11, 2026

Abstract

What role did industrial policy play in the East Asian growth miracle? We focus on export-oriented industrial policies and explore how the turn to these new policies coincided with aggregate economic development in postwar Asia. We combine historical research that dates and defines export-oriented industrial policy with a synthetic difference-in-differences design to estimate the turn's impact. This lets us compare the evolution of Japan, Taiwan, South Korea, and Singapore, the first wave of adopters, with that of synthetic counterfactual economies. Using cohort-specific versions of this design, with historically grounded estimation windows and placebo-based inference, we estimate an average cohort effect—the average over the first two decades after the turn—of 0.449 log points on GDP per capita and 1.356 log points on exports per capita. Ten years after adoption, the estimated effects are 0.417 log points on GDP per capita and 1.317 log points on exports per capita. We show that growth in exports and income also coincides with investment, capital deepening, and employment, but we do not detect improvement in common indices of human capital. Using simple measures of total factor productivity (TFP), we find evidence of productivity gains, though less robust than the output effects, and show that TFP may have contributed positively to growth under several quantitative assumptions.

Keywords: industrial policy, economic development, East Asia, developmental state, export-oriented industrialisation

JEL Classification: O14, O25, O53, P11

* LSE, CESifo, Email: N.Lane@lse.ac.uk

† University of Oslo, Email: max.marczinek@econ.uio.no. The author gratefully acknowledges funding by Norwegian Research Council project number 359741.

‡ Oxford, Email: dilan.yang@nuffield.ox.ac.uk

1 Introduction

In the decades following World War II, Japan, Taiwan, South Korea, and Singapore saw rapid gains in industrial production, exports, and living standards. Few questions in growth and development have generated more persistent disagreement than the sources of East Asia's ascent (Rodrik, 1994). Classic accounts place the developmental state—in particular, its selective interventions to accelerate industrialisation—at the centre of this transformation (Johnson, 1982; Amsden, 1989; Wade, 1990); for a recent restatement, see Cherif and Hasanov (2019). Other interpretations emphasise macroeconomic stability (World Bank, 1993), human-capital accumulation (Young, 1995), and favourable geopolitical conditions (Woo, 1991; Irwin, 2021a). That these economies became trade powerhouses is not in dispute, but the roles of exports and trade policy, including the orientation of trade policy, remain unsettled (Balassa, 1978; Westphal, 1978; Lall, 1997; Frankel and Romer, 1999; Lawrence and Weinstein, 2001). Industrial policy has returned to prominence (Fernandes and Reed, 2026), yet its contribution to the East Asian miracle remains contested.

This paper revisits the relationship between industrial policy and growth in East Asia. We focus on export-oriented industrial policy (EOIP), a form of industrial policy that, we argue, broadly typified industrial strategy across the region. EOIP departed from the import-substitution strategies that dominated postwar development economics.[1] The shift to EOIP spread across the region in stages. Japan was the pioneer. Taiwan and South Korea adopted, and in part domesticated, similar strategies in their emerging postwar economies. Singapore, a newly independent entrepôt economy, came next, followed by the rest of Southeast Asia.

We study the relationship between these shifts to export-oriented industrial policy and aggregate development in Japan, Taiwan, South Korea, and Singapore, which we call the "early adopters".[2] Drawing on the comparative historical literature, we document and date each transition to EOIP. We use differences in the timing of these transitions to study economic outcomes before and after each policy turn. Our synthetic difference-in-differences estimator compares each treated country cohort with a counterfactual trajectory constructed from donor economies outside East and Southeast Asia. We estimate large positive effects of the turn on GDP and exports per capita, alongside parallel movements in factor accumulation.

Debates over East Asia's growth miracle face two related problems: how to characterise industrial policy and how to estimate its effects. Conceptually, scholars disagree over whether East Asia's reforms constituted industrial policy at all and, if so, what kind. Studies of South Korea's industrial policy, for example, have variously characterised the same episodes as liberalisation (Connolly and Yi, 2015; Billmeier and Nannicini, 2013),

[1] Countries pursued export-promoting industrial policies to varying degrees, but import substitution was arguably the dominant strategy of the period, including in many East Asian economies.

[2] We exclude Hong Kong, which was administered directly by the United Kingdom as a Crown colony and famously eschewed such interventions. Although largely absent from studies of industrial policy, Hong Kong illustrates one of our conclusions: policy experience across East Asia was heterogeneous.

export pushes (Westphal, 1990), and sectoral dirigisme (Amsden, 1989).[3] Empirically, the sources of East Asia's growth remain contested.[4] The small-*N* setting, data limitations, and difficulties in operationalising treatment constrain empirical work (Rodrik, 2012; Lane, 2020). The question "what is the counterfactual?" permeates the literature.

We make an incremental contribution to these debates. We identify EOIP as a shared feature of postwar East Asian industrial strategy. Broadly, we understand EOIP as a diverse class of industrial-policy regimes centred on export performance and entry into external markets. A particularly important feature was sector- or activity-specific interventions that conditioned access to policy incentives (subsidies, preferential finance, tax treatment, imported inputs, or other forms of support) on export performance. We distinguish industrial policy from broad macroeconomic reforms *per se*, following the literature (Juhász, Lane, and Rodrik, 2024; Juhász, Lane, Oehlsen, and Perez, 2025; Fernandes and Reed, 2026). Such reforms may promote exports, but they aim at macroeconomic objectives rather than at the composition of economic activity.[5] Nor is export promotion equivalent to laissez-faire trade policy. As the East Asian cases show, export-promotion policies may simultaneously expose firms to competition in external markets and involve deliberate state intervention.

We examine the transition to EOIP among the first wave of East Asian economies. In many cases, this meant moving away from import-substitution industrialisation, which had dominated development strategy. These transitions shifted the weight of policy support from broad import protection towards incentives contingent on a specific activity: exporting. Each turn corresponds to the adoption of a policy package, or bundle: export-conditioned incentives formed its industrial-policy core, supported by complementary measures, such as currency devaluations, that are not industrial policy on their own (Section 2). EOIP could also coexist with protectionism (Cherif and Hasanov, 2024). Both the early adopters and the four later adopters combined export-focused interventions with varying degrees of domestic market protection.

We first reconstruct the timing of the turn towards export orientation. We define this turn as the point at which a country's industrial-policy regime shifted towards EOIP (instruments conditioned on export performance). Typically, this shift involved moving away from explicitly import-substituting trade regimes. Following Irwin (2021a) and Yoo (2017), we date each turn by policy decisions rather than economic outcomes. To identify these decisions, we construct a new database from primary policy documents and academic accounts. We code policies by hand, then check each entry against its sources using a purpose-built retrieval-augmented generation system powered by a large language model. Each observation is a policy-sector (three-digit industry, or ISIC) pair. Our database covers

[3] Even South Korea's later heavy and chemical industry drive has been characterised by unprecedented import barriers (Lall, 1997), while other work highlights subsidies for imported intermediates (Lane, 2025).

[4] These include debates in growth accounting (e.g., Young, 1992; Hsieh, 2002) and in regression analysis (e.g., Easterly, 1995; Rodríguez and Rodrik, 2001; Rodrik, 2012).

[5] For instance, although exchange-rate devaluations encourage exports and were essential components of EOIP regimes (Juhász and Lane, 2024), they do not, on their own, constitute industrial policy or EOIP. They belong instead to the policy package adopted at each turn (Section 2).

the four early adopters and the four comparison economies (China, India, Malaysia, and Thailand), whose transitions we also date.[6]

We then study the relationship between the export-orientation turn and country-level economic development using a synthetic difference-in-differences (SDID) estimator (Arkhangelsky et al., 2021). SDID combines features of synthetic control methods (SCM) and difference-in-differences, offering advantages in our empirical setting. Although SCM is a powerful tool for case studies with few treated units, its "contextual requirements" (Abadie, 2021) can be difficult to meet in postwar East Asia. In particular, SCM benefits from longer pre-treatment periods, which are often unavailable for newly independent postwar economies. Short pre-treatment periods raise the risk of overfitting, while turbulent outcomes among treated economies increase the risk of fitting noise. Constructing a counterfactual also becomes difficult when these economies are unusual in their pre-treatment outcomes or covariates, or in their combination.[7] SDID relaxes some of these constraints, for example by allowing persistent level differences.

We apply SDID in our staggered-adoption setting to study the differential evolution of the four early adopters relative to a set of donor economies.[8] We use SDID to estimate how output, factors, and productivity respond to the export-promotion turn. We estimate a separate effect for each treated economy (cohort), using country-specific donor pools and observation windows, and then pool the estimates. For each outcome, we report the average cohort effect alongside its horizon profile at 10, 15, and 20 years after adoption.

We then ask: how much of the estimated output response to the policy turn is accounted for by the corresponding estimated employment and capital responses? We use our SDID estimates for output and factors in an accounting exercise that separates the part of estimated growth accounted for by measured factor inputs from the part left to a residual (which includes forces such as TFP). In doing so, we provide a simple bridge between modern econometric estimands and growth accounting.

We highlight three findings. First, the estimated effect of the export-orientation turn on GDP per capita is large and positive. The average cohort effect is 0.449 log points, and all four cohort effects are positive. The horizon profile shows the effect accumulating over time rather than as a one-time level shift immediately after adoption. Second, the effect on exports per capita is larger than that on GDP per capita, both on average and at every reported horizon. The average cohort effect on exports per capita is 1.356 log points. The larger estimated proportional export response is consistent with the policy's export orientation. Importantly, the signs of the estimates do not rest on any one economy. All four cohort effects on exports per capita are also positive, and every leave-one-out average retains its positive sign for both outcomes (Section 10).

Third, we find significant movements in traditional factor inputs—employment and capital—and *TFP*, and find that each potentially plays a sizable role in output growth. We fail to detect effects on human capital using common indices. Our baseline calculations for GDP per capita find that employment accounts for approximately 29%, capital deepening

[6] It comprises 922 observations across 104 industries between 1949 and 1990.

[7] These concerns are not theoretical; they recur in studies of liberalisation and macroeconomic reforms (e.g., Billmeier and Nannicini, 2013).

[8] We exclude neighbouring Asian economies, such as Indonesia and the former South Vietnam, to limit spillovers.

30%, and the residual 41%. Although employment and capital together account for a majority of the per-capita response, they still leave a sizable role for the residual. Our TFP estimates track the magnitude of the residual.

We contribute to three related literatures. First, we revisit empirical work on the sources of East Asian economic growth. This literature debates the role of policy interventions (World Bank, 1993; Amsden, 1989; Wade, 1990; Rodrik, 1995) and the contributions of factor accumulation and productivity (Young, 1995; Hsieh, 2002). Building on early descriptive cross-country studies, we combine historically grounded treatment dates with econometric counterfactual analysis (Lane, 2020). By estimating positive effects of turns to EOIP-centred policy packages on economic growth, we challenge claims that East Asian economies grew "despite industrial policy".

Our results also suggest that changes in factor accumulation and total factor productivity (TFP) may be downstream responses to export promotion. We thus speak to the longstanding debate over factor accumulation versus productivity as sources of East Asian growth.[9]

Second, we contribute to the literature on industrial policy in development, growth, and structural change (Juhász, Lane, and Rodrik, 2024; Fernandes and Reed, 2026; European Bank for Reconstruction and Development, 2024). We engage particularly with work on the role of export activity and industrial policy (Reed, 2024; Cherif and Hasanov, 2024). Rather than classifying countries using aggregate openness indicators or relying on a few canonical programmes, we code individual policies at the sector level in a new database and trace changes in their orientation over time. The data document both the shift towards export-promoting policies and the heterogeneity of policy packages. They show why EOIP should not be equated with general liberalisation: export-promoting measures frequently operated alongside selective protection and other forms of domestic intervention.

Third, we contribute to the literature on trade policy, export policy, and growth (Wacziarg and Welch, 2008; Connolly and Yi, 2015; Irwin, 2021c; Cherif and Hasanov, 2024) by dating the turn from government policy decisions rather than export outcomes or aggregate openness. We evaluate the policy turn in historical cross-country settings using SDID with one cohort per estimate. This allows each comparison window to begin at the corresponding historical policy turn rather than imposing a common start date on all four early adopters. We measure each economy over the same span—the first two decades after its turn—and count each once. The horizon profile traces how the estimated effect accumulates from one decade to two (Section 6.2). This design yields statistically significant estimates of the effects of the turn on GDP and exports per capita.

The remainder of the paper proceeds as follows. Section 2 defines industrial policy and EOIP. Section 3 explains the dating of the four export-orientation turns. Section 4 introduces the historical policy corpus and macroeconomic data. Section 5 uses the policy database to document how the level and composition of intervention changed around those dates. Section 6 sets out the estimator, the estimand, the identifying assumption, and placebo inference. Section 7 reports the effects on GDP per capita and exports per capita. Section 8

[9] The debate is unresolved. For instance, Young (1995) emphasises factor accumulation using primal growth accounting, whereas Hsieh (2002) uses dual accounting methods and estimates higher productivity growth for Taiwan and Singapore than primal estimates imply. Cherif and Hasanov (2019) also dispute the factor-accumulation reading, pointing to sustained TFP growth.

estimates the effects on factor inputs and TFP. Section 9 then accounts for the output response, using an accounting exercise combining the estimates from Sections 7 and 8. Section 10 presents the robustness exercises. Section 11 concludes.

2 Concepts and Definitions

Industrial Policy We define industrial policy (IP), most broadly, as state policy that intentionally seeks to change the structure of economic activity (Juhász, Lane, and Rodrik, 2024). To qualify as IP, a policy must meet two requirements: intentionality, reflecting a deliberate strategy, and specificity, targeting a particular sector or activity.

Industrial policy has familiar theoretical rationales, including the correction of market failures (Juhász, Lane, and Rodrik, 2024).[10] Critiques of industrial policy are nevertheless numerous, many arising from debates over the East Asian growth miracle (Pack, 2000; Pack and Saggi, 2006). A theoretical rationale for IP does not, however, resolve the distinct problems of implementation.

Export-Oriented Industrial Policy as a Special Case We treat export-oriented industrial policy (EOIP) as a special case of IP that targets export expansion. Intentionality excludes policies directly imposed on a government by a foreign power (Clifford, 1994; Woo, 1991), but not a government's strategic response to external pressure or national challenges. Such policies remain the government's own choice. The desired outcomes of EOIP are commonly expressed as growth targets for individual sectors or the whole economy, most strikingly in Japan's Income Doubling Plan of 1960 (Shinohara, 1964). Common instruments include export targets, export subsidies, trade finance, and preferential schemes tied to exporting.

This list is not exhaustive. EOIP encompasses diverse interventions aimed at export expansion, with access to policy support often conditioned on export performance. Such interventions frequently require complementary measures and form part of broader policy packages, including other forms of IP and policies outside IP. This was particularly true in East Asia, as we discuss below.

East Asia and the Export Turn The export-orientation turn marks the transition towards industrial strategies centred on export development. We argue that this policy regime broadly typified East Asian industrial strategy (Westphal, 1978; Westphal, 1990; Lane, 2025). EOIP prioritised export activity and made export performance, on both the extensive and intensive margins, central to policy (Balassa, 1982). In practice, this often meant conditioning incentives and other forms of support on export outcomes. The turn was often, but not always, a transition from import substitution (Irwin, 2021c).

[10] Policymakers' motives extend beyond correcting market failures to include national security, self-sufficiency in critical industries, and geopolitical power.

While we take EOIP as the defining industrial-policy orientation of the East Asian model of industrial development, it took various forms and was embedded in broader policy packages. Complementary macroeconomic policies, such as currency devaluations, also supported the turn. The early adopters combined EOIP with a diverse range of other institutions and policy measures.

We therefore treat each turn as the adoption of a policy package, or bundle. Export-conditioned incentives form its industrial-policy core. The complementary measures adopted with them, such as devaluations, are not industrial policy on their own, but they belong to the package. Thus, our estimates capture the package.

3 Dating the Export-Orientation Turn

We date the export-orientation turn to the year in which a government chooses to pursue export-oriented industrial policy, often following a period of import substitution. The principal criterion is the decision to adopt export-focused measures, such as export-contingent policies. We use this ex ante policy decision rather than subsequent export or growth responses (ex post criteria). We adjudicate country-level dates using modal dates drawn from expert assessments in primary and secondary scholarship on these regime changes.

The turn may involve policy changes in only a few sectors, so import-substituting and export-oriented policies can coexist after it. We identify two typical patterns. The first, exemplified by South Korea (Yoo, 2017; Kim, 2011), may be called "trial and error". Experimentation or a policy's unforeseen success can lead governments to embrace export-oriented policies fully, culminating in the export-orientation turn. The second pattern is policy copying: some East Asian policymakers looked to Japan and adopted export-oriented industrial policies without a trial. In these cases, initial adoption marks the policy turn. We explain our dating decisions for each early adopter.

Japan We divide Japan's economic trajectory into three phases. The initial phase of reconstruction was followed by a second phase, which began with the election of the first Liberal Democratic Party (LDP) government in 1955 and the promulgation of the First Five-Year Plan (Yamada, 1956). The Income Doubling Plan of 1960 marks the third phase (Shinohara, 1964). The idea of a Japanese "economic miracle" was already being discussed in 1962 (Johnson, 1982), leaving a 17-year window in which to date the policy turn.

The first major export-related policy in postwar Japan was the Dodge Line of 1949, which fixed the yen-to-dollar exchange rate (Johnson, 1982). The creation of the Ministry of International Trade and Industry (MITI) that same year, together with its intent to promote industrial exports (Johnson, 1982), makes 1949 a candidate for the policy turn. The case for 1955 rests on the first LDP cabinet and the draft of the First Five-Year Plan, which included explicit export growth targets and recognised that a "rapid expansion of export[s] is necessary and the competitive power of [exports] must be strengthened by means of industrial rationalization" (Yamada, 1956). The plan set a target of increasing exports by 166 per cent by 1960 relative to 1954 levels. While the Income Doubling Plan of 1960 may have

been the boldest initiative, it came too late: high export targets were a feature of the 1955 plan, and the underlying institutions, notably MITI, were already in place. Following Johnson (1982), we argue that although the institutions of high-speed growth and export orientation began to be forged in 1949, they only became fully operative in 1955. We therefore code 1955 as the policy turn.

South Korea South Korean policy in the 1950s was marked by import substitution under Syngman Rhee (Seth, 2017) and limited trade liberalisation under an interim government (Haggard, Kim, and Moon, 1991). Park Chung Hee came to power through a coup in 1961 and shifted policy towards export orientation. In the mid-1960s, the government revised the First Five-Year Plan and implemented export-promotion policies more rigorously (Yoo, 2017). Devaluations followed in 1964. In 1967, the government adopted the Second Five-Year Plan, maintaining a focus on expanding industrial capacity and exports. From 1971, more targeted measures, such as the heavy and chemical industry drive, replaced general export-promotion policies. This set of policies ended with Park Chung Hee's assassination in 1979.

We date the turn to the point at which the Park government committed its planning apparatus to export orientation. The First Five-Year Plan of 1962 drew on the 1961 Five-Year Plan of Chang Myon's interim government (Wolf, 1962) and still reflected earlier import substitution. The 1961 devaluation was also largely a legacy policy. In 1962, however, Park requested a revision of the First Five-Year Plan that laid the groundwork for export orientation. The same year, the government instituted full-scale annual export targets: the Ministry of Commerce and Industry set them by commodity, region, and destination, tracked performance against them in an "export situation room", and reviewed them at an Expanded Export Conference chaired by the President (Frank, Kim, and Westphal, 1975, pp. 46–47; see also Westphal, 1978; Westphal, 1990). Export subsidies and tax relief for exporters had begun in 1961 (Frank, Kim, and Westphal, 1975, p. 46). The government built on emerging export successes to develop an increasingly explicit export-oriented growth policy. Other accounts place the turn in the early to mid-1960s (Kim, 2004; Kim, 2011; Irwin, 2021a). Dating it to the mid-1960s, however, would credit the outcome rather than the decision that set it in motion. We therefore code 1962, when export targets entered the planning apparatus, as the policy turn.

Taiwan While Taiwan was a largely agricultural economy at the onset of industrialisation, it stood as the most industrially advanced Chinese province in 1945 (Wade, 1990). Taiwan's postwar experience started with the retreat of the Chinese Nationalist Party (Kuomintang) after its defeat by the Communist Party in the Chinese Civil War in 1949. Taiwan adopted import-substitution policies in the early 1950s. Measures to promote exports began to take shape in the mid-1950s. In 1954, the government introduced import-duty and commodity-tax rebates for raw materials used in export production. In 1957, the government set up a special export loan programme to lower the cost of funds for exporters (Rodrik, 1995). In 1958, the government adopted an export-promotion package that devalued the currency by 25 per cent, unified exchange rates, introduced export incentives, and removed quantitative import restrictions (World Bank, 1993). The change in policy was accompanied by an export increase of 50 per cent in 1959 (Irwin, 2021a).

Although export promotion began in the mid-1950s, these early measures were limited in scale (Kuo and Ranis, 1982). The 1958 reforms, by contrast, combined export incentives with exchange-rate reform and import liberalisation, marking a broader policy commitment to export orientation. We therefore date the policy turn to 1958 (Tsai, 1999; Ho, 1987; Vogel, 1990).

Singapore Singapore passed the Industrial Expansion Ordinance in 1959, allowing the government to grant tax holidays to industrial firms (Lee, 1973; van Elkan, 1995). In 1961, the government established the Economic Development Board, which would become the principal agency responsible for industrial policy (Lim, 1995; Lee, 1973). Singapore began raising tariffs in 1962; by 1965, some 230 commodities were subject to import licensing and quantitative restrictions (van Elkan, 1995). In 1963, Singapore joined the newly formed Malaysia, from which it was expelled in 1965 (Suryanarayan, 1969). Britain's announcement in 1967 that it would withdraw its troops from the region by 1971 further increased the urgency of shifting from import substitution to export orientation (Soon and Stoever, 1996). That same year, Singapore passed the Economic Expansion Incentives Act, granting tax breaks and promotional benefits to firms in pioneer industries, particularly those with export potential (van Elkan, 1995).

Dating Singapore's policy turn is relatively clear-cut. Scholars broadly agree that Singapore's expulsion from Malaysia and the announced withdrawal of British troops compelled the government to change course. These developments prompted a shift towards labour-intensive, export-oriented industrialisation (Lim, 1995). We therefore date the policy turn to 1967, when the Economic Expansion Incentives Act extended industrial incentives to firms with export potential, marking a more explicit commitment to export orientation.

Table 1 presents these adoption dates alongside the estimation sample periods for each cohort. Each sample begins before the corresponding adoption year; the start years mark the estimation windows, not alternative policy dates.

Four Comparison Economies Online Appendix OA-A discusses the dating of the export-orientation turn in four comparison economies. We date these transitions to 1968 for Malaysia, 1972 for Thailand, 1978 for China, and 1991 for India.

4 Data

4.1 Industrial policy corpus data

We focus on four early adopters of export-oriented industrial policy: Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). For our policy dating analyses, we also code policies for four later adopters: Malaysia (1968), Thailand (1972), China (1978), and India (1991). See Section 3. These comparison economies are not the subject of our main empirical analysis, nor are they an exhaustive list of later adopters.

We build a database of industrial policies from a corpus of primary policy documents and academic accounts, in two stages. In the first stage, we code policies from the corpus by hand, following a documented coding scheme. Each row is a policy-sector pair (ISIC-coded) recording the policy lever, the targeting metric, and one or more direction flags: export-promoting (EXP), import-limiting (IMP), domestic (DOM), or export-banning (EXPBAN). Table A.2 describes the resulting database. The policy data cover 1949 to 1990; coverage is not uniform, with the earliest years available only for a subset of countries.

In the second stage, we test the hand-coded database for two kinds of error with a retrieval-augmented language model built on the same corpus. To detect misclassification, the system retrieves for each coded policy row the passages most relevant to the policy and its sector, and the model judges each direction flag against them. To detect omissions, we screen every passage that resembles no coded policy for measures the hand coding missed. Appendix A.2.4 describes the corpus, the retrieval system, and the review. The verification largely confirms the hand-coded record, with some corrections and additions.

4.2 Macroeconomic data

The estimation samples differ by cohort and together span 1949 to 1987, well before the Asian financial crisis and China's accession to the WTO. Each sample ends 20 years after adoption; Table 1 gives the adoption dates and sample coverage by cohort.

We construct a country-year panel of macroeconomic data from five sources: the Maddison Project Database (Bolt and van Zanden, 2024), the Penn World Table (PWT) (Feenstra, Inklaar, and Timmer, 2015), TRADHIST (Fouquin and Hugot, 2016), the Barro-Lee educational attainment dataset (Barro and Lee, 2013), and the Global Macro Database (Müller, Xu, Lehbib, and Chen, 2025). A country enters the merged panel only if it has usable coverage in at least the Maddison and TRADHIST sources. We drop countries present in only one source or with no trade coverage. Appendix A.1 details the outcome variables, predictors, missing data, and sample restrictions. Table A.1 summarises the resulting variables and their sources.

5 Disaggregating the Miracle in Policy Data

Using our database of industrial policies coded at the policy-sector level across eight economies from 1949 to 1990, we trace changes in the policy mix during the East Asian miracle. We focus on the turn to export-oriented industrial policy, which we identify as the industrial-policy core of the East Asian model. We show that, in Japan, Taiwan, and South Korea, the historically dated turns coincide with increased adoption of export-promoting policies. This corroborates our treatment dates with detailed policy data.

We make four choices in constructing the comparisons that follow. First, we express time relative to each country's treatment year, displaying 12 years before and after treatment. Second, rather than report raw policy counts, we report cumulative shares of policies implemented or withdrawn by each year, using 1990 as the endpoint. These series need not

begin at zero or end at one. Third, we restrict attention to trade-related policy codes, since larger economies such as Japan employ relatively more domestic policies unrelated to trade. Finally, for this exercise only, we include Malaysia, Thailand, China, and India as comparison economies, using the historically derived treatment years from Section 3. Although we do not construct a counterfactual in this analysis, the descriptive comparison allows us to examine whether these economies also shifted their industrial policies in comparable ways around their respective turns to export orientation.

Figure 1 shows a common pattern in Japan, Taiwan, and South Korea: the adoption of export-promoting policies increases around their historically dated turns to export orientation. The pattern is weakest in Japan, where only 7 of its 163 coded rows are export-promoting. Singapore's turn is not visible in our coded data, which record no export-promoting policies for the country. Its treatment date therefore rests on the historical record.

Our data capture policy adoptions and withdrawals, but not the stock of institutions through which industrial policy operates. They therefore cannot measure established export-contingent institutions, such as export-import banks and trading companies, whose coordinating role in export-led industrial policy Reed (2024) emphasises. Examples include the Export-Import Bank of Japan (JEXIM) and Singapore's International Trading Company (INTRACO).

Unlike export promotion, import policy shows no uniform pattern around the turn. This is consistent with accounts that describe the East Asian miracles as combining export promotion with tariff protection (Cherif and Hasanov, 2024). In Taiwan and South Korea, withdrawals of import-limiting policies also increase around the export-orientation turn. In Singapore, by contrast, all import-limiting policies are withdrawn before treatment. Japan, the earliest adopter, withdraws none before treatment and liberalises import protection later. For South Korea, these patterns are consistent with the historical record, which documents a gradual decline in import restrictions during the Park-era export boom (Lane, 2025).

Figure A.1 shows a similar rise in export-promoting policy adoption in Malaysia, China, and Thailand. India's 1991 turn falls outside the study period (1949–1990), so we exclude it from this analysis. Taken together, the historically dated turns coincide with substantial policy adoption and withdrawal, with a pronounced shift towards export-promoting policies and a weaker tendency towards the withdrawal of import-limiting policies.

Beyond policy adoption and withdrawal, Figure A.2 examines how the composition of active policies changes around the turn. In Japan, Taiwan, and South Korea, the policy mix shifts markedly away from purely import-limiting policies and towards policies combining import limits with export promotion, as well as purely export-promoting policies. Singapore's turn remains absent from our coded data. In China, Malaysia, and Thailand, the share of sectors under partial or pure export promotion also rises, but the share under purely import-limiting policies shows no clear downward trend. This is consistent with Malaysia's combination of import substitution in automobiles and export orientation in electronics (Cherif and Hasanov, 2024).

Additional descriptive evidence appears in Appendix A. Figures A.3 and A.4 show the frequency of policies over time by country and type, while Figure A.5 presents the raw policy coding database.

The policy data yield three main findings. First, the historically dated turns in Japan, Taiwan, and South Korea coincide with increased adoption of export-promoting policies,

corroborating the historical dating with systematic evidence. This complements the wider historical literature, which has relied largely on qualitative assessments of the policy record. Singapore's treatment date, by contrast, rests on historical evidence that our coding does not capture. Second, Malaysia, China, and Thailand also show a rise in export promotion around their respective turns. Third, the early adopters show varying degrees of movement away from import protection, whereas this tendency is less apparent among the comparison economies. We now turn to the empirical analysis using our historically dated treatment variable.

6 Empirical Strategy

We estimate the effects of East Asia's export-orientation turns on aggregate development using synthetic difference-in-differences. Our four early adopters are thus referred to as our treated economies. Using the dated turns towards EOIP, we construct a separate weighted comparison for each treated economy based on its pre-adoption growth path. We report the average cohort effect across the four treated economies and a horizon profile showing how effects evolve over time. We describe the choice and implementation of SDID in Section 6.1, define the estimands in Section 6.2, discuss the identification conditions and limitations in Section 6.3, and explain placebo inference in Section 6.4.

6.1 Estimation strategy

Our empirical setting motivates our choice of estimator. We consider a small number of economies that turned towards export orientation, each at a different date. These four treated economies lie near the lower end of the pre-adoption outcome distribution, with outcome levels substantially below those of otherwise plausible donor economies. Their usable pre-adoption histories are also short, reflecting their nascent postwar political development.

These features complicate the construction of credible counterfactuals. Conventional difference-in-differences (DiD) is less appealing when an unweighted donor average is unlikely to follow the treated economies' pre-adoption trajectories. Synthetic control methods address differences in trajectories by constructing a weighted combination of donor economies. However, their requirement to match both outcome levels and trajectories can make a satisfactory fit infeasible. They may also concentrate weight on a few donors.[11]

[11] For the comparison with DiD and synthetic control, see Arkhangelsky et al. (2021) and Abadie, Diamond, and Hainmueller (2010). Factor-model alternatives, generalised synthetic control (Xu, 2017) and matrix completion (Athey et al., 2021), estimate the factor structure of untreated outcomes directly, which needs more pre-adoption years than the six or seven available here. As a robustness check, we implement synthetic control in its augmented form (ASCM) (Ben-Michael, Feller, and Rothstein, 2021).

Synthetic difference-in-differences separates the level from the path (Arkhangelsky et al., 2021) and requires neither parallel trends in the unweighted donor average nor level matching. Instead, it weights donors so that their combination tracks the treated economy's growth path. Because it weights both years and donors, SDID has a double robustness property and generally requires fewer pre-adoption years than synthetic control (Clarke, Pailañir, Athey, and Imbens, 2024).

We estimate a separate SDID effect for each treated economy, using its own donor pool and estimation window. Donor weights reproduce the treated economy's pre-adoption trajectory while allowing for a persistent difference in outcome levels. Time weights are fitted on donors alone, so that weighted pre-adoption outcomes approximate the donors' post-adoption averages; the same weights define the treated economy's baseline. A weighted regression with economy and year fixed effects then compares the treated economy's change from this baseline to its post-adoption average with the corresponding change among its weighted donors. Appendix B.1 provides the specification and explains its relationship to the standard staggered-adoption implementation.

Each cohort's sample start is set by historical circumstances and available data, and the sample ends 20 years after adoption. Table 1 reports the dates, while Appendix A.1 describes donor eligibility and Appendix B.3 reports donor-pool sizes.

Our headline specification adjusts outcomes for three covariates: logged pre-adoption population, the logged human-capital index, and inflation. We interact each covariate's pre-adoption mean with a linear time trend and remove its fitted contribution before constructing the SDID weights, following Kranz (2022) and Clarke, Pailañir, Athey, and Imbens (2024). Appendix B.2 gives the specification.

6.2 Estimand

Our main estimand, $\bar{\tau}$, is the *average cohort effect* across Japan, Taiwan, South Korea, and Singapore. For each economy c, with adoption year a_c and event time $h = t - a_c$ (t is the calendar year), the *cohort effect*, τ_c, is the average effect over the adoption year and the following 20 years. Because each turn corresponds to the adoption of a policy package (Section 2), τ_c is the effect of that package as a whole. We then average the four cohort effects with equal weights.[12] We report this summary for every outcome, including our two core outcomes, GDP per capita and exports per capita.

Equal weights give each economy's policy turn the same importance, while the common 21-year period makes the effects comparable across cohorts with different adoption dates.[13]

To show how effects evolve, we also report a *horizon profile*. Here, $\tau_c(h)$ is the effect for economy c at h years after adoption, and $\bar{\tau}(h)$ is its average across the four economies. We report this average at 10, 15, and 20 years: each point averages event-year effects from the

[12] Hats indicate estimates of the corresponding estimands: $\hat{\tau}_c(h)$, $\hat{\tau}_c$, $\hat{\bar{\tau}}$, and $\hat{\bar{\tau}}(h)$ estimate $\tau_c(h)$, τ_c, $\bar{\tau}$, and $\bar{\tau}(h)$, respectively.

[13] The economy-year weighted average treatment effect on the treated, the default aggregation in common implementations (e.g., codesynthdid R package) equals the average cohort effect here because all cohorts contribute the same number of treated years (Callaway and Sant'Anna, 2021).

same full-window fits, rather than estimates from separately refitted windows or cumulative averages. By contrast, $\bar{\tau}$ summarises the full period. The estimates are the following

$$\hat{\tau}_c = \tfrac{1}{21}\sum_{h=0}^{20}\hat{\tau}_c(h), \qquad \hat{\bar{\tau}} = \tfrac{1}{4}\sum_{c}\hat{\tau}_c, \qquad \hat{\bar{\tau}}(h) = \tfrac{1}{4}\sum_{c}\hat{\tau}_c(h) \tag{1}$$

where $h \in \{10, 15, 20\}$ is the number of years after adoption. Appendix B.1 defines the annual effect estimates underlying these summaries. We report the individual cohort effects alongside the overall average in Section 7 to show each economy's contribution.

6.3 Identification

Our identification rests on the assumption that, without the policy turn, each treated economy would have followed the path of the counterfactual, or the weighted donors. After covariate adjustment, the gap measured before adoption must therefore remain stable in the absence of treatment. Under this assumption, subsequent changes in that gap identify the effect of the turn. We assess the plausibility of this assumption through pre-adoption fit and historical evidence on the policy turns.

Pre-adoption fit. The first question is whether the weighted donors reproduce the treated economies' trajectories before adoption. After allowing for differences in outcome levels, the weighted donors closely track those trajectories (Table B.1). For exports, the treated economies show some upward movement relative to their weighted donors in the three years before adoption. Trade volumes fluctuate more than GDP in response to shocks (Engel and Wang, 2011), which may help explain this weaker match. The close fit for GDP supports the credibility of its counterfactual.

Selection into idiosyncratic gain. Even a close pre-adoption fit leaves open the possibility that governments changed policy regime because their economies were already poised for faster growth. The historical evidence helps assess this concern about adoption timing. Prospective reductions in US assistance encouraged reform in Taiwan, while declining aid and US policy pressure contributed to South Korea's transition towards export promotion (Irwin, 2021b; Irwin, 2021a). In Singapore, separation from Malaysia and Britain's announced military withdrawal increased the urgency of adopting an export-oriented strategy (Lim, 1995; Soon and Stoever, 1996). These episodes suggest that external pressures helped determine when governments changed policy, alongside domestic economic conditions (Section 3). They provide some evidence against the interpretation that adoption simply marked the attainment of the capabilities needed for faster growth. They may also ease the concern about anticipation effects, as reforms prompted by external events, such as aid reductions or Singapore's separation from Malaysia, were harder for domestic actors to foresee than reforms planned internally.

No interference. The credibility of the weighted donor comparison also rests on donor economies remaining unaffected by the treatment or similar interventions. We therefore exclude East and Southeast Asian economies, because regional trade, investment, and policy diffusion could transmit the effects of the policy turns to them (Appendix A.1).

Contemporaneous events. Identifying the effects of the turn also requires excluding other developments as explanations for the subsequent changes in outcomes. The historical discussion in Section 3 documents the measures surrounding each turn and helps distinguish components of the export-oriented policy package from separate reforms. Measures that implement the policy turn form part of the treatment being evaluated, whereas independent reforms and macroeconomic shocks remain potential competing explanations.

Dynamics. Finally, earlier horizons allow less time for subsequent policies and shocks to influence the comparison, although they remain exposed to events coinciding with the turn. Later horizons capture effects that may emerge gradually through investment and structural change, while requiring the counterfactual relationship to remain credible for longer. We therefore report the dynamic path alongside the average cohort effect (Figure 2 and table 2). Reading the two together shows how the estimated effects develop over time, though longer-horizon estimates are more exposed to competing explanations from unrelated policies and shocks.

6.4 Inference

We conduct inference by comparing the estimated policy effects with placebo effects in donor economies (i.e., placebo inference). For each placebo, we assign a treated economy's adoption year to an eligible donor and repeat the full estimation, including covariate adjustment and the donor and time weights. We exclude the treated economies from these comparisons. This procedure shows how large an estimated effect can arise among donors without the policy turn. It adapts the staggered placebo procedure of Clarke, Pailañir, Athey, and Imbens (2024) to cohort-specific windows and donor pools. The placebo standard error is the standard deviation of the estimates across placebo assignments (Arkhangelsky et al., 2021, Algorithm 4).[14]

For each cohort, we run every eligible donor once as the placebo against the rest of the cohort's donor pool. We compute the placebo standard error across these n donors and compare the observed cohort estimate $\hat{\tau}_c$ with their placebo effects. Following Abadie, Diamond, and Hainmueller (2010) and Abadie, Diamond, and Hainmueller (2015), the two-sided rank p-value of $\hat{\tau}_c$ is $(1 + k)/(1 + n)$, where k counts placebo effects at least as large as $\hat{\tau}_c$ in absolute value, ties included. The smallest attainable value is $1/(n + 1)$. Because this reference distribution uses every eligible donor, it has no simulation error. See Appendix C.1 for details and Table C.3 for the rank and pool.

[14] The jackknife is undefined with one treated economy per cohort. We reject the bootstrap because resampling countries leaves at least one of the four treated economies out of most resamples. Most bootstrap estimates therefore average over fewer than four cohorts and describe a different quantity.

For aggregate estimates, which combine economies, each placebo assignment allocates the four adoption years to four distinct eligible donors. We estimate each placebo cohort and combine the results using the same procedure as for the observed effects. Enumerating every such assignment is infeasible, so we draw 1,000 assignments at random. A pooled p-value therefore carries simulation error: a different random draw would move it slightly. Table C.1 reports its range. The two-sided Monte Carlo p-value is $(1 + k)/(1 + R)$, where R is the number of seeded assignments per outcome and k counts placebo statistics at least as large as the observed statistic in absolute value, ties included (Phipson and Smyth, 2010). The smallest attainable value is $1/(1 + R)$, which prints as 0.001, a floor, not zero. Leave-one-out reuses the same assignments (Section 10).

We follow the literature (Arkhangelsky et al., 2021; Clarke, Pailañir, Athey, and Imbens, 2024) and, for reference, report approximate 95 per cent Gaussian confidence intervals: the estimate plus or minus 1.96 placebo standard errors. The figures' each have pointwise bands, drawn for each event year. Our claims rest on the p-values, while we use bands for visual guidance. For p-values, we use the five per cent level ($p < 0.05$). For core outcomes, our average cohort effects use Holm-adjusted p-values (Table C.2). All other p-values, such as those of the horizon profile and the cohort effects, support secondary claims and are unadjusted.[15]

Every standard error, interval, and p-value rests on the equal-noise assumption that treated economies are as noisy as their donors. If the treated economies are noisier, uncertainty is understated (Arkhangelsky et al., 2021; Conley and Taber, 2011; Ferman and Pinto, 2019). The pooled statistics combine four neighbouring economies, the case in which such tests can over-reject. If the four economies' estimation errors are more correlated than those of randomly drawn donors, the pooled standard errors are too small.[16] The treated economies were not chosen at random, so this is placebo and not randomisation inference. No p-value is exact.

7 Results: GDP and Exports

We focus first on the average effects of the turn on the two core outcomes. The turn is the adoption of a policy package with EOIP at its core (Section 2). We estimate average cohort effects of 0.449 log points for GDP per capita (equivalent to 57 per cent) and 1.356 log points for exports per capita (equivalent to 288 per cent). The average cohort effect covers each cohort's adoption year and the 20 years that follow. We report these estimates beside the horizon profile at 10, 15, and 20 years (Section 6.2). Throughout this section—and the paper—"average effect" refers to the average cohort estimate. Below we examine dynamics,

[15] Holm's (1979) step-down adjustment controls the probability of falsely rejecting either null hypothesis at the stated level, whatever the dependence between them. Adjusted p-values are never smaller than unadjusted ones. The adjustment inherits the approximations of the placebo p-values.

[16] See the correlation-bound column of Table C.1.

average effects, cohort effects, and magnitudes for output growth and exports after the export turn.

Dynamic paths. Figure 2 plots the event-year effects for GDP per capita and exports per capita, averaged over the four cohorts, with pointwise bands. The pre-adoption paths show how closely each outcome matches its synthetic comparison, though they are not a test of identification (Section 6.3 and Appendix B.4). GDP's pre-adoption path stays close to zero within its band, indicating a close match. Exports rise relative to the synthetic comparison in the three years before adoption. The exports counterfactual therefore rests on a slightly looser match than GDP's (Section 6.3 and Table B.1). Only two cohorts contribute at the earliest pre-adoption year, compared with four thereafter, so the earliest point does not describe all four cohorts.

At adoption, the GDP and exports point estimates are small, and both (pointwise) bands include zero. Yet the effects are first rapid and then steady through the two decades after the regime shift. From the first year after adoption through year 20, both bands lie above zero. The SDID paths for GDP per capita and exports per capita follow similar qualitative patterns (Figure 2), but the export effects are much larger in scale. This difference is particularly salient in the estimates of the total average impact, which we turn to next.

Average effects. Tables 2 and C.2 boil the SDID estimates down to a few summary effects: the total average (cohort) effect over the study period, and snapshots at the 10-, 15-, and 20-year horizons (h). A decade following the regime shift, at $h = 10$, the estimated effect on GDP per capita is 0.417 log points (equivalent to 52 per cent), and the average cohort effect is 0.449 log points (Holm-adjusted placebo $p = 0.002$).

The GDP horizon profile in Tables 2 and C.2 builds to 0.707 log points at 15 years and 0.868 at 20. For exports per capita, the 10-year effect is 1.317 log points (equivalent to 273 per cent), and the average cohort effect is 1.356 log points (Holm-adjusted placebo $p = 0.002$). The exports profile also builds, reaching 2.148 log points at 15 years and 2.425 at 20.

The estimated exports effect exceeds the GDP effect on average and at every horizon of the profile. Across the average and horizon estimates in Tables 2 and C.2, the exports effect tends to be three times the GDP effect in log points. The larger proportional export response is perhaps consistent with the policy's export orientation.

We show these estimates are robust. Section 10 reports the leave-one-out tests, the ASCM comparison, the covariate forms, and the correlation bound.

Cohort effects. We decompose the average effects on GDP per capita and exports per capita by treated economy, to show how they vary across the four East Asian adopters. Figure 3 compares the paths for Japan, Taiwan, South Korea, and Singapore on a common event-time axis for each outcome. Figures C.1 and C.2 plot each cohort's path separately with its pointwise band.

All four GDP cohort point estimates are positive, ranging from Taiwan at 0.303 log points to Singapore at 0.615 (Figure 4 and Table C.3). All four exports cohort point estimates are also positive, ranging from Singapore at 0.737 log points to South Korea at 2.429. Leaving out South Korea moves the exports average the most, lowering it to 0.998 log points. No single omission moves the GDP average by more than 0.056 log points, and every leave-one-out average retains its positive sign (Table E.1).

Magnitudes. To put these estimates in context, we compare the EOIP reforms with studies of broad policy liberalisation, particularly trade liberalisation. Our estimates are more than four times a cross-country panel benchmark but close to synthetic-control estimates for South Korea and Singapore.

Wacziarg and Welch (2008) provide a benchmark based on within-country changes. The authors regress annual GDP per capita growth on a liberalisation indicator with two-way fixed effects for 133 countries over 1950–1998. Growth is 1.42 percentage points higher after liberalisation. Their timing regression shows that the gain arrives with a lag: relative to the years before $T - 3$, growth is 0.30 percentage points higher from T to $T + 2$ (not statistically significant), 1.44 points higher from $T + 3$ to $T + 6$, and 1.02 points higher from $T + 7$ onward.

Cumulating these annual growth coefficients implies an approximate log-level effect of 0.10 at 10 years after liberalisation, 0.15 at 15 years, and 0.20 at 20 years. Each is just under a quarter of our estimate at the corresponding horizon. Their liberalisation dates for South Korea, Taiwan, and Japan are 6, 5, and 9 years later than ours, respectively, which accounts for a large part of the difference in magnitudes.[17]

Billmeier and Nannicini (2013) apply synthetic control methods to the same liberalisation dates but find magnitudes comparable to ours. Ten years after treatment, the log gap between treated and synthetic GDP per capita is 0.649 for South Korea and 0.480 for Singapore, compared with 0.19 to 0.45 for their Latin American episodes.[18] Our 10-year effect of 0.417 log points lies just below their East Asian estimates. Our cohort effects, which range from 0.303 to 0.615, reach into their East Asian range at the top and their Latin American range at the bottom.

Connolly and Yi (2015) instead take a structural route by calibrating a two-country Eaton–Kortum model of South Korea and the G7 with multistage production and traded investment goods. They simulate South Korea's tariff cuts between 1962 and 1989, its duty exemption on imported inputs and capital goods, and the GATT reductions. These policies raise South Korea's manufacturing value added per worker relative to the G7 by 0.143 log points, consumption per worker by 0.191, and the export share of manufacturing value added by 1.81 log points, close to the observed rise of 1.84. Since log exports per capita equals the log export share of GDP plus log GDP per capita, our averages imply an export-share effect of about $1.356 - 0.449$ log points, roughly half of what they attribute to trade policy. The two shares have different denominators, GDP in ours and manufacturing value added in theirs, so this comparison is merely indicative. Their output effect, which measures manufacturing value added per worker relative to the G7, is about a third of our average cohort effect on GDP per capita. Trade policy in a neoclassical model therefore accounts for only part of the income gain we estimate.

In short, our GDP effects are roughly three to four times the average within-country effect of liberalisation and the trade-policy effect from a calibrated model, but they are of the same order as the synthetic-control estimates for the East Asian episodes.

---

[17] They date South Korea to 1968, Taiwan to 1963, Singapore to 1965, and Japan to 1964.

[18] Our calculation from the means of real GDP per capita at $T_0 + 10$ for the treated country and its synthetic control in their Tables 5 and 6.

8 Results: Factor Accumulation and Productivity

Having examined output and export expansion above, we now turn to how factors and productivity responded to the shift towards export-oriented industrial policy. Section 8.1 reports the effects on inputs, and Section 8.2 estimates the effect on TFP directly. Beyond showing how each responded to the policy shift, these estimates also feed the growth accounting in Section 9.

8.1 Factor responses effects

We use our headline SDID specification (Section 6.1) to estimate the effects of the export-orientation turn on factor inputs and investment. These results point to increases in employment per capita, capital, capital per worker, and the investment share of GDP. Table 3 reports average cohort effects of 0.132, 0.489, 0.398, and 0.468 log points, respectively. The table's horizon profile reports effects at $h = 10$, 15, and 20 years after adoption. Figure 5 plots average event-year effects for each outcome.

We highlight three strong, positive patterns seen across labour and capital outcomes. First, employment expansion coincides with strong capital deepening. Our estimates indicate higher employment per capita and more capital per worker relative to the weighted donors. Second, the positive investment-share estimate suggests that investment rose relative to output following the policy turn. Yet while capital accumulation and capital deepening tend to be precisely estimated, the investment-share estimates are noisy. Third, the capital-per-worker effect increases across the three horizons, whereas the employment-per-capita and investment-share effects peak at year 15 and are lower at year 20.

The positive patterns seen for employment and capital are also broad. That is, all four factor outcomes have positive point estimates in every economy, but the balance between employment expansion and capital deepening differs. Taiwan has the smallest employment-per-capita response alongside a larger capital-per-worker response (Figure D.1), suggesting that capital deepening is the more prominent feature of its estimated factor adjustment. Singapore has the largest point estimates for both outcomes, combining higher employment per capita with more capital per worker. The responses thus share a direction but differ in their mix of employment expansion and capital deepening. Panel B of Table C.1 and Panel B of Table C.3 report placebo inference for the factor outcomes.

Yet while traditional factors—labour and capital—responded positively, the evidence around human capital is less clear. Although we use SDID, we report estimates without controls because human capital is also a core control in the headline specification.

The human-capital response, measured by the PWT index, is small and uneven across the treated economies. The average effect is 0.013 log points, with a positive point estimate only for South Korea and negative estimates for Japan, Taiwan, and Singapore (Figure D.2 and table D.1). Importantly, South Korea's estimate—the only positive estimate—relies entirely on a single donor (France), and thus may be brittle.

Although these human capital results (Panel B of Table C.1 and Panel B of Table C.3) provide limited support for a human-capital response across the four economies, but they are not exhaustive. For instance, we use a single measure and operationalisation of human

capital. Theories of human capital and East Asian growth may also consider other forms of human capital accumulation embodied in skilled labour.

Yet, beyond factor accumulation, the role of TFP is central to debates around growth and East Asian industrial development, which we turn to next.

8.2 Measured productivity effect

How does total factor productivity respond to the export-oriented shift? TFP is central to debates about East Asian growth, and early proponents of export promotion emphasised the potential productivity-enhancing effects of East Asian development strategy (e.g., Balassa, 1988). We explore the relationship between TFP and the export turn, using simple aggregate measures of productivity.

We follow basic conventions and measure TFP as the part of output per worker that capital per worker and human capital do not explain:

$$\ln A \;=\; \ln(Y/L) \;-\; \alpha\, \ln(K/L) \;-\; (1-\alpha)\, \ln H, \tag{2}$$

where Y is real GDP, L is employment (persons engaged), K is the capital stock, and H is the PWT human-capital index. The first term, $\ln(Y/L)$, is output per worker.[19] Given the vintage of the data, we cannot rely on pre-built TFP series.

We calculate TFP using Equation (2) in two ways. The first, our benchmark measure (TFP), sets the capital elasticity α to 1/3, which we then vary in our later analyses. Second, we compute a more empirically grounded measure, denoted "TFP (labour share)", which uses α_i, country i's average elasticity derived from PWT's labour income share measure $(1 - \alpha_i)$. PWT's labour share elasticities are highly stable within each country over the sample period. The last term of Equation (2) removes human capital's contribution from TFP. Both measures of TFP are built from PWT inputs and are distinct from PWT's own TFP series.

TFP responds positively to the shift to EOIP, and these effects are significant for aggregate estimates. Table 4 reports an average TFP estimates from our SDID estimates of 0.197 log points. As with human capital SDID estimates above, we omit covariate adjustment to stay close to unadjusted growth accounting and because human capital is itself part of the covariate. However, these results are robust to covariate sets. Our results are also robust across our two simple measures of TFP.

Figure 6 plots the average event-year effects for both TFP measures. Importantly, the estimates track the same dynamic pattern as outputs (exports) and factor outcomes. TFP's horizon profile rises from 0.193 at $h = 10$ to 0.303 at $h = 15$ and 0.361 at $h = 20$ years after adoption. Table D.2 reports placebo inference for both TFP measures. Using each economy's own PWT labour share (capital shares 0.25 to 0.60) in place of one-third, the average TFP effect is 0.170 log points against 0.197, positive at every horizon. This is, in part, driven by

[19] Both measures rest on standard assumptions: Cobb–Douglas technology with constant returns to scale, $Y = AK^{\alpha}(HL)^{1-\alpha}$, where A is Hicks-neutral technology and H is labour-augmenting human capital. The value 1/3 is a simple calibration benchmark, while α_i treats capital's income share as the elasticity.

the availability of donors with complete TFP estimates using empirical elasticities, which are more sparse in early postwar data. For Singapore, where the TFP effects are weakest, the effect falls from 0.255 to 0.050, on a smaller donor pool.

9 Accounting for the Output Response

We next ask how much of the estimated gain in GDP per capita can be accounted for by the gains in employment per capita and capital deepening. We split the GDP-per-capita effect into three parts: the part accounted for by employment per capita, the part accounted for by capital per worker, and a residual, which neither accounts for.

9.1 Treatment effects meet accounting

We use our estimates thus far to perform a type of accounting exercise. Start with a residual r, which is the part of the output-per-worker response not accounted for by capital deepening at an assumed capital elasticity α. Formally,

$$r \equiv \tau_{Y/N} - \tau_{L/N} - \alpha\tau_{K/L} = \tau_A + (1-\alpha)\tau_H + \tau_\xi, \tag{3}$$

where the residual r subtracts the employment-per-capita effect, $\tau_{L/N}$, and the contribution assigned to capital deepening, $\alpha\tau_{K/L}$, from the GDP-per-capita effect, $\tau_{Y/N}$. Each τ_X denotes the average treatment effect on $\ln X$, and each effect use the same treated economies, periods and aggregation weights.[20] The first equality in Equation (3) defines our target quantity as a combination of treatment effects. The second equality gives this quantity as the sum of the TFP effect, the contribution of human capital, and a term accounting for the difference between the two output series.[21] This equality holds by construction from the definition of TFP when employment, capital and human capital are measured consistently and α is common and fixed.

Empirically, we estimate r by substituting the SDID estimates of the output and factor effects. This yields a "plug-in" estimate $\widehat{r}$,

$$\widehat{r} \equiv \widehat{\tau}_{Y/N} - \widehat{\tau}_{L/N} - \alpha\widehat{\tau}_{K/L}. \tag{4}$$

[20] For horizon-specific calculations, each effect instead refers to the same event-time horizon.

[21] Equation (2) uses PWT output, whereas Y/N here is the core outcome. Define $\xi = \ln(Y/N) - \ln(Y^{\text{PWT}}/N)$, using the same population denominator. Here τ_ξ denotes the average treatment effect on this log ratio, with the same aggregation as the other effects.

The three component estimates come from the previous sections: $\widehat{\tau}_{Y/N}$ is the estimated effect on log GDP per capita, $\widehat{\tau}_{L/N}$ the estimated effect on log employment per capita, and $\widehat{\tau}_{K/L}$ the estimated effect on log capital per worker. Multiplying $\widehat{\tau}_{K/L}$ by α gives the contribution assigned to capital deepening. The separately estimated TFP and human-capital effects are now comparators, and not ingredients required to calculate this residual. Although the SDID comparisons do not enforce an accounting identity, their combination estimates a useful, well-defined quantity: the output-per-worker effect remaining after the contribution assigned to capital deepening at a given elasticity.

We therefore analyze the following equation:

$$\underbrace{\widehat{\tau}_{Y/N}}_{\text{GDP-per-capita}} = \underbrace{\widehat{\tau}_{L/N}}_{\text{employment-per-capita}} + \underbrace{\alpha\widehat{\tau}_{K/L}}_{\text{assigned to capital deepening}} + \underbrace{\widehat{r}}_{\text{residual}}. \tag{5}$$

Equation (5) follows from Equation (4), and we use it to assess whether the estimated factor responses can account for the estimated output response. The calculated residual need not equal the sum of the TFP effect and the human-capital effect.[22] We compare the accounting residual with the separately estimated TFP effect to check whether the two exercises give a similar picture of the response beyond measured capital deepening. The residual combines the estimated output and factor effects, whereas the TFP estimate applies SDID directly to an index that removes capital deepening and human capital from output per worker.

In our analysis of Equation (5), we vary α. The residual is defined algebraically, while interpreting $\alpha\tau_{K/L}$ as the contribution assigned to capital deepening uses a constant-elasticity Cobb–Douglas benchmark. Our elasticity grid applies a common, fixed α. By contrast, TFP (labour share) uses country-specific average elasticities, relaxing cross-country commonality but retaining a fixed elasticity within each country. Thus, we use the estimated elasticities as a separate sensitivity exercise rather than a point on the grid.[23]

9.2 Results: Accounting for output

How much of the estimated output response to the policy turn is accounted for by the corresponding estimated employment and capital responses? Table 5 shows that employment and capital deepening potentially account for sizeable changes in output per capita over the period, but they still leave a substantial role for the residual. Under many parameterisations, estimated employment and capital responses are considerable, yet not large enough to account fully for the estimated output response.

Equation (6) makes the factor contributions concrete under the basic assumptions of Table 5. Here, estimates from measured employment and capital deepening account for about three-fifths of the GDP-per-capita effect at our baseline calibration, $\alpha = 1/3$, and the

[22] For instance, the residual retains human capital, and the exercises use different output series, covariate adjustment, and fitted comparisons.

[23] The accounting calculation does not require competitive factor markets, though interpreting capital's income share as its output elasticity requires additional assumptions.

residual for about two-fifths. This residual measures the part of the response not accounted for by the measured factor inputs.[24]

$$\underbrace{0.449}_{\text{GDP-per-capita}} = \underbrace{0.132}_{\text{employment-per-capita}} + \underbrace{\alpha \times 0.398}_{\text{assigned to capital deepening}} + \underbrace{0.184}_{\text{residual}} \tag{6}$$

The positive residual in Table 5 does not depend on the one-third capital elasticity, and is significant across the α grid. Assigning a larger elasticity to capital reduces the residual, but it remains positive throughout the grid, which covers the economies' own capital shares behind TFP (labour share), for both the average cohort effect and every reported horizon. The break-even elasticities in Table 5 show how far the calibration must move to eliminate the residual: on average, the capital elasticity would need to reach 0.80.

It is also useful to compare our implied estimates in Table 5 with our direct estimates in Table 4 and note two patterns. First, the direct and implied output-per-worker estimates are remarkably close: 0.318 log points versus 0.317 on average (with differences of 0.012 or less at reported horizons). This agreement is reassuring because the two tables are built differently: Table 4 uses PWT output and separately fitted comparisons, whereas Table 5 combines the core GDP-per-capita and employment-per-capita estimates.

Second, our direct TFP estimates (Table 4) resemble the residual (Table 5) in magnitudes and pattern. The average TFP estimate is 0.197, compared with a residual of 0.184 at $\alpha = 1/3$, and both increase between horizons 10 and 20. Because the two share α and the capital measure, this agreement is a consistency check, not a test of α. The residual also includes the human-capital term, but the gap between the two cannot be attributed to that term, since each outcome has its own weights and donor pool.

The question remains, however: what accounts for this residual? This exercise is, of course, limited by strong assumptions and cannot decompose the residual. Early work on growth accounting and East Asia pointed out the compatibility of factor accumulation and TFP growth. For instance, learning, organisational change, and structural transformation can generate both accumulation and productivity improvements (Nelson and Pack, 1999). Our combination of positive capital responses and a positive residual is compatible with such accounts, but cannot identify them. Our patterns may also be compatible with export mechanisms. Reallocation towards more productive firms and sectors is one possible reading of the positive residual (e.g., Pavcnik, 2002). Learning by exporting and technology absorption could be another explanation (e.g., Nelson and Pack, 1999; De Loecker, 2007).

[24] Employment is a headcount and capital is a stock measure. The residual can therefore include hours worked, human capital, utilisation, capital quality, reallocation, and measurement error. And it is neither an independently identified TFP effect nor a causal mediation share.

10 Robustness

Leave one cohort out. For each left-out cohort, the *leave-one-out delta* is the average cohort effect minus the average of the other three cohort effects. Its placebo p-value tests whether that cohort stands apart from the other three more than placebo cohorts do, using the same joint placebo assignments as the pooled estimates (Section 6.4). For GDP per capita, every delta has $p \geq 0.215$. For exports per capita, leaving out South Korea gives a delta of 0.358 log points ($p = 0.001$); leaving out Singapore gives a delta of -0.206 log points ($p = 0.067$; Table E.1).

Augmented synthetic control. We use the augmented form because matching outcome levels and paths simultaneously can be infeasible for synthetic control, as Section 6.1 explains. The point estimates agree across SDID and ASCM: every ASCM cohort effect has SDID's sign, and both ASCM averages lie inside SDID's Gaussian confidence intervals. The ASCM average cohort effects are 0.416 log points for GDP per capita and 1.235 for exports per capita (Table E.2). As with SDID (Section 6.3), ASCM's pre-period root mean squared prediction error (RMSPE) indicates closer pre-adoption fit for GDP per capita than for exports per capita in every cohort. Figures E.1 and E.2 show the dynamic comparisons; the GDP and exports path grids are in the Online Appendix OA-D.

Covariate adjustment. When we add a quadratic term to the covariate adjustment's linear time trend or replace the trend with year dummies, the average cohort effects are, respectively, 0.449 and 0.446 log points for GDP per capita and 1.347 and 1.294 for exports per capita (Appendix E.3), compared with the linear-trend averages in Section 7.

Cross-cohort correlation bound. The bound gives the pooled placebo standard error that would result if the four cohorts' estimation errors were perfectly positively dependent (bound column of Table C.1). Perfect positive dependence is the worst case that Alvarez and Ferman (2026) use for one of their conservative tests in difference-in-differences with few treated units. Both average cohort effects, on GDP per capita and on exports per capita, still exceed 1.96 times this standard error. The rows of that table whose effects do not exceed this threshold are employment per capita (h = 20), capital per worker (h = 10), investment (average cohort effect), investment (h = 10), investment (h = 15), and investment (h = 20). In Table D.2, neither TFP nor TFP (labour share) exceeds 1.96 times its bound, on average or at any reported horizon.

11 Conclusion

The East Asian miracle is often invoked either as evidence for industrial policy or as evidence that outward orientation and market discipline matter more than state intervention. Our findings suggest that this opposition is misleading. Japan, Taiwan, South Korea, and Singapore adopted policy regimes that combined extensive state intervention with exposure

to international markets. We relate this industrial-policy mix to their postwar growth spurts. The distinctive feature of these regimes was not the absence of protection or the wholesale liberalisation of domestic markets, but the increasing use of export performance to allocate and discipline industrial support.

We revisit that experience by treating export-oriented industrial policy as a class of industrial-policy regimes—one pursued at different moments across Asia. These regimes prioritised industrial development through export market expansion. The transition to EOIP across postwar economies often meant departing from import-substitution industrialisation and from the broad nominal protection of output markets. While the portfolio of interventions was diverse, we argue that export-promoting interventions were a shared, essential feature of national development strategies in postwar Asian economies.

We date the shift to export-oriented industrial policy to 1955 in Japan, 1958 in Taiwan, 1962 in South Korea, and 1967 in Singapore. These dates correspond to changes in national development strategy, articulated through de jure and de facto shifts in policy levers and institutions, and are reflected in the periodisation of qualitative scholarship.

We assess our country-level dating using new sector-country-year data on East Asian industrial policy. For Japan, Taiwan, and South Korea in particular, the data show the "turn" as a discernible shift towards export-promoting policy. Our dates broadly correspond to the adoption of export-contingent policies, but these policies can coexist with varying degrees of import protection. Export promotion in postwar economies took many forms.

Our estimates indicate that this turn was followed by substantial improvements in both income and exports. These results should not be interpreted as showing that every export-promoting intervention is effective, or that the East Asian policy packages can be reproduced by adopting isolated export subsidies. The historical regimes combined finance, tax treatment, access to imported inputs, administrative coordination, investment promotion, and sectoral targeting. Their effectiveness plausibly depended on the state's ability to observe performance, enforce conditions, and withdraw support, or at least threaten to do so. The relevant policy lesson is therefore about the design and governance of intervention: support tied to an observable external performance criterion differs from unconditional protection or assistance based solely on domestic political influence.

Our findings carry four qualifications. First, the treatment episodes were not randomly assigned. Although our estimation strategy constructs comparison groups that reproduce pre-treatment trajectories, it cannot rule out every contemporaneous political, geopolitical, or institutional change that coincided with the export-orientation turns. Nor can it separate the industrial-policy core of each package from the complementary measures adopted with it: the estimates measure the effect of the package as a whole. The estimates should therefore be understood as evidence from historically structured comparative cases rather than as the outcome of a literal experiment.

Second, inference remains difficult with four early adopters. Our joint placebo procedure retains dependence across cohort-level estimates and places each country in the one-treated-unit setting for which placebo inference has the clearest justification. Nevertheless, the four early adopters are geographically close and economically interconnected. If their untreated shocks are more correlated than those of the donor economies in a joint placebo assignment, the placebo distribution may be too narrow and the reported p-values too small. Placebo assignments that preserve regional correlation would be a useful extension.

Third, the historical policy record is necessarily uneven. Some countries and periods are described more densely than others, and policy withdrawals are less completely observed than policy introductions. The policy data are consequently better suited to identifying broad changes in composition than to serving as a complete annual inventory of government activity. Continued audit and expansion of the underlying corpus will strengthen the measurement of both the timing and intensity of intervention.

Finally, aggregate GDP and export outcomes do not by themselves identify the mechanisms through which EOIP operated. The estimates are consistent with several channels, including learning through exporting, scale economies, cost discovery, capital accumulation, and the reallocation of resources towards more productive activities. The aggregate factor responses and the accounting residual cannot distinguish among these channels; doing so requires sector-level outcomes, along with closer attention to distributional effects and to the firms and sectors that lost support during the transition.

Subject to these qualifications, the evidence supports a more precise interpretation of the East Asian experience. Governments deliberately reorganised selective interventions around participation in foreign markets and used export performance, to varying degrees, as a test of the firms and industries they supported. The resulting policy turns are visible in the historical record and are followed by sustained divergences in exports and income relative to synthetic counterfactuals. The East Asian miracle therefore points less to a general choice between states and markets than to the importance of institutional arrangements that use market outcomes to discipline state support.

References

Abadie, A. (2021). "Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects". In: *Journal of Economic Literature* 59.2, pp. 391–425. DOI: 10.1257/jel.20191450.

Abadie, A., A. Diamond, and J. Hainmueller (2010). "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program". In: *Journal of the American Statistical Association* 105.490, pp. 493–505. DOI: 10.1198/jasa.2009.ap08746.

— (2015). "Comparative Politics and the Synthetic Control Method". In: *American Journal of Political Science* 59.2, pp. 495–510. DOI: 10.1111/ajps.12116.

Alvarez, L. and B. Ferman (2026). *Inference in Difference-in-Differences with Few Treated Units and Spatial Correlation*. arXiv: 2006.16997v8.

Amsden, A. H. (1989). *Asia's Next Giant: South Korea and Late Industrialization*. New York: Oxford University Press.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). "Synthetic Difference-in-Differences". In: *American Economic Review* 111.12, pp. 4088–4118. DOI: 10.1257/aer.20190159.

— (2024). *synthdid: Synthetic Difference-in-Difference Estimation*. Version 0.0.9. R package. URL: https://github.com/synth-inference/synthdid/tree/70c1ce3eac58e28c30b67435ca377bb48baa9b8a.

Athey, S., M. Bayati, N. Doudchenko, G. Imbens, and K. Khosravi (2021). "Matrix Completion Methods for Causal Panel Data Models". In: *Journal of the American Statistical Association* 116.536, pp. 1716–1730. DOI: 10.1080/01621459.2021.1891924.

Balassa, B. (1978). "Exports and Economic Growth: Further Evidence". In: *Journal of Development Economics* 5.2, pp. 181–189. DOI: 10.1016/0304-3878(78)90006-8.

Balassa, B. (1982). *Development Strategies in Semi-Industrial Economies*. Published for the World Bank. Baltimore: Johns Hopkins University Press.
— (1988). "The Lessons of East Asian Development: An Overview". In: *Economic Development And Cultural Change* 36.3, pp. 273–290.
Barro, R. J. and J.-W. Lee (2013). "A New Data Set of Educational Attainment in the World, 1950–2010". In: *Journal of Development Economics* 104, pp. 184–198. DOI: 10.1016/j.jdeveco.2012.10.001.
Ben-Michael, E., A. Feller, and J. Rothstein (2021). "The Augmented Synthetic Control Method". In: *Journal of the American Statistical Association* 116.536, pp. 1789–1803. DOI: 10.1080/01621459.2021.1929245.
Billmeier, A. and T. Nannicini (2013). "Assessing Economic Liberalization Episodes: A Synthetic Control Approach". In: *The Review of Economics and Statistics* 95.3, pp. 983–1001. DOI: 10.1162/REST_a_00324.
Bolt, J. and J. L. van Zanden (2024). "Maddison Style Estimates of the Evolution of the World Economy: A New 2023 Update". In: *Journal of Economic Surveys* 39.2. Maddison Project Database 2023, pp. 631–671. DOI: 10.1111/joes.12618.
Callaway, B. and P. H. C. Sant'Anna (2021). "Difference-in-Differences with Multiple Time Periods". In: *Journal of Econometrics* 225.2, pp. 200–230. DOI: 10.1016/j.jeconom.2020.12.001.
Cherif, R. and F. Hasanov (2019). *The Return of the Policy That Shall Not Be Named: Principles of Industrial Policy*. IMF Working Paper 19/74. Washington, DC: International Monetary Fund. DOI: 10.5089/9781498305402.001.
— (2024). "The Pitfalls of Protectionism: Import Substitution vs. Export-Oriented Industrial Policy". In: *Journal of Industry, Competition and Trade* 24.1. Article 14, pp. 1–34.
Clarke, D., D. Pailañir, S. Athey, and G. Imbens (2024). "On Synthetic Difference-in-Differences and Related Estimation Methods in Stata". In: *The Stata Journal* 24.4, pp. 557–598. DOI: 10.1177/1536867X241297914.
Clifford, M. L. (1994). *Troubled Tiger: Business, Bureaucrats, and Generals in South Korea*. Armonk, NY: M.E. Sharpe, Inc.
Conley, T. G. and C. R. Taber (2011). "Inference with "Difference in Differences" with a Small Number of Policy Changes". In: *The Review of Economics and Statistics* 93.1, pp. 113–125. DOI: 10.1162/REST_a_00049.
Connolly, M. and K.-M. Yi (2015). "How Much of South Korea's Growth Miracle Can Be Explained by Trade Policy?" In: *American Economic Journal: Macroeconomics* 7.4, pp. 188–221. DOI: 10.1257/mac.20120197. URL: https://www.aeaweb.org/articles?id=10.1257/mac.20120197.
De Loecker, J. (2007). "Do Exports Generate Higher Productivity? Evidence from Slovenia". In: *Journal of International Economics* 73.1, pp. 69–98. DOI: 10.1016/j.jinteco.2007.03.003.
Easterly, W. (1995). "Explaining Miracles: Growth Regressions Meet the Gang of Four". In: *Growth Theories in Light of the East Asian Experience*. Ed. by T. Ito and A. O. Krueger. Vol. 4. NBER-East Asia Seminar on Economics. Chicago: University of Chicago Press, pp. 267–299. URL: http://www.nber.org/chapters/c8553.
Engel, C. and J. Wang (2011). "International Trade in Durable Goods: Understanding Volatility, Cyclicality, and Elasticities". In: *Journal of International Economics* 83.1, pp. 37–52. DOI: 10.1016/j.jinteco.2010.08.007.
European Bank for Reconstruction and Development (2024). *Transition Report 2024–25: Navigating Industrial Policy*. London: European Bank for Reconstruction and Development.
Feenstra, R. C., R. Inklaar, and M. P. Timmer (2015). "The Next Generation of the Penn World Table". In: *American Economic Review* 105.10, pp. 3150–3182. DOI: 10.1257/aer.20130954.
Ferman, B. and C. Pinto (2019). "Inference in Differences-in-Differences with Few Treated Groups and Heteroskedasticity". In: *The Review of Economics and Statistics* 101.3, pp. 452–467. DOI: 10.1162/rest_a_00759.

Fernandes, A. M. and T. Reed (2026). *Industrial Policy for Development: Approaches in the 21st Century*. A World Bank Group Policy Research Report. Washington, DC: World Bank. DOI: 10.1596/978-1-4648-2276-6.

Fouquin, M. and J. Hugot (2016). *Two Centuries of Bilateral Trade and Gravity Data: 1827–2014*. CEPII Working Paper 2016-14. CEPII.

Frank Jr., C. R., K. S. Kim, and L. E. Westphal (1975). *Foreign Trade Regimes and Economic Development: South Korea*. New York: National Bureau of Economic Research. URL: http://www.nber.org/books/fran75-1.

Frankel, J. A. and D. Romer (1999). "Does Trade Cause Growth?" In: *American Economic Review* 89.3, pp. 379–399. DOI: 10.1257/aer.89.3.379.

Haggard, S., B.-K. Kim, and C.-i. Moon (1991). "The Transition to Export-led Growth in South Korea: 1954–1966". In: *Journal of Asian Studies* 50.4, pp. 850–73. DOI: 10.2307/2058544.

Ho, S. P. S. (1987). "Economics, Economic Bureaucracy, and Taiwan's Economic Development". In: *Pacific Affairs* 60.2, pp. 226–247. DOI: 10.2307/2758133.

Holm, S. (1979). "A Simple Sequentially Rejective Multiple Test Procedure". In: *Scandinavian Journal of Statistics* 6.2, pp. 65–70. URL: https://www.jstor.org/stable/4615733.

Hsieh, C.-T. (2002). "What Explains the Industrial Revolution in East Asia? Evidence from the Factor Markets". In: *American Economic Review* 92.3, pp. 502–526. DOI: 10.1257/00028280260136372.

Irwin, D. A. (2021a). *From Hermit Kingdom to Miracle on the Han: Policy Decisions that Transformed South Korea into an Export Powerhouse*. Working Paper 29299. National Bureau of Economic Research. DOI: 10.3386/w29299.

— (2021b). *How Economic Ideas Led to Taiwan's Shift to Export Promotion in the 1950s*. Working Paper 29298. National Bureau of Economic Research. DOI: 10.3386/w29298.

— (2021c). "The Rise and Fall of Import Substitution". In: *World Development* 139, p. 105306. DOI: 10.1016/j.worlddev.2020.105306.

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

Juhász, R. and N. Lane (2024). "The Political Economy of Industrial Policy". In: *Journal of Economic Perspectives* 38.4, pp. 27–54. DOI: 10.1257/jep.38.4.27.

Juhász, R., N. Lane, E. Oehlsen, and V. C. Perez (2025). *Measuring Industrial Policy: A Text-Based Approach*. NBER Working Paper 33895. Revised July 2025. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w33895. URL: http://www.nber.org/papers/w33895.

Juhász, R., N. Lane, and D. Rodrik (2024). "The New Economics of Industrial Policy". In: *Annual Review of Economics* 16, pp. 213–242. DOI: 10.1146/annurev-economics-081023-024638.

Kim, B.-K. (2011). "The Leviathan: Economic Bureaucracy under Park". In: *The Park Chung Hee Era: The Transformation of South Korea*. Ed. by B.-K. Kim and E. F. Vogel. Cambridge, MA: Harvard University Press, pp. 200–232. DOI: 10.4159/harvard.9780674061064.c7.

Kim, H.-A. (2004). *Korea's Development Under Park Chung Hee*. RoutledgeCurzon.

Kranz, S. (2022). "Synthetic Difference-in-Differences with Time-Varying Covariates". Ulm University. First version 2021-12-16. URL: https://github.com/skranz/xsynthdid/blob/b79dc81b02bfc76b36b94442f555eae68055101f/paper/synthdid_with_covariates.pdf.

Kuo, S. W. Y. and G. Ranis (1982). *The Taiwan Success Story: Rapid Growth With Improved Distribution In The Republic Of China, 1952–1979*. Routledge.

Lall, S. (1997). *Selective Policies for Export Promotion: Lessons from the Asian Tigers*. Research for Action 43. Helsinki: UNU World Institute for Development Economics Research (UNU/WIDER). URL: https://ageconsearch.umn.edu/record/295338/files/RFA43.pdf.

Lane, N. (2020). "The New Empirics of Industrial Policy". In: *Journal of Industry, Competition and Trade* 20.2, pp. 209–234. DOI: 10.1007/s10842-019-00323-2.

— (2025). "Manufacturing Revolutions: Industrial Policy and Industrialization in South Korea". In: *The Quarterly Journal of Economics* 140.3, pp. 1683–1741. DOI: 10.1093/qje/qjaf025.

Lawrence, R. Z. and D. E. Weinstein (2001). "Trade and Growth: Import Led or Export Led? Evidence from Japan and Korea". In: *Rethinking the East Asian Miracle*. Ed. by J. E. Stiglitz and S. Yusuf. New York: World Bank and Oxford University Press. Chap. 10, pp. 379–408.
Lee, S. A. (1973). *Industrialization in Singapore*. Longman Australia.
Lim, L. (1995). "Foreign Investment, the State and Industrial Policy in Singapore". In: *Asian Industrialization and Africa: Studies in Policy Alternatives to Structural Adjustment*. Ed. by H. Stein. St. Martin's Press, pp. 205–238.
Müller, K., C. Xu, M. Lehbib, and Z. Chen (2025). *The Global Macro Database: A New International Macroeconomic Dataset*. NBER Working Paper 33714. Cambridge, MA: National Bureau of Economic Research. DOI: 10.3386/w33714.
Nelson, R. R. and H. Pack (1999). "The Asian Miracle and Modern Growth Theory". In: *The Economic Journal* 109.457, pp. 416–436. DOI: 10.1111/1468-0297.00455.
Pack, H. (2000). "Industrial Policy: Growth Elixir or Poison?" In: *World Bank Research Observer* 15.1, pp. 47–67. DOI: 10.1093/wbro/15.1.47.
Pack, H. and K. Saggi (2006). "Is there a case for industrial policy? A critical survey". In: *The World Bank Research Observer* 21.2, pp. 267–297. DOI: 10.1093/wbro/lkl001.
Pavcnik, N. (2002). "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants". In: *The Review of Economic Studies* 69.1, pp. 245–276. DOI: 10.1111/1467-937X.00205.
Phipson, B. and G. K. Smyth (2010). "Permutation *P*-values Should Never Be Zero: Calculating Exact *P*-values When Permutations Are Randomly Drawn". In: *Statistical Applications in Genetics and Molecular Biology* 9.1, 39. DOI: 10.2202/1544-6115.1585.
Reed, T. (2024). "Export-Led Industrial Policy for Developing Countries: Is There a Way to Pick Winners?" In: *Journal of Economic Perspectives* 38.4, pp. 3–26. DOI: 10.1257/jep.38.4.3.
Rodríguez, F. and D. Rodrik (2001). "Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence". In: *NBER Macroeconomics Annual 2000* 15, pp. 261–338. DOI: 10.2307/3585399.
Rodrik, D. (1994). *King Kong Meets Godzilla: The World Bank and the East Asian Miracle*. CEPR Discussion Paper 944. London: Centre for Economic Policy Research.
— (1995). "Getting Interventions Right: How South Korea and Taiwan Grew Rich". In: *Economic Policy* 10.20, pp. 53–107. DOI: 10.2307/1344538.
— (2012). "Why We Learn Nothing from Regressing Economic Growth on Policies". In: *Seoul Journal of Economics* 25.2, pp. 137–151.
Seth, M. J. (2017). *South Korea's Economic Development, 1948–1996*. Oxford University Press.
Shinohara, M. (1964). "Evaluation of Economic Plans in the Japanese Economy". In: *Weltwirtschaftliches Archiv* 92, pp. 208–21.
Soon, T.-W. and W. A. Stoever (1996). "Foreign Investment and Economic Development in Singapore: A Policy-Oriented Approach". In: *The Journal of Developing Areas* 30.3, pp. 317–340.
Suryanarayan, V. (1969). "Singapore in Malaysia". In: *International Studies* 11.1, pp. 1–40. DOI: 10.1177/002088176901100101.
Tsai, P.-L. (1999). "Explaining Taiwan's Economic Miracle: Are the Revisionists Right?" In: *Agenda: A Journal of Policy Analysis and Reform* 6.1, pp. 69–82.
Van Elkan, R. (1995). "Accounting for Growth in Singapore". In: *Singapore: A Case Study in Rapid Development*. Ed. by K. Bercuson. International Monetary Fund, pp. 4–10. DOI: 10.5089/9781557754639.084.
Vogel, E. F. (1990). *The Four Little Dragons: The Spread of Industrialization in East Asia*. Harvard University Press.
Wacziarg, R. and K. H. Welch (2008). "Trade Liberalization and Growth: New Evidence". In: *The World Bank Economic Review* 22.2, pp. 187–231. DOI: 10.1093/wber/lhn007.
Wade, R. (1990). *Governing the Market: Economic Theory and the Role of Government in East Asian Industrialization*. Princeton University Press.

Westphal, L. E. (1978). "The republic of Korea's experience with export-led industrial development". In: *World Development* 6.3, pp. 347–82. DOI: 10.1016/0305-750X(78)90113-4.
— (1990). "Industrial Policy in an Export-Propelled Economy: Lessons from South Korea's Experience". In: *Journal of Economic Perspectives* 4.3, pp. 41–59. DOI: 10.1257/jep.4.3.41.
Wolf, C. (1962). "Economic Planning in Korea". In: *Asian Survey* 2.10, pp. 22–28. DOI: 10.2307/3023492.
Woo, J.-e. (1991). *Race to the Swift: State and Finance in Korean Industrialization*. Columbia University Press.
World Bank (1993). *The East Asian Miracle: Economic Growth and Public Policy*. A World Bank Policy Research Report. New York: Oxford University Press. URL: https://documents.worldbank.org/curated/en/975081468244550798.
Xu, Y. (2017). "Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models". In: *Political Analysis* 25.1, pp. 57–76. DOI: 10.1017/pan.2016.2.
Yamada, Y. (1956). "On the Five Year Economic Plan in Japan". In: *The Annals of the Hitotsubashi Academy* 7.1, pp. 33–45.
Yoo, J. (2017). "Korea's Rapid Export Expansion in the 1960s: How It Began". In: *KDI Journal of Economic Policy* 39.2, pp. 1–23. DOI: 10.23895/kdijep.2017.39.2.1.
Young, A. (1992). "A Tale of Two Cities: Factor Accumulation and Technical Change in Hong Kong and Singapore". In: *NBER Macroeconomics Annual 1992*. Ed. by O. J. Blanchard and S. Fischer. Cambridge, MA: MIT Press, pp. 13–64. DOI: 10.2307/3584993.
— (1995). "The Tyranny of Numbers: Confronting the Statistical Realities of the East Asian Growth Experience". In: *The Quarterly Journal of Economics* 110.3, pp. 641–680. DOI: 10.2307/2946695.

Table 1. Estimation windows and adoption years by cohort

Cohort	Adoption year	Pre-adoption years	Sample coverage
Japan	1955	1949–1954	1949–1975
Taiwan	1958	1952–1957	1952–1978
South Korea	1962	1955–1961	1955–1982
Singapore	1967	1960–1966	1960–1987

Notes: The windows apply to cohort-level synthetic difference-in-differences (SDID) estimates of log GDP per capita and log exports per capita. The sample starts reflect historical circumstances and available data: Dodge Line stabilisation for Japan, post-retreat currency stabilisation for Taiwan, and the post-armistice period for South Korea. Singapore's sample starts in 1960 to provide a pre-adoption period supported by the available historical data, including years before independence. Here c indexes cohorts and a_c denotes the adoption year. The pre-adoption observation count, T_c^-, is six for Japan and Taiwan and seven for South Korea and Singapore. Sample coverage includes the adoption year ($h = 0$) and the following 20 years, giving 21 annual observations from adoption onward.

Table 2. Average cohort effects and horizon profiles: core outcomes

Outcome	Average cohort effect	h = 10	h = 15	h = 20
GDP per capita	0.449	0.417	0.707	0.868
	(0.080)	(0.092)	(0.125)	(0.154)
Exports per capita	1.356	1.317	2.148	2.425
	(0.178)	(0.207)	(0.276)	(0.328)

Notes: The table reports synthetic difference-in-differences estimates, in log points, of the effect of the export-orientation turn. Placebo standard errors are in parentheses; Table C.1 reports placebo p-values, simulation ranges and approximate intervals.

Sample: The treated cohorts, with their adoption years, are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967).

Variables: The average cohort effect is the mean of the cohort effects, each averaged over event times 0 to 20 (event time 0 is the adoption year). Each h column gives the mean of the cohort effects at event time h.

Method: Inference compares each estimate with 1,000 joint placebo assignments, and the placebo standard error is the standard deviation of the statistic across its placebos (Arkhangelsky et al., 2021, Algorithm 4). It also compares each cohort effect with runs that use every eligible donor once as the placebo. Estimates use the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted (see Section 6.1). The headline family is the average cohort effect on GDP per capita and on exports per capita; its Holm-adjusted placebo p-values are in Table C.2.

Table 3. Average cohort effects and horizon profiles: factor outcomes

Outcome	Average cohort effect	h = 10	h = 15	h = 20
Employment per capita	0.132	0.128	0.234	0.213
	(0.032)	(0.033)	(0.049)	(0.064)
Capital	0.489	0.437	0.761	1.079
	(0.084)	(0.091)	(0.134)	(0.163)
Capital per worker	0.398	0.353	0.586	0.946
	(0.089)	(0.102)	(0.138)	(0.168)
Investment	0.468	0.401	0.666	0.575
	(0.136)	(0.182)	(0.196)	(0.204)

Notes: The table reports pooled synthetic difference-in-differences (SDID) effects of the export-orientation turn on the factor outcomes, in log points. Parentheses contain placebo standard errors; the estimate plus or minus 1.96 standard errors is an approximate 95 per cent Gaussian confidence interval.

Sample: The treated cohorts and adoption years are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). Their windows begin in 1950, 1952, 1955, and 1960 and end in 1975, 1978, 1982, and 1987; donor pools vary by cohort and outcome.

Variables: The rows are employment per capita, the capital stock, capital per worker, and the investment share of gross domestic product (GDP); each outcome is in logs. Capital is the Penn World Table (PWT) 10.0 `rnna` stock at constant 2017 national prices; capital per worker divides it by same-vintage PWT persons engaged (`emp`). `rnna` measures a stock, not capital services. The average cohort effect is the mean of the four cohort effects, each averaged over event times 0 to 20, with adoption at event time 0. The h columns report the horizon profile: the mean of the four event-year effects at the listed years since adoption.

Method: Estimates use the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted. Each pooled statistic is compared with 1,000 joint placebo assignments. A placebo standard error is the standard deviation of the statistic across its placebos.

Table 4. Growth-accounting effects and horizon profiles

Outcome	Average cohort effect	h = 10	h = 15	h = 20
Output per worker	0.318	0.301	0.477	0.655
	(0.093)	(0.105)	(0.139)	(0.163)
Human capital	0.013	0.008	0.013	0.036
	(0.018)	(0.020)	(0.029)	(0.038)
TFP	0.197	0.193	0.303	0.361
	(0.071)	(0.081)	(0.103)	(0.121)
TFP (labour share)	0.170	0.160	0.275	0.322
	(0.070)	(0.080)	(0.098)	(0.117)

Notes: The table reports synthetic difference-in-differences estimates, in log points, of the effect of the export-orientation turn on the growth-accounting outcomes. The effects are separately estimated, each with its own unit weights, time weights, window and donor pool, and are not an additive decomposition of the estimated output effect. Placebo standard errors are in parentheses; Table D.2 reports placebo p-values, simulation ranges and approximate intervals.

Sample: The treated cohorts, with their adoption years, are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967).

Variables: The average cohort effect is the mean of the cohort effects, each averaged over event times 0 to 20 (event time 0 is the adoption year). Each h column gives the mean of the cohort effects at event time h. Log output per worker is $\ln(Y/L)$, where Y is real GDP at constant national prices (`rgdpna`) and L is same-vintage PWT persons engaged (`emp`). Log human capital is $\ln H$, where H is the PWT human-capital index (`hc`). Log TFP (labour share) is $\ln[(Y/L)(K/L)^{-\alpha_i}H^{-(1-\alpha_i)}]$, where K is the PWT 10.x `rnna` stock at constant 2017 national prices, Y is real GDP at constant national prices (`rgdpna`), L is same-vintage PWT persons engaged (`emp`), and H is the PWT human-capital index (`hc`). The economy-constant capital share α_i is one minus the mean of the first five admissible labour shares on or after 1950 (strictly between zero and one).

Method: Inference compares each estimate with 1,000 joint placebo assignments, and the placebo standard error is the standard deviation of the statistic across its placebos (Arkhangelsky et al., 2021, Algorithm 4). The effects use no controls: the specification projects no covariates out of the outcome (see Section 8). The outcomes lie outside the headline family (the average cohort effect on GDP per capita and on exports per capita), whose placebo p-values carry a Holm adjustment.

Table 5. Growth-accounting residual under alternative capital elasticities α

	Average cohort effect	$h = 10$	$h = 15$	$h = 20$
$\Delta \ln(Y/L)$	0.317	0.289	0.473	0.655
$\Delta \ln(K/L)$	0.398	0.353	0.586	0.946
Growth residual $\Delta \ln(Y/L) - \alpha \, \Delta \ln(K/L)$, with α the capital elasticity				
$\alpha = 0$	0.317	0.289	0.473	0.655
$\alpha = 0.30$	0.198	0.183	0.297	0.371
$\alpha = 1/3$	0.184	0.171	0.278	0.340
$\alpha = 0.40$	0.158	0.148	0.239	0.277
$\alpha = 0.50$	0.118	0.113	0.180	0.182
$\alpha = 0.60$	0.078	0.077	0.121	0.087
Break-even α	0.80	0.82	0.81	0.69
Share of GDP-per-worker effect ($\alpha = 1/3$)	0.58	0.59	0.59	0.52
Share of GDP-per-capita effect ($\alpha = 1/3$)	0.41	0.41	0.39	0.39

Notes: The table reports the growth-accounting residual implied by the per-worker decomposition at each capital elasticity; it is not a separate estimate of total factor productivity (TFP) or a causal mediation share.

Sample: Columns give the pooled average-cohort effects and pooled effects at $h = 10$, 15 and 20 for the four treated cohorts.

Variables: $\Delta \ln(Y/L)$ is the GDP-per-capita effect minus the employment-per-capita effect; $\Delta \ln(K/L)$ is the direct capital-per-worker estimate. Capital is the Penn World Table (PWT) 10.0 stock at constant 2017 national prices (`rnna`), not capital services; capital per worker divides that stock by same-vintage PWT persons engaged (`emp`).

Method: The calculations use the point estimates in Tables 2 and 3. They assume constant returns to scale, Hicks-neutral technology, a common capital elasticity in treated and control economies, and an adequate first-order approximation. The residual does not identify a causal mediation share.

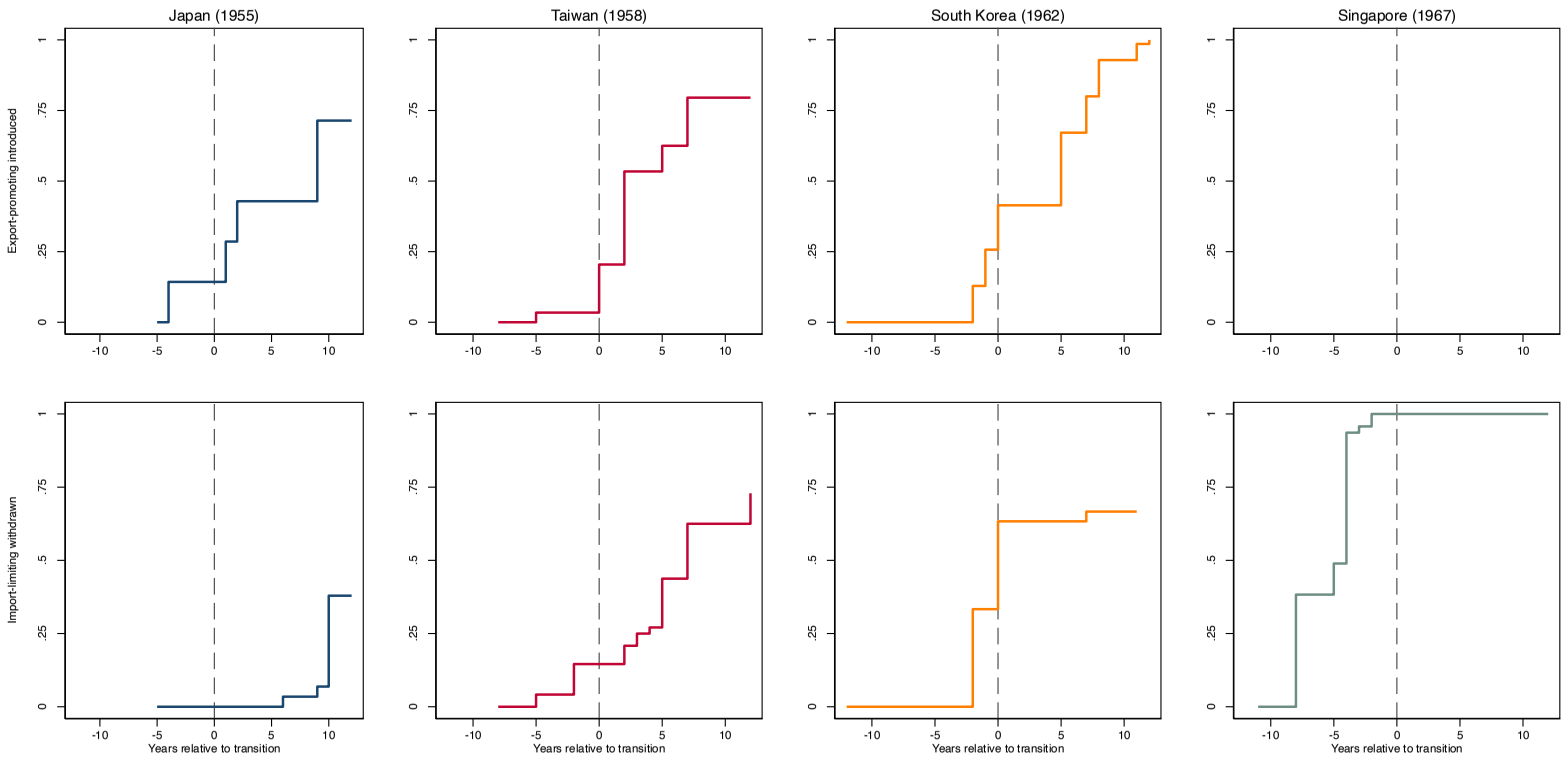


Figure 1. Timing of trade policy change: the early adopters

Notes: Each panel plots the cumulative share of a country's own policies of the stated type and direction that had been recorded by year t, with the country aligned on its own transition year (vertical dashed line). The denominator is the country's full total over the 1949–1990 sample window, so every series runs on a 0–1 scale irrespective of how many policies the country has. The series are step functions of the raw event dates, with no smoothing and no moving window. A curve need not start at 0 or end at 1 inside the plotted ±12-year range: events before $t - 12$ are already accumulated at the left edge, and events after $t + 12$ enter the denominator but are not yet accumulated at the right edge, so a curve ending below 1 indicates that part of the country's record lies outside the window. The withdrawal panels cover *recorded* withdrawals only: a policy with no coded end year generates no withdrawal event and is absent from numerator and denominator alike, so these are timing distributions of the withdrawals the sources describe, not shares of introduced policies that were later removed. Singapore records no export-promoting policies, so that panel is empty; its import-limiting withdrawals are concentrated in 1959–1965 and precede its 1967 date. Indonesia is excluded throughout, as it has no transition date.

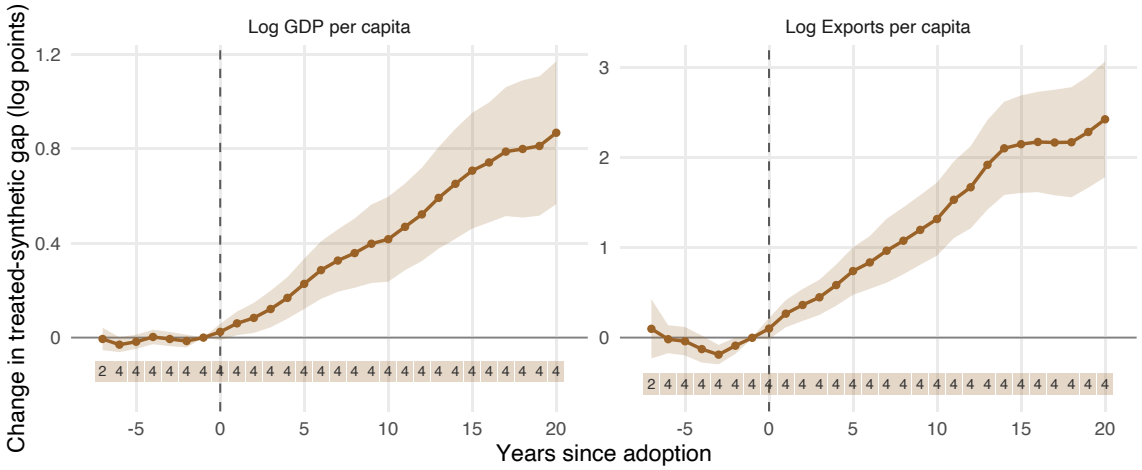


Figure 2. Average event-year effects, by outcome

Notes: Each panel plots the average event-year effect on one outcome, the mean of the available cohorts' event-year effects at each event time, in log points, from synthetic difference-in-differences (SDID). The panels are log GDP per capita and log exports per capita, in that order, and each has its own vertical scale. Event time is the calendar year relative to adoption; the vertical line marks adoption at event time 0. Each cohort's event-year effect is the change in the treated–synthetic gap at that event time relative to its fitted time-weighted pre-adoption baseline, under the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted (see Section 6.1). The cohorts are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967); their fitted windows end in 1975, 1978, 1982, and 1987, respectively, and the panels plot event times −7 to 20. Donors exclude the economies of Eastern Asia and South-eastern Asia, and India. The shaded band is a pointwise approximate 95 per cent magnitude band: the effect plus or minus 1.96 times the standard deviation of the placebo effects at that event time, from joint placebo assignments averaged over the same cohorts and event times. Each placebo path is measured from its own fit's time-weighted baseline; the pre-adoption band is a fit diagnostic, not a pre-trend test. The boxed strip along the bottom of each panel gives the number of cohorts contributing to the mean at each event time; all four cohorts contribute at every event time from 0 to 20.

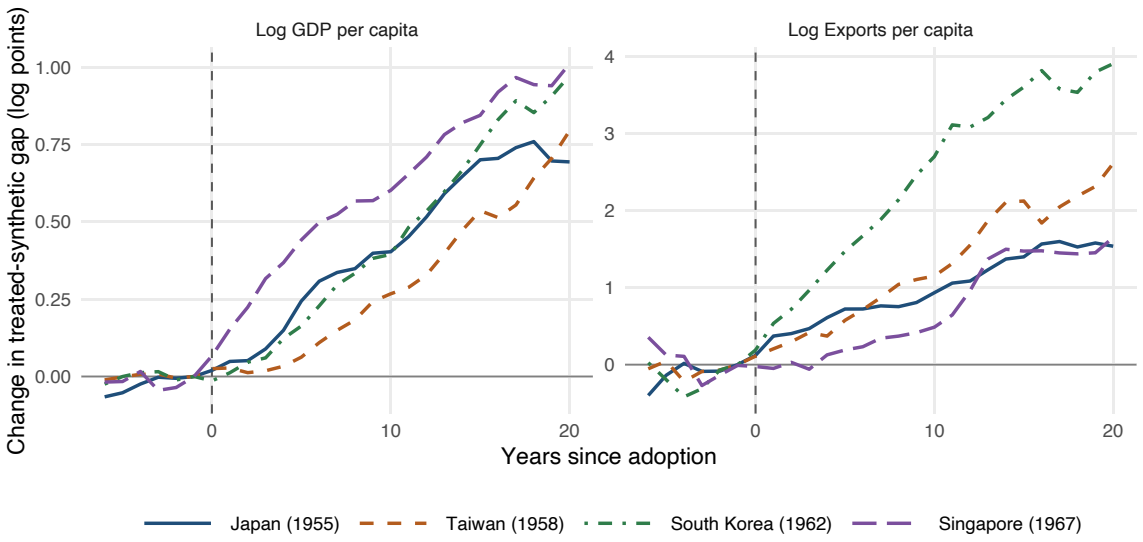


Figure 3. Cohort effect paths in event time, by outcome

Notes: The left and right panels overlay cohort-specific synthetic difference-in-differences (SDID) effects on log GDP per capita and log exports per capita, respectively, in log points against years relative to adoption. Each effect is the change in the treated–synthetic gap relative to its fitted time-weighted pre-adoption baseline under the headline SDID specification, which residualises the outcome on three covariates (logged pre-adoption population, the logged human-capital index and standardised inflation) before the weights are fitted. The vertical dashed line marks adoption at zero; colour and line type in the legend jointly identify Japan (1955), Taiwan (1958), South Korea (1962) and Singapore (1967). Their cohort windows are 1949–1975, 1952–1978, 1955–1982 and 1960–1987, respectively.

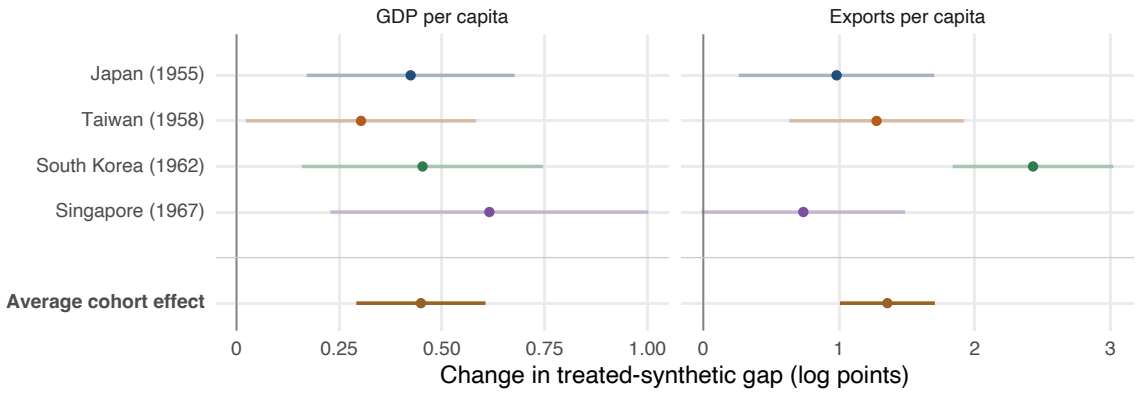


Figure 4. Cohort effects and their average, with approximate intervals, by outcome

Notes: Each panel reports synthetic difference-in-differences (SDID) effects on one outcome, in log points. The panels are GDP per capita and exports per capita, in that order, and each has its own horizontal scale; the vertical line marks zero. Estimates are for the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted (see Section 6.1). The cohorts are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). For each cohort the dot is its average effect over event times 0 to 20, where event time 0 is the adoption year; below the separator is the average cohort effect, the mean of the cohort effects. Dots are estimates and whiskers are approximate 95 per cent magnitude intervals, the estimate plus or minus 1.96 placebo standard errors, where a placebo standard error is the standard deviation of the statistic across its placebos. The cohorts' whiskers are drawn lighter than the average cohort effect's. The figure prints no p-values: the cohorts' rank placebo p-values, over runs that use every eligible donor once as the placebo, are in Table C.3, and the average cohort effect's placebo p-value, over joint placebo assignments, is in Table C.1.

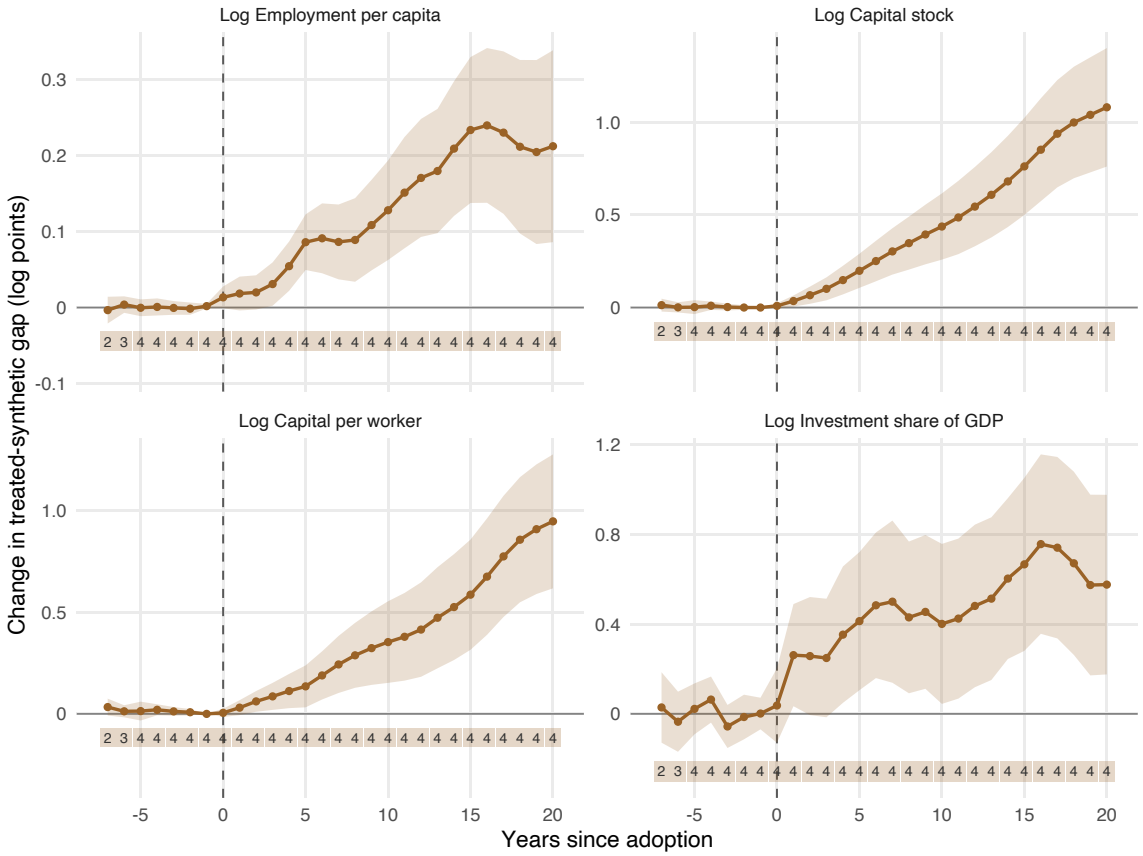


Figure 5. Average event-year effects, by factor outcome

Notes: Each panel plots the mean of the available cohort-specific synthetic difference-in-differences (SDID) effects at each event time for employment per capita, the capital stock, capital per worker, or the investment share of gross domestic product (GDP). Capital is the Penn World Table (PWT) 10.0 `rnna` stock at constant 2017 national prices; capital per worker divides it by same-vintage PWT persons engaged (`emp`). The vertical axis is the change in the treated–synthetic gap, in log points, relative to its fitted time-weighted pre-adoption baseline; the horizontal axis is years since adoption, and the dashed vertical line marks adoption. The shaded bands are pointwise approximate 95 per cent intervals: the effect plus or minus 1.96 times the standard deviation of placebo effects at that event time from 1,000 joint placebo assignments. The strip below each panel gives the number of contributing cohorts at each event time. The treated cohorts and adoption years are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967); their windows begin in 1950, 1952, 1955, and 1960 and end in 1975, 1978, 1982, and 1987. Estimates use the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted.

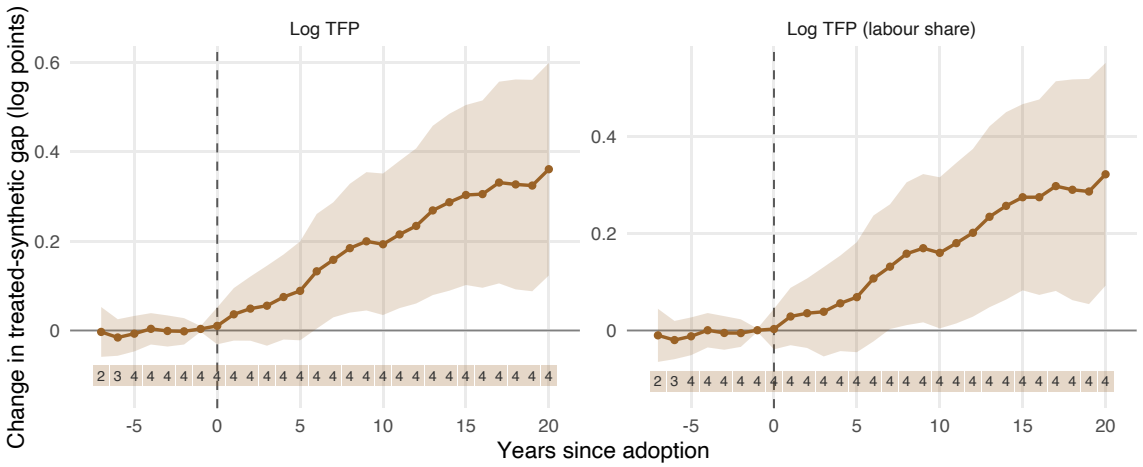


Figure 6. Average event-year effects, by total factor productivity (TFP) measure

Notes: Each panel plots one outcome's average event-year effect, in log points: the mean of the available cohorts' synthetic difference-in-differences (SDID) event-year effects at each event time. The panels are log TFP and log TFP (labour share), in that order, and each has its own vertical scale. Event time 0 is the adoption year (vertical line). A cohort's event-year effect is the change in the treated–synthetic gap relative to its fitted time-weighted pre-adoption baseline, under the no-controls specification, which projects no covariates out of the outcome (see Section 8). The cohorts are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967); their fitted windows end in 1975, 1978, 1982, and 1987, respectively, and the panels plot event times −7 to 20. The shaded band is a pointwise approximate 95 per cent magnitude band: the effect plus or minus 1.96 times the standard deviation of the placebo effects at that event time, from joint placebo assignments averaged over the same cohorts. Each placebo path is measured from its own fit's time-weighted baseline; the pre-adoption band is a fit diagnostic, not a pre-trend test. The strip below each panel gives the number of cohorts contributing at each event time; all four cohorts contribute from event time 0 to 20. Log TFP (labour share) uses each economy's own capital elasticity, taken from its average PWT labour share; see Table D.2.

A. Data and Policy Coding Appendix

A.1 Variables and sample construction

Outcome variables. Our two core outcomes are GDP per capita and exports per capita, both expressed in logs. We take GDP per capita directly from Maddison, which reports real GDP per capita in 2011 international dollars (purchasing-power-parity adjusted), giving us a single, internally consistent numeraire across countries and time.[25] We construct exports per capita from TRADHIST bilateral trade flows, deflated to 2011 prices using a UK GDP deflator and converted to US dollars at the 2011 average exchange rate, then divided by Maddison population. Because TRADHIST's coverage of Taiwan and Singapore is unreliable in this period—both economies' political status led to their exclusion from many contemporaneous trade databases—we replace their export series with the corresponding series derived from the Penn World Table (PWT) wherever the PWT series is available, retaining TRADHIST only where PWT is missing. We construct the PWT series from PWT's GDP and expenditure-share variables and deflate it from 2017 to 2011 US dollars using a US GDP deflator.[26]

Predictors. The headline synthetic difference-in-differences specification projects three covariates out of the outcome before fitting the weights: logged pre-adoption population (Maddison), the logged human-capital index (PWT), and standardised inflation (Global Macro Database). Each covariate is interacted with a time trend.

Table A.1 summarises the data.

Missing data. Not every country-year cell is observed in the underlying sources. We address this conservatively rather than through model-based extrapolation. We first expand each country's series to a balanced 1945–1996 skeleton so that gap length can be measured correctly, then linearly interpolate only short *internal* gaps—at most two years for the annual macroeconomic series (GDP, population, trade, capital, total factor productivity (TFP), inflation) and at most nine years, in logit space, for the quinquennial Barro-Lee education shares. We never extrapolate beyond the first or last observed year of a series, and we never interpolate across a country's treatment year, so that no imputed value can smuggle post-treatment information into a pre-treatment predictor or vice versa. We flag every interpolated value. Panel construction stops if a treated country lacks a required $t-1$ outcome, population, or inflation value; analysis-side screens record and skip specifications lacking other required predictor support.

Sample restrictions. No economy in the UN M49 regions Eastern Asia or South-eastern Asia is a donor, and we exclude India by name. For each cohort and outcome, a donor may miss at most two raw outcome years between the window start and the year before adoption; the treated economy is exempt from this screen. A donor must have the outcome in every year of the window. It must also have at least two observed pre-adoption years of each covariate in the headline specification. For the factor outcomes (capital, capital per worker, employment per capita, and investment), estimation windows start no earlier than 1950.

---

[25] We do not use the Global Macro Database's GDP series as an outcome, since it is reported in national-currency units rather than a common cross-country numeraire.

[26] The two series use different sources and deflators. The headline specification interacts each economy's pre-adoption mean inflation with a linear time trend (Section 6.1), and full sets of estimates that use the PWT-derived series, or TRADHIST, for every economy show that the results are robust to this choice.

Table A.1. Variables and sources in the empirical analysis

Variable	Source
CPI inflation	Global Macro Database
Output and trade	
GDP per capita	Maddison Project Database
Exports per capita	TRADHIST, PWT and Maddison Project Database
Population and employment	
Population	Maddison Project Database
Employment per capita	PWT
Capital and investment	
Capital stock	PWT
Capital per worker	PWT
Investment share of GDP	PWT
Productivity and human capital	
Output per worker	PWT
Human capital index, via Princeton redistribution	PWT via Princeton redistribution
Human capital index, direct PWT source	PWT
Labour share of GDP	PWT
Constructed total factor productivity	Constructed from PWT inputs

Notes: Reported outcomes enter in logs. GDP per capita is in 2011 international dollars and is supplied directly by the Maddison Project Database, rather than divided locally by the separately listed population series. Exports per capita are in thousands of 2011 US dollars per person and use population from the Maddison Project Database; PWT supplies exports for Taiwan and Singapore where available. Population is in thousands. CPI inflation is the annual percentage change in the consumer price index. Employment per capita is PWT persons engaged divided by PWT population. PWT output and capital measures are at constant 2017 national prices. Capital stock is in millions of US dollars. Output and capital per worker use PWT persons engaged. Investment is a share of GDP and the human capital measures are indices. PWT is Penn World Table. Constructed total factor productivity is defined in Equation (2).

A.2 Policy coding

A.2.1 Coverage of the coding

Table A.2 reports, for each economy, the number of coded policy-sector rows, the distinct sectors and years they cover, and the count of each direction flag.

A.2.2 Policy change around the turn

Figure A.1 plots the adoption of export-promoting policies and the withdrawal of import-limiting policies in the comparison economies, each aligned on the date of its own export-orientation turn. Figure A.2 plots the share of sectors under export-promoting, import-limiting, or both kinds of policy in each economy and year.

A.2.3 Policy counts by country and type

Figures A.3 and A.4 show the frequency of policies by country and type over time. Figure A.5 shows the raw policy coding database.

A.2.4 Verifying and screening the coding against the corpus

We test the hand-coded database for two kinds of error, misclassified policies and omitted ones, with a retrieval-augmented language model built on the corpus we coded from. Table A.3 reports the corpus, the verification, our review, and the omission screen for each economy.

The corpus for each economy collects its academic literature and policy documents, with scanned pages converted to text by optical character recognition. We divide each document into overlapping passages of about 400 words, which the system indexes and retrieves. For each coded row, the system combines semantic and keyword search, reranks the results by relevance to the policy and its International Standard Industrial Classification (ISIC) sector, and keeps the eight highest-ranked passages. A language model then judges each of the row's eight coding fields against the row's own policy text and, separately, against the passages. The model is Claude Sonnet 4.6 for the treated economies and Claude Sonnet 5.5 for the comparison economies.

A judgement drawn from the literature must quote the passage it rests on. Code checks each quote against the cited passage, and we record a judgement whose quote fails as having no literature support. The literature supports or partly supports 90.7 per cent of verified rows and fails to support one, in Taiwan. Where the supported share is low, as for China, the retrieved passages do not discuss the policy.

We review every policy with a row in which the model judges a field incorrectly coded or a verified quote contradicts the coding. For each, we accept the model's proposed value, substitute our own, or keep the original; where one row recorded both the strengthening and the weakening of an instrument, we split it into two rows. The review recoded 55 of the 66 flagged policies. The review also removed duplicate rows, which leaves the 922 rows of Table A.2 from the 1,000 coded before review.

The omission screen inverts the search. We compare every passage with every coded policy text and keep the passages whose highest embedding similarity to any coded policy falls below 0.35. The model reads them in windows of up to three adjacent passages and proposes a coding when a window describes a policy. From 4,096 windows, it proposed eight candidates. After review we added two, the Korea Trade Promotion Corporation and Singapore's National Wages Council; Korea's second candidate duplicates the first.

Table A.2. Coverage of the industrial policy coding, 1949–1990

Country	Rows	Distinct ISIC3	Years	Most active year	Policy type			
					EXP	IMP	DOM	EXPBAN
Japan	163	24	1950–1983	1950 (30)	7	31	132	2
					(4.1%)	(18.0%)	(76.7%)	(1.2%)
Taiwan	182	37	1949–1990	1960 (33)	88	57	102	0
					(35.6%)	(23.1%)	(41.3%)	(0.0%)
South Korea	128	29	1954–1986	1970 (29)	79	34	62	0
					(45.1%)	(19.4%)	(35.4%)	(0.0%)
Singapore	142	47	1959–1972	1963 (63)	0	47	95	0
					(0.0%)	(33.1%)	(66.9%)	(0.0%)
China	44	30	1970–1990	1980 (13)	16	15	40	0
					(22.5%)	(21.1%)	(56.3%)	(0.0%)
India	124	54	1950–1978	1965 (36)	88	44	49	1
					(48.4%)	(24.2%)	(26.9%)	(0.5%)
Malaysia	47	22	1950–1979	1969 (10)	9	16	36	1
					(14.5%)	(25.8%)	(58.1%)	(1.6%)
Thailand	92	38	1950–1990	1962 (31)	23	75	75	1
					(13.2%)	(43.1%)	(43.1%)	(0.6%)
All	922	104	1949–1990	1963 (71)	310	319	591	5
					(25.3%)	(26.0%)	(48.2%)	(0.4%)

Notes: The unit of observation is a policy-sector row: one coded policy text paired with one 3-digit International Standard Industrial Classification (ISIC) industry. "Most active year" reports the year with the largest number of policy-sector rows, with that number in parentheses. A row may carry more than one policy-type flag (641 rows carry one, 259 two, and 22 three), so the four type columns sum to more than the row count; percentages in the second line are shares of all policy-type flags in that row.

Variables: "Distinct ISIC3" counts distinct 3-digit ISIC sector codes. EXP is export-promoting; IMP, import-limiting; DOM, domestic (not directly trade-related); and EXPBAN, export ban or restriction.

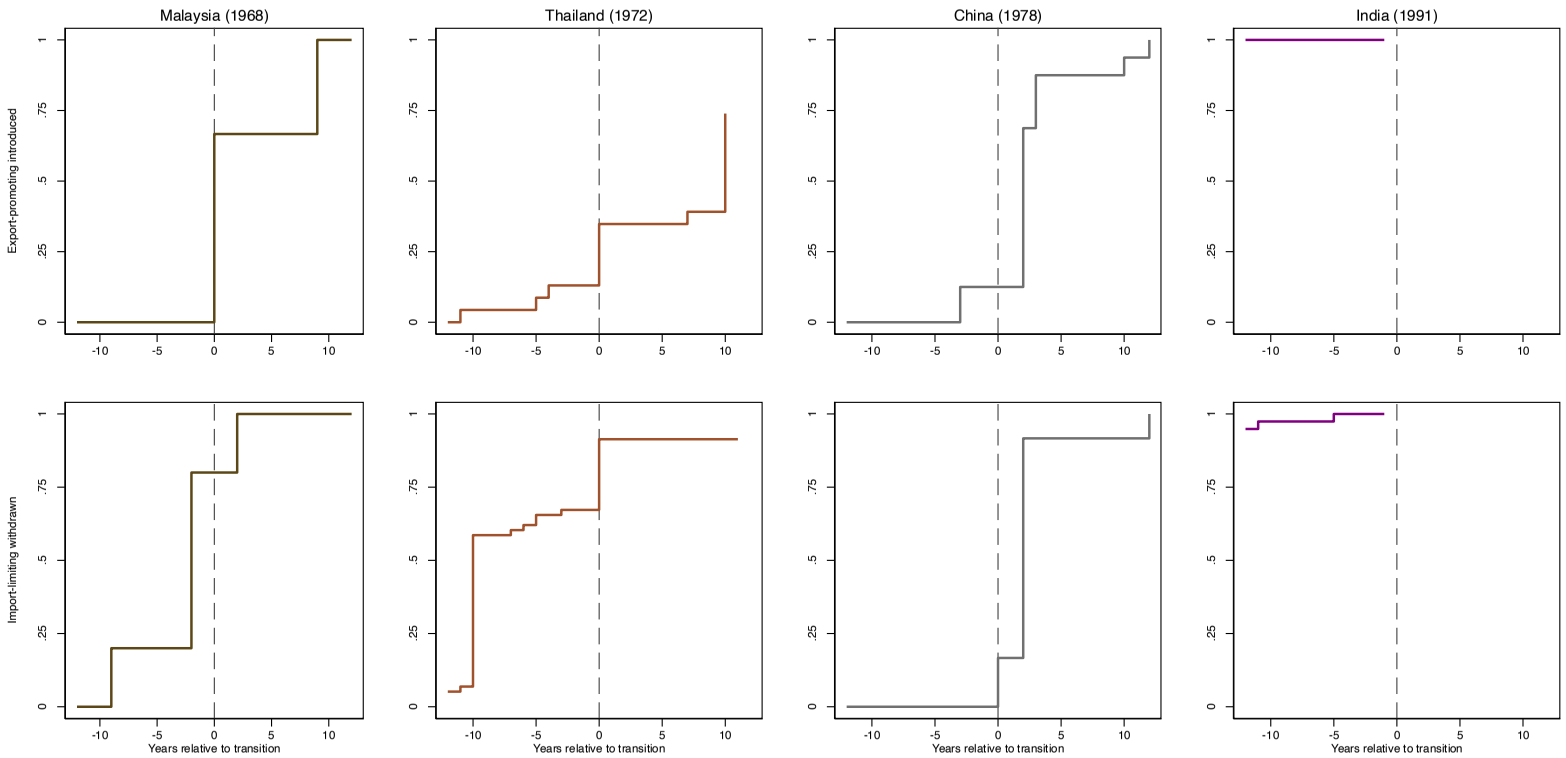


Figure A.1. Timing of trade policy change: the comparison economies, on their own transition dates

Notes: Each panel plots the cumulative share of a comparison economy's own export-promoting policies introduced (top row) or import-limiting policies withdrawn (bottom row) that had been recorded by year t, with the country aligned on its own transition date (vertical dashed line): Malaysia 1968, Thailand 1972, China 1978, India 1991. The denominator is the country's full total over the 1949–1990 sample window, so every series runs on a 0–1 scale irrespective of how many policies the country has. The series are step functions of the raw event dates, with no smoothing and no moving window. A curve need not start at 0 or end at 1 inside the plotted ±12-year range: events before $t - 12$ are already accumulated at the left edge, and events after $t + 12$ enter the denominator but are not yet accumulated at the right edge. The withdrawal panels cover *recorded* withdrawals only: a policy with no coded end year generates no withdrawal event and is absent from numerator and denominator alike. India's transition date lies just outside the sample window, so its panels show the pre-transition period only and its denominators exclude events after 1990; its export-promoting record is in any case complete by 1971, two decades before the date. China's record begins in 1970, truncating its pre-transition side at $t - 8$.

Table A.3. Verification and omission screen of the policy coding, by economy

	Treated economies				Comparison economies				
	Japan	Taiwan	South Korea	Singapore	China	India	Malaysia	Thailand	*All*
Panel A: Corpus									
Documents	124	98	138	82	28	21	63	142	696
Pages	7,201	5,127	14,192	5,648	1,836	1,635	1,863	8,865	46,367
Passages	7,743	5,469	14,527	5,872	1,793	2,395	2,278	9,311	49,388
Panel B: Verification									
Policy-sector rows	165	195	127	141	65	125	87	95	1,000
Distinct policies	112	56	39	3	30	29	19	93	381
Supported (%)	98.2	96.4	98.4	100.0	35.4	83.2	82.8	96.8	90.7
No evidence (%)	1.8	3.1	1.6	0.0	64.6	16.8	17.2	3.2	9.2
Quotes verified (%)	92.2	94.5	94.0	88.4	100.0	100.0	100.0	99.7	94.8
Panel C: Review									
Policies flagged	3	16	5	2	5	12	7	16	66
Policies recoded	2	10	3	0	5	12	7	16	55
Panel D: Omission screen									
Windows screened	299	148	1,433	924	106	207	361	618	4,096
Candidate policies	0	0	2	4	0	2	0	0	8
Policies added	0	0	1	1	0	0	0	0	2

Notes: All pools the eight economies; its shares are computed over pooled rows or quotes.

Sample: Policy-sector rows of the hand-coded dataset before review. A distinct policy is one coded policy text, which can span several sectors.

Variables: Pages exclude blank pages and reference matter. Passages are overlapping segments of about 400 words, the unit the system retrieves. Supported: share of rows whose passages support or partly support the coding; no evidence: share whose passages do not discuss the policy; one Taiwan row is not supported. Quotes verified: share of the model's quotes that match the cited passage at a similarity of at least 0.80. Flagged: at least one row with a field the model judges incorrect or a quote contradicts. Recoded: coding changed in review. Windows: groups of up to three adjacent passages with embedding similarity below 0.35 to every coded policy. Candidate policies: windows the model classifies as a policy; Korea's two are the same policy.

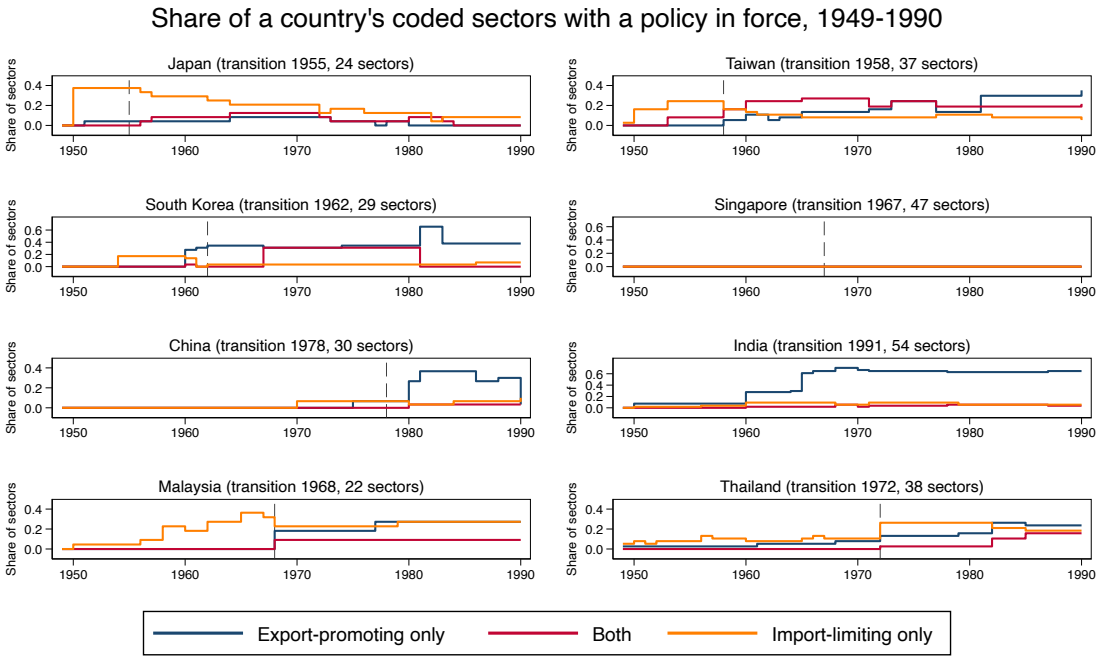


Figure A.2. Sectors with a policy in force, by combination, 1949–1990

Notes: Each panel plots, for one country and year, the share of its coded sectors in which a policy is in force, split into three mutually exclusive groups: sectors with an export-promoting policy and no import-limiting policy in force (navy), sectors with both in force (red), and sectors with an import-limiting policy and no export-promoting policy in force (orange). A policy is in force in year t if its coded start year is at most t and its coded end year at least t. A policy with no coded end year is treated as in force through 1990. The unit is the 3-digit International Standard Industrial Classification (ISIC) sector, and the denominator is the number of distinct sectors in which the country has at least one coded policy of any type introduced in the sample window, reported in each panel title. The vertical dashed line marks the country's transition date. The series are step functions of the raw coded years.

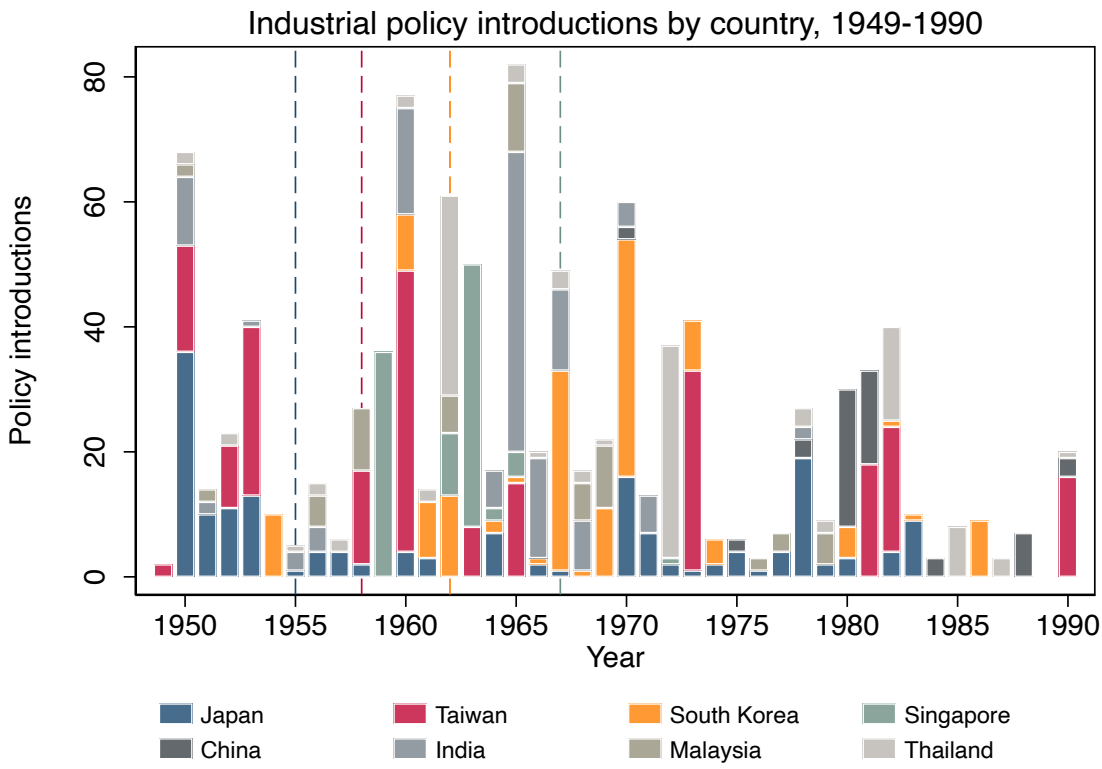


Figure A.3. Industrial policy introductions by country, 1949–1990

Notes: Bars stack the number of policy–sector rows each country switches on in a given year; counts are unsmoothed. The four early adopters are drawn in colour, with dashed vertical lines at their dated turns, and the four comparison economies in grey. Levels are not comparable across countries as they reflect how densely the underlying sources describe policy.

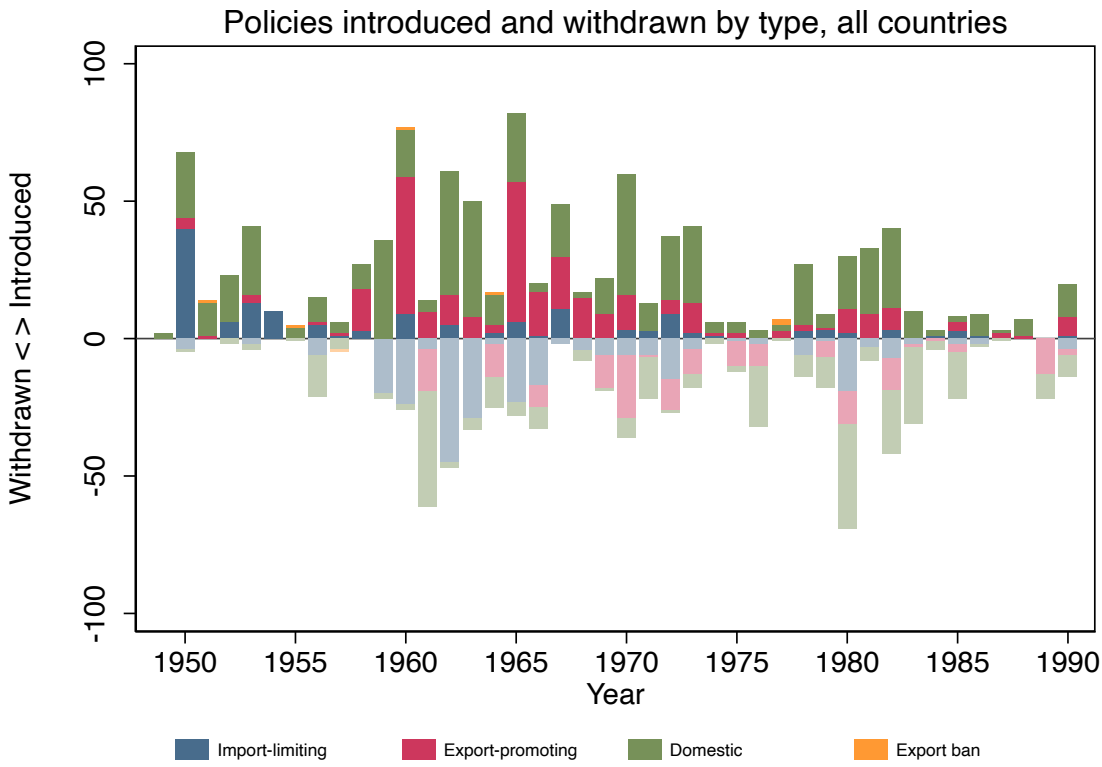


Figure A.4. Policy introductions and withdrawals by type, 1949–1990

Notes: Annual policy–sector event counts pooled across the eight countries and stacked by policy type; counts are unsmoothed. Introductions appear above zero and withdrawals below zero. A row can carry more than one policy-type flag and therefore enter more than one series. Rows without an end year generate no withdrawal event.

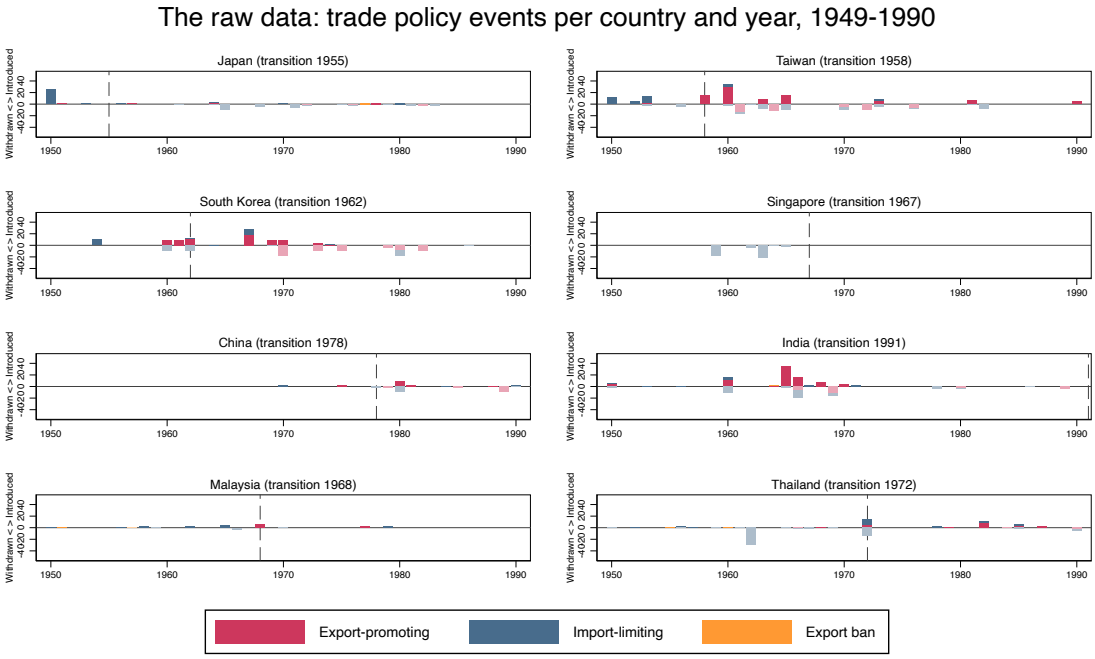


Figure A.5. Trade policy events by country and year, 1949–1990

Notes: Panels report unsmoothed annual counts for the eight countries. Export-promoting, import-limiting, and export-ban events are stacked by type; introductions appear above zero and withdrawals below zero in a paler shade. Dashed vertical lines mark each country's dated turn. Domestic-policy rows are not plotted.

B. Estimation Appendix

This appendix describes the implementation and diagnostics that support the main analysis.

B.1 Estimator

Synthetic difference-in-differences (SDID) fits unit weights ω to match the treated economy's pre-adoption path up to a level difference, and time weights λ to choose its pre-adoption baseline. For cohort c, let $\mathcal{D}_c$ be its donor pool of size n_c. Its pre- and post-adoption years are $\mathcal{T}_c^-$ and $\mathcal{T}_c^+$, with lengths T_c^- and T_c^+. Write $\widetilde{Y}_{it}$ for the covariate-adjusted outcome defined in Appendix B.2. With Δ_m denoting the non-negative weights of length m that sum to one, the unit-weight problem is (Arkhangelsky et al., 2021, eq. (4))

$$(\hat{\omega}_{0c}, \hat{\omega}_c) = \arg \min_{\omega_0 \in \mathbb{R},\, \omega \in \Delta_{n_c}} \left\{ \sum_{t \in \mathcal{T}_c^-} \left(\omega_0 + \sum_{i \in \mathcal{D}_c} \omega_i \widetilde{Y}_{it} - \widetilde{Y}_{ct} \right)^2 + \zeta_{\omega c}^2 T_c^- \|\omega\|_2^2 \right\}. \tag{7}$$

The intercept permits a level difference. The time-weight problem uses donors alone to match each donor's post-adoption average with its weighted pre-adoption outcomes (Arkhangelsky et al., 2021, eq. (6)):

$$(\hat{\lambda}_{0c}, \hat{\lambda}_c) = \arg \min_{\lambda_0 \in \mathbb{R},\, \lambda \in \Delta_{T_c^-}} \left\{ \sum_{i \in \mathcal{D}_c} \left(\lambda_0 + \sum_{t \in \mathcal{T}_c^-} \lambda_t \widetilde{Y}_{it} - \frac{1}{T_c^+} \sum_{t \in \mathcal{T}_c^+} \widetilde{Y}_{it} \right)^2 + \zeta_{\lambda c}^2 n_c \|\lambda\|_2^2 \right\}. \tag{8}$$

We use the package's default regularisation: $\zeta_{\omega c}$ scales the unit-weight penalty, while $\zeta_{\lambda c}$ supplies the numerical ridge term the package adds to the time-weight problem. The fitted donor and pre-adoption weights, together with weight one on the treated economy and uniform weights over post-adoption years, enter the weighted two-way fixed-effects regression in Equation (9) for $\hat{\tau}_c$, the cohort effect.

For cohort c, we also write c for its treated economy. The fitted vectors $\hat{\omega}_c$ and $\hat{\lambda}_c$ supply the donor and pre-adoption weights; extend them to give the treated economy weight $\hat{\omega}_{cc} = 1$ and each post-adoption year weight $\hat{\lambda}_{tc} = 1/T_c^+$. The treatment indicator W_{it} equals one for economy c from adoption onward and zero otherwise; μ_c is an intercept, while α_{ic} and β_{tc} are unit and year fixed effects. We estimate the cohort effect τ_c in the form of Arkhangelsky et al. (2021, eq. (1)):

$$(\hat{\tau}_c, \hat{\mu}_c, \hat{\alpha}_c, \hat{\beta}_c) = \arg \min_{\tau_c, \mu_c, \alpha_c, \beta_c} \sum_{i \in \mathcal{D}_c \cup \{c\}} \sum_{t \in \mathcal{T}_c^- \cup \mathcal{T}_c^+} \left(\widetilde{Y}_{it} - \mu_c - \alpha_{ic} - \beta_{tc} - \tau_c W_{it} \right)^2 \hat{\omega}_{ic} \hat{\lambda}_{tc}. \tag{9}$$

Let a_c be economy c's adoption year, t the calendar year, and $h = t - a_c$ event time, with $h = 0$ denoting the adoption year. The estimated event-year effect is the treated-minus-synthetic gap in year $a_c + h$, net of the time-weighted pre-adoption gap:

$$\hat{\tau}_c(h) = \left(\widetilde{Y}_{c,a_c+h} - \sum_{i \in \mathcal{D}_c} \hat{\omega}_{ic} \widetilde{Y}_{i,a_c+h} \right) - \sum_{t \in \mathcal{T}_c^-} \hat{\lambda}_{tc} \left(\widetilde{Y}_{ct} - \sum_{i \in \mathcal{D}_c} \hat{\omega}_{ic} \widetilde{Y}_{it} \right). \tag{10}$$

Averaging these annual effect estimates over the adoption year and the following 20 years gives $\hat{\tau}_c$ in Equation (9), whose post-adoption weights are uniform. Equation (1) defines the cohort average, the average across cohorts, and the horizon profile from these same annual estimates.

The anchor-year column of Table B.1 reports where the time weights concentrate. This concentration follows from how SDID chooses the time weights $\hat{\lambda}$. SDID fits them so that the $\hat{\lambda}$-weighted pre-adoption average of the donors' outcomes tracks the donors' post-adoption average. With only five to seven pre-adoption years of strongly trending series, the most recent year predicts the post-adoption average best, so the fit puts the weight there. Online Appendix OA-B reports the weights and the levels plots, which show the consequence: the synthetic line is aligned on the $\hat{\lambda}$-weighted pre-adoption gap; where $\hat{\lambda}$ is on $t-1$, this is the gap at $t-1$.

With a common estimation window, the staggered estimator is equivalent to estimating each adoption cohort separately and averaging its effect with the estimator's aggregation weights (Clarke, Pailañir, Athey, and Imbens, 2024, section 2.3). Our cohort-specific windows are the one deliberate departure from that computation.[27] We implement each cohort estimate with the `synthdid` R package, version 0.0.9 (Arkhangelsky et al., 2024).

B.2 Covariate adjustment

Population, human capital, and a macroeconomic variable are standard predictors of the outcomes we study. We use their pre-adoption values to avoid conditioning on variables that treatment could itself change; their definitions and sources are in Appendix A.1. We remove their fitted contribution by residualisation, following Kranz (2022) and Clarke, Pailañir, Athey, and Imbens (2024).

Each covariate enters as its unit's pre-adoption mean multiplied by calendar year. A pre-adoption mean alone is constant within a unit and would be absorbed by unit fixed effects; interacting it with time lets it account for differences in trends (Xu, 2017, Remark 3).

Within each cohort window, we regress the outcome on these covariate-by-year terms and unit and year fixed effects, using never-treated units only. Let $\bar{X}^{-}_{ic}$ collect unit i's pre-adoption covariate means for cohort c, and let $\hat{\gamma}_c$ be the slopes from that regression. We subtract the fitted covariate contribution from every unit, including the treated economy:

$$\widetilde{Y}_{it} = Y_{it} - (t\bar{X}^{-}_{ic})'\hat{\gamma}_c. \tag{11}$$

We retain the unit and year fixed effects in the outcome and fit both weight vectors to $\widetilde{Y}_{it}$.

Residualisation estimates the slopes on untreated outcomes before fitting the weights. Fitting slopes and weights jointly can make the covariate adjustment sensitive to the optimisation and can fail to recover the covariate–outcome relationship in some settings. The fixed-effects regression provides the separate adjustment proposed by Kranz (2022) and discussed by Clarke, Pailañir, Athey, and Imbens (2024).

Every placebo assignment repeats the residualisation on its own untreated sample before fitting new unit and time weights. We exclude the placebo economy from the slope regression, just as we exclude the actual treated economy. This carries the covariate adjustment through the inference procedure in Appendix C.1.

We choose the linear interaction for its pre-period fit. Appendix E.3 reports the quadratic and year-dummy alternatives.

[27] Per-cohort estimation is also called block estimation.

B.3 Donor pools

For GDP per capita, the donor pools contain 41, 60, 77, and 82 economies for Japan, Taiwan, South Korea, and Singapore, respectively. For exports per capita, the corresponding sizes are 39, 57, 73, and 79.

Japan's donor pool is the smallest of the four cohorts for both core outcomes. Unit fixed effects absorb level differences, so the donors are asked to match Japan's path rather than its income. Japan's fitted donor weights for both outcomes are in Online Appendix OA-B.

B.4 Fit and weighting diagnostics

The credibility of the estimate rests on both weight vectors. We report the fit and the weights directly rather than summarising them, so that the reader can see which donors and which pre-treatment years the comparison rests on.

For reasons of space, Online Appendix OA-B reports the estimated time weights and donor weights for GDP per capita and exports per capita, alongside those for the factor outcomes. The time weights determine how heavily each pre-adoption year enters the comparison; the donor weights underlie each cohort's synthetic control. In the donor-weight tables, a dash marks a donor that was ineligible for that cohort under the sample screens, while `0.000` marks an eligible donor whose fitted weight rounds to zero. The distinction matters, because the first reflects the design and the second reflects the optimisation.

Table B.1 summarises the fit and weighting diagnostics for every cohort and outcome in one place.

Online Appendix OA-B shows the levels trajectory for each cohort against its weighted synthetic control, for GDP per capita and exports per capita. These plots are the primary visual checks on pre-treatment fit.

Table B.1. Fit and weighting diagnostics by cohort and outcome

Cohort	Outcome	Pre-periods	Donors	Max omega	Effective donors	Anchor year	Effective pre-periods	Pre-fit SD ratio
Japan	GDP	6	41	0.207	7.195	1954	1.000	0.259
Taiwan	GDP	6	60	0.062	36.063	1957	1.000	0.114
South Korea	GDP	7	77	0.056	42.870	1961	1.000	0.311
Singapore	GDP	7	82	0.089	31.712	1966	1.000	0.321
Japan	Exports	6	39	0.133	14.159	1954	1.025	0.396
Taiwan	Exports	6	57	0.073	23.506	1957	1.000	0.674
South Korea	Exports	7	73	0.194	9.175	1961	1.000	0.574
Singapore	Exports	7	79	0.170	8.646	1966	1.143	0.670
Japan	Employment per capita	5	43	0.225	7.849	1954	1.000	0.404
Taiwan	Employment per capita	6	49	0.224	7.358	1957	1.000	0.247
South Korea	Employment per capita	7	57	0.035	46.070	1960	2.608	0.259
Singapore	Employment per capita	7	71	0.076	25.758	1965	1.645	0.119
Japan	Capital	5	47	0.045	36.633	1954	1.000	0.076
Taiwan	Capital	6	53	0.060	36.312	1957	1.000	0.052
South Korea	Capital	7	61	0.304	5.255	1961	1.000	0.491
Singapore	Capital	7	81	0.037	52.009	1966	1.000	0.068
Japan	Capital per worker	5	43	0.086	21.929	1954	1.000	2.271
Taiwan	Capital per worker	6	49	0.045	41.761	1957	1.000	0.391
South Korea	Capital per worker	7	57	0.507	3.098	1961	1.000	0.171
Singapore	Capital per worker	7	71	0.039	47.616	1966	1.000	0.283
Japan	Investment	5	47	0.075	22.749	1954	1.868	0.444
Taiwan	Investment	6	53	0.154	19.447	1957	1.769	0.502
South Korea	Investment	7	60	0.102	24.040	1961	2.720	0.564
Singapore	Investment	7	80	0.052	46.601	1966	3.036	0.360

Notes: The table reports unit-weight, time-weight and pre-adoption fit diagnostics of the synthetic difference-in-differences (SDID) fits, one row for each cohort and outcome.

Sample: Rows cover Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967) for six outcomes: GDP, Exports, Employment per capita, Capital, Capital per worker, and Investment. Cohort windows end in 1975, 1978, 1982, and 1987, respectively. Donors exclude the economies of Eastern Asia and South-eastern Asia, and India.

Variables: All six outcomes are in logs; capital is the PWT 10.x `rnna` stock at constant 2017 national prices, and capital per worker divides it by same-vintage PWT persons engaged (`emp`). Investment is the investment share of GDP. "Pre-periods" counts pre-adoption years, and "Donors" counts eligible donor units. "Max omega" is the largest unit weight; "Effective donors" is the inverse of the sum of squared unit weights. "Anchor year" is the earliest year attaining the maximum time weight; "Effective pre-periods" is the inverse of the sum of squared time weights. "Pre-fit SD ratio" is the sample standard deviation of the pre-adoption treated–synthetic gap divided by that of the covariate-adjusted treated outcome.

Method: The diagnostics come from the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted (see Section 6.1).

C. Results Appendix

This appendix sets out the inference procedures (Appendix C.1), then reports the average effects (Appendix C.2) and cohort effects (Appendix C.3).

C.1 Inference procedures

For pooled inference we set aside the real treated economies and sample uniformly from valid joint placebo assignments: each assigns the four adoption dates to four distinct donors, each eligible for its cohort slot. All assigned donors leave every control pool. Within each cohort's window, we repeat the residualisation on the remaining controls, fit new unit and time weights, and estimate the placebo cohort effect. We then recompute the same pooled statistic, with the same cohort weights or horizon as the observed statistic, following the staggered placebo procedure of Clarke, Pailañir, Athey, and Imbens (2024). For any reported pooled statistic S, let $\hat{S}$ be its observed value, S_r^* its value in assignment r, R the number of seeded assignments per outcome, here 1,000, and $\bar{S}^* = R^{-1}\sum_{r=1}^{R} S_r^*$. The placebo standard error and two-sided p-value are

$$\widehat{\mathrm{SE}}(S) = \left[\frac{1}{R}\sum_{r=1}^{R}(S_r^* - \bar{S}^*)^2\right]^{1/2}, \qquad p_S = \frac{1 + k_S}{1 + R}, \quad k_S = \sum_{r=1}^{R}\mathbf{1}\{|S_r^*| \geq |\hat{S}|\}. \tag{12}$$

The divisor is R, as in Arkhangelsky et al. (2021, Algorithm 4), and the added one follows Phipson and Smyth (2010). This is not an exact test: the actual treated economies are excluded from the placebo assignments and were not randomly assigned treatment, so the observed statistic is not interchangeable with its placebo reference. The simulation range is the 95 per cent Clopper–Pearson interval for the exceedance probability, computed from k_S and R. If that range, evaluated before rounding, includes 0.05, we make no five per cent claim.

For cohort inference we instead run every eligible donor once as the placebo, against the cohort's pool minus itself, with the real treated economies excluded. Each run repeats the residualisation and fits both sets of weights anew. If $\hat{\tau}_c$ is the observed cohort effect, $\hat{\tau}_{cj}^*$ is donor j's placebo effect, n_c is the number of eligible donors, and $\bar{\hat{\tau}}_c^* = n_c^{-1}\sum_{j=1}^{n_c}\hat{\tau}_{cj}^*$, then

$$\widehat{\mathrm{SE}}(\hat{\tau}_c) = \left[\frac{1}{n_c}\sum_{j=1}^{n_c}(\hat{\tau}_{cj}^* - \bar{\hat{\tau}}_c^*)^2\right]^{1/2}, \qquad p_c = \frac{1 + k_c}{1 + n_c}, \quad k_c = \sum_{j=1}^{n_c}\mathbf{1}\{|\hat{\tau}_{cj}^*| \geq |\hat{\tau}_c|\}. \tag{13}$$

The standard-error divisor is n_c, not $n_c - 1$. The p-value is the observed effect's absolute rank, $1 + k_c$, divided by the pool including the treated economy, $1 + n_c$; we print both rank and pool beside it in Table C.3 (Abadie, Diamond, and Hainmueller, 2010; Abadie, Diamond, and Hainmueller, 2015; Abadie, 2021). This gives each cohort a standalone reference using every donor, with no simulation error and a resolution set by the pool size. It does not make the test exact. Where the approximate Gaussian interval excludes zero but the rank p-value is not below 0.05, the p-value governs and no cell is marked. The disagreeing cells are Japan employment per capita, South Korea capital, South Korea investment, and Taiwan investment.

Jackknife inference requires at least two treated units per treatment period, whereas each of our cohorts contains one. We reject the cluster bootstrap because each of the four treated economies is absent from a resample with probability about 0.37, and about 84 per cent of accepted replicates omit at least one cohort. Those replicates estimate a different pooled quantity.

The *cross-cohort correlation bound* is the pooled placebo standard error that would result if the four cohorts' estimation errors were perfectly correlated: the weighted average of the cohort placebo standard errors for the same statistic. The correlation is between estimation errors, not outcomes. The bound holds only if the cohort standard errors are themselves valid.

C.2 Average effects

Table C.1 gives each pooled estimate with its placebo standard error, approximate interval, placebo p-value, simulation range, and cross-cohort correlation bound. Table C.2 gives the unadjusted and Holm-adjusted p-values of the headline family.

C.3 Cohort effects

Table C.3 gives each cohort effect with its placebo standard error, approximate interval, and rank placebo p-value. Figure C.1 plots cohort-level effect paths for GDP per capita with pointwise bands, and Figure C.2 gives the corresponding paths for exports per capita.

Table C.1. Pooled effects with placebo inference, by outcome

Outcome and statistic	Estimate	Placebo SE	Approximate 95% interval	Placebo p-value	Simulation range	Correlation bound
Panel A: Core outcomes						
GDP per capita						
Average cohort effect	0.449	(0.080)	[0.292, 0.606]	0.001	[0.000, 0.004]	0.153
h = 10	0.417	(0.092)	[0.237, 0.597]	0.001	[0.000, 0.004]	0.174
h = 15	0.707	(0.125)	[0.462, 0.953]	0.001	[0.000, 0.004]	0.237
h = 20	0.868	(0.154)	[0.566, 1.171]	0.001	[0.000, 0.004]	0.294
Exports per capita						
Average cohort effect	1.356	(0.178)	[1.007, 1.705]	0.001	[0.000, 0.004]	0.350
h = 10	1.317	(0.207)	[0.911, 1.723]	0.001	[0.000, 0.004]	0.407
h = 15	2.148	(0.276)	[1.607, 2.690]	0.001	[0.000, 0.004]	0.542
h = 20	2.425	(0.328)	[1.782, 3.068]	0.001	[0.000, 0.004]	0.648
Panel B: Additional outcomes						
Employment per capita						
Average cohort effect	0.132	(0.032)	[0.070, 0.194]	0.001	[0.000, 0.004]	0.061
h = 10	0.128	(0.033)	[0.063, 0.194]	0.001	[0.000, 0.004]	0.064
h = 15	0.234	(0.049)	[0.138, 0.330]	0.001	[0.000, 0.004]	0.095
h = 20	0.213	(0.064)	[0.086, 0.339]	0.001	[0.000, 0.004]	0.130
Capital						
Average cohort effect	0.489	(0.084)	[0.325, 0.653]	0.001	[0.000, 0.004]	0.161
h = 10	0.437	(0.091)	[0.258, 0.615]	0.001	[0.000, 0.004]	0.177
h = 15	0.761	(0.134)	[0.499, 1.022]	0.001	[0.000, 0.004]	0.255
h = 20	1.079	(0.163)	[0.760, 1.399]	0.001	[0.000, 0.004]	0.310
Capital per worker						
Average cohort effect	0.398	(0.089)	[0.223, 0.573]	0.001	[0.000, 0.004]	0.176
h = 10	0.353	(0.102)	[0.152, 0.553]	0.001	[0.000, 0.004]	0.200
h = 15	0.586	(0.138)	[0.315, 0.857]	0.001	[0.000, 0.004]	0.271
h = 20	0.946	(0.168)	[0.617, 1.275]	0.001	[0.000, 0.004]	0.330
Investment						
Average cohort effect	0.468	(0.136)	[0.201, 0.736]	0.002	[0.000, 0.006]	0.276
h = 10	0.401	(0.182)	[0.044, 0.757]	0.031	[0.020, 0.043]	0.364
h = 15	0.666	(0.196)	[0.281, 1.051]	0.002	[0.000, 0.006]	0.389
h = 20	0.575	(0.204)	[0.176, 0.975]	0.004	[0.001, 0.009]	0.404

Notes: The table reports pooled synthetic difference-in-differences (SDID) estimates in log points, with placebo standard errors (SEs) in parentheses; the approximate 95 per cent interval is the estimate plus or minus 1.96 SE. Claims about how unusual an effect is rest on the placebo p-value, not on the interval; where the p-value's simulation range includes 0.05, no five per cent claim is made.

Variables: The average cohort effect is the mean of the four cohort effects, each averaged over event times 0 to 20, where event time 0 is the adoption year. The cohorts are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). The h rows give the average cohort effect at event time h. All outcomes are in logs. The simulation range is the 95 per cent Clopper–Pearson interval for the probability that a placebo statistic is at least as large in absolute value as the observed one. The correlation bound is the placebo SE if the cohorts' estimation errors were perfectly correlated: the average of the cohorts' placebo SEs from the same joint placebo assignments, for the same statistic.

Method: The placebo SE is the standard deviation of the statistic across its placebos, taken over joint placebo assignments, each giving the four adoption dates to four different eligible donors. A two-sided placebo p-value is $(1 + k)/(1 + R)$, where R is the number of joint placebo assignments per outcome, 1,000 here, and k of them give a statistic at least as large in absolute value as the observed one.

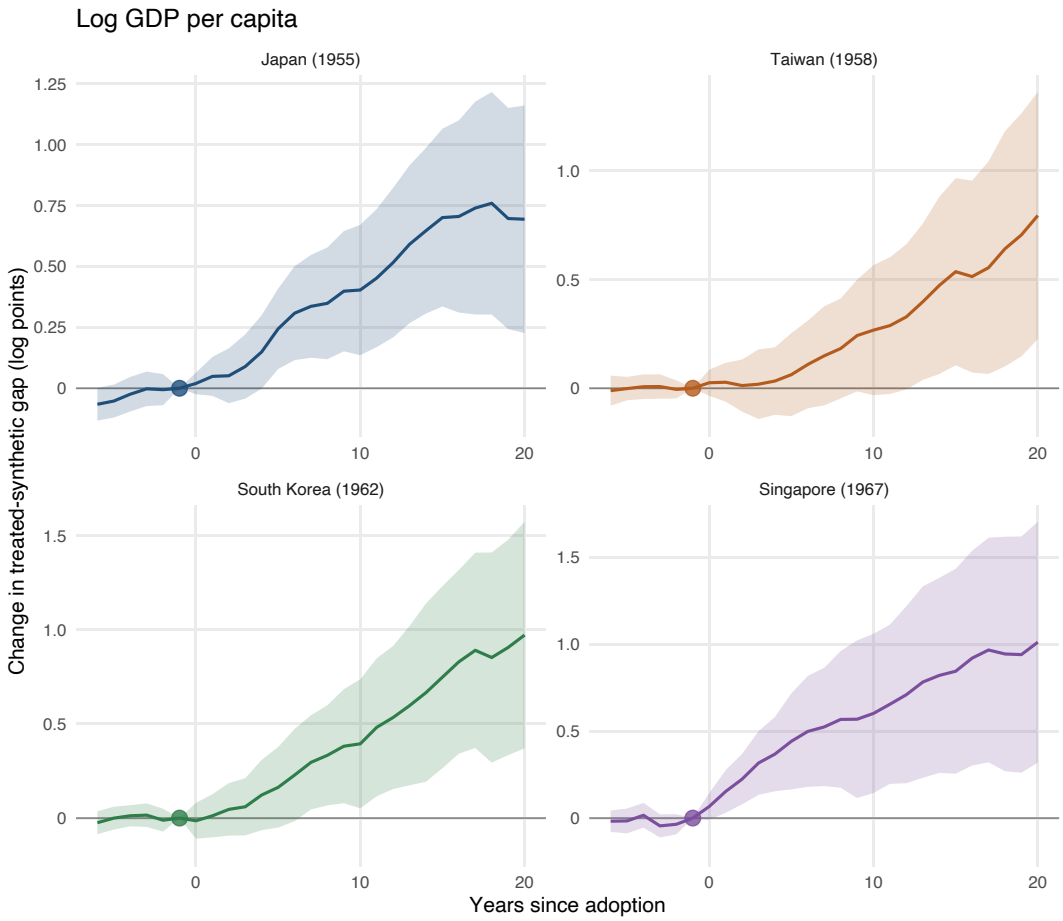


Figure C.1. Cohort-level effect paths: log GDP per capita

Notes: Each panel plots a cohort-specific synthetic difference-in-differences (SDID) effect on log GDP per capita, in log points, against years since adoption, where event time 0 is the adoption year: the change in the treated–synthetic gap relative to its fitted time-weighted pre-adoption baseline, under the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted (see Section 6.1). Panels are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). Cohort windows begin in 1949, 1952, 1955, and 1960 and end in 1975, 1978, 1982, and 1987, respectively; each panel plots event times −6 to 20, the span supported in every window. Donors exclude the economies of Eastern Asia and South-eastern Asia, and India. Dots on the zero line identify plotted pre-adoption years with positive fitted time weight, and a larger dot marks a larger weight. The shaded band is a pointwise approximate 95 per cent magnitude band: the effect plus or minus 1.96 standard deviations of the cohort's placebo effects that year, from runs that use every eligible donor once as the placebo. Each placebo path is measured from its own fit's time-weighted baseline; the pre-adoption band is a fit diagnostic, not a pre-trend test. Each panel has its own horizontal and vertical scales.

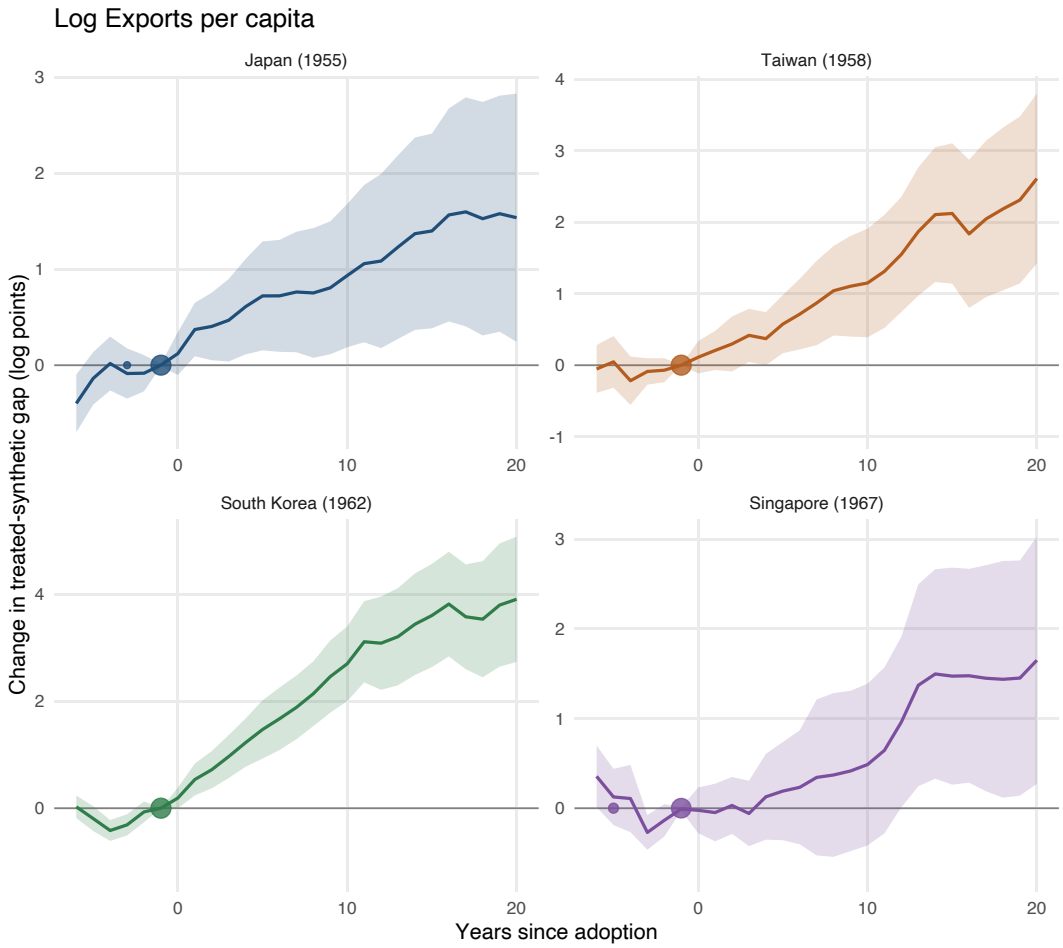


Figure C.2. Cohort-level effect paths: log exports per capita

Notes: Each panel plots a cohort-specific synthetic difference-in-differences (SDID) effect on log exports per capita, in log points, against years since adoption, where event time 0 is the adoption year: the change in the treated–synthetic gap relative to its fitted time-weighted pre-adoption baseline, under the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted (see Section 6.1). Panels are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). Cohort windows begin in 1949, 1952, 1955, and 1960 and end in 1975, 1978, 1982, and 1987, respectively; each panel plots event times −6 to 20, the span supported in every window. Donors exclude the economies of Eastern Asia and South-eastern Asia, and India. Dots on the zero line identify plotted pre-adoption years with positive fitted time weight, and a larger dot marks a larger weight. The shaded band is a pointwise approximate 95 per cent magnitude band: the effect plus or minus 1.96 standard deviations of the cohort's placebo effects that year, from runs that use every eligible donor once as the placebo. Each placebo path is measured from its own fit's time-weighted baseline; the pre-adoption band is a fit diagnostic, not a pre-trend test. Each panel has its own horizontal and vertical scales.

Table C.2. Headline family: unadjusted and Holm-adjusted placebo p-values

Outcome	Statistic	Placebo p-value	Holm-adjusted placebo p-value
GDP per capita	Average cohort effect	0.001	0.002
Exports per capita	Average cohort effect	0.001	0.002

Notes: The table reports the placebo p-values of the headline family, unadjusted and Holm-adjusted. The headline family is the average cohort effect on GDP per capita and on exports per capita; all other p-values are unadjusted.

Variables: The average cohort effect is the mean of the four cohort effects, each averaged over event times 0 to 20, where event time 0 is the adoption year. The outcomes are in logs.

Method: A two-sided placebo p-value is $(1 + k)/(1 + R)$, where R is the number of joint placebo assignments per outcome, 1,000 here, and k of them give a statistic at least as large in absolute value as the observed one. The Holm step-down adjustment over the family's two tests multiplies the smallest placebo p-value by the number of tests, the next smallest by one fewer, and so on; it then takes the running maximum in that order and caps each value at one.

Table C.3. Cohort effects with placebo inference, by outcome

Cohort (adoption year)	Estimate	Placebo SE	Approximate 95% interval	Rank placebo p-value	Rank and pool
Panel A: Core outcomes					
GDP per capita					
Japan (1955)	0.424	(0.129)	[0.171, 0.677]	0.048	2 of 42
Taiwan (1958)	0.303	(0.143)	[0.023, 0.583]	0.049	3 of 61
South Korea (1962)	0.453	(0.150)	[0.159, 0.746]	0.013	1 of 78
Singapore (1967)	0.615	(0.197)	[0.228, 1.002]	0.012	1 of 83
Exports per capita					
Japan (1955)	0.982	(0.368)	[0.261, 1.702]	0.025	1 of 40
Taiwan (1958)	1.276	(0.328)	[0.634, 1.919]	0.017	1 of 58
South Korea (1962)	2.429	(0.302)	[1.837, 3.021]	0.014	1 of 74
Singapore (1967)	0.737	(0.383)	[-0.013, 1.487]	0.050	4 of 80
Panel B: Additional outcomes					
Employment per capita					
Japan (1955)	0.103	(0.049)	[0.006, 0.199]	0.114	5 of 44
Taiwan (1958)	0.033	(0.048)	[-0.060, 0.126]	0.440	22 of 50
South Korea (1962)	0.180	(0.068)	[0.046, 0.313]	0.034	2 of 58
Singapore (1967)	0.213	(0.076)	[0.063, 0.363]	0.014	1 of 72
Capital					
Japan (1955)	0.461	(0.146)	[0.175, 0.748]	0.021	1 of 48
Taiwan (1958)	0.498	(0.144)	[0.215, 0.780]	0.019	1 of 54
South Korea (1962)	0.380	(0.167)	[0.054, 0.707]	0.065	4 of 62
Singapore (1967)	0.617	(0.192)	[0.241, 0.993]	0.024	2 of 82
Capital per worker					
Japan (1955)	0.425	(0.161)	[0.110, 0.740]	0.023	1 of 44
Taiwan (1958)	0.387	(0.164)	[0.067, 0.708]	0.040	2 of 50
South Korea (1962)	0.303	(0.187)	[-0.064, 0.670]	0.103	6 of 58
Singapore (1967)	0.477	(0.197)	[0.091, 0.863]	0.042	3 of 72
Investment					
Japan (1955)	0.471	(0.278)	[-0.074, 1.017]	0.146	7 of 48
Taiwan (1958)	0.535	(0.232)	[0.081, 0.989]	0.056	3 of 54
South Korea (1962)	0.541	(0.253)	[0.046, 1.036]	0.098	6 of 61
Singapore (1967)	0.326	(0.324)	[-0.308, 0.961]	0.272	22 of 81

Notes: Each estimate is a cohort's synthetic difference-in-differences (SDID) effect averaged over event times 0 to 20, where event time 0 is the adoption year, in log points; parentheses contain placebo standard errors (SEs). The approximate 95 per cent interval is the estimate plus or minus 1.96 SE. Claims about how unusual an effect is rest on the placebo p-value, not on the interval, and no cell is marked where the two disagree.

Sample: Each row is one treated economy for one outcome: GDP per capita and exports per capita in Panel A, and employment per capita, capital, capital per worker, and investment in Panel B. Donor pools are cohort- and outcome-specific. For Japan, Taiwan, South Korea, and Singapore, respectively, cohort windows begin in 1949, 1952, 1955, and 1960 in Panel A and 1950, 1952, 1955, and 1960 in Panel B, and end in 1975, 1978, 1982, and 1987.

Variables: All outcomes are in logs; capital is the PWT 10.x `rnna` stock at constant 2017 national prices, and capital per worker divides it by same-vintage PWT persons engaged (`emp`). Investment is the investment share of GDP. "Rank and pool" gives the rank of the observed effect in absolute value among itself and the placebo effects, and the number ranked.

Method: Estimates are for the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted (see Section 6.1). Each cohort effect is compared with runs that use every eligible donor once as the placebo, each against the cohort's pool minus itself. A placebo SE is the standard deviation of the statistic across its placebos; the rank placebo p-value is $(1 + k)/(1 + n)$, where k of the n eligible donors have a placebo effect at least as large in absolute value as the observed effect, that is, the rank divided by the pool.

D. Factor Accumulation and Accounting Appendix

D.1 Factor responses

Figure D.1 shows the cohort effects on the four factor outcomes.

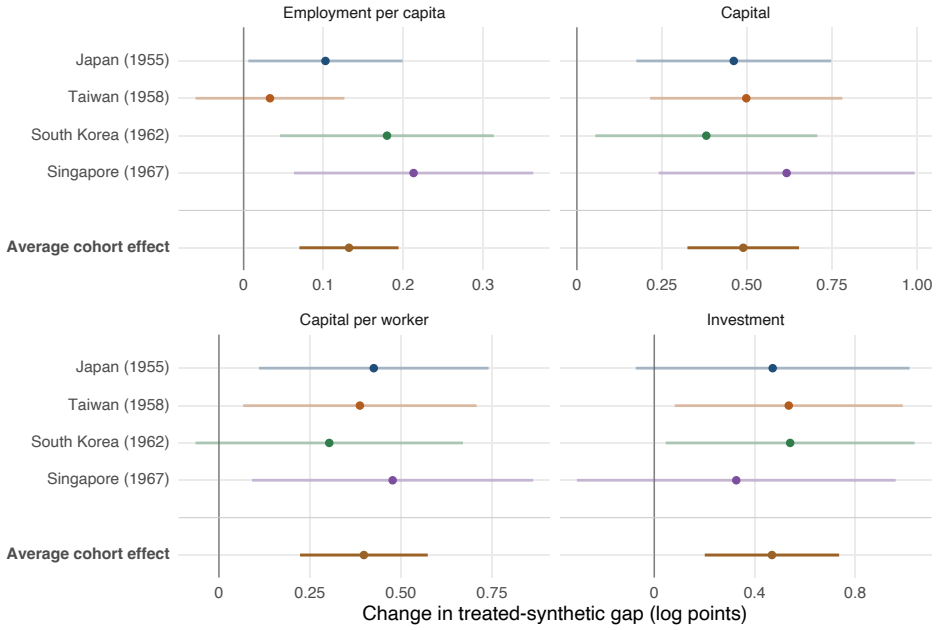


Figure D.1. Cohort effects and their average, with approximate intervals: factor outcomes

Notes: Each panel reports synthetic difference-in-differences (SDID) effects on one factor outcome, in log points. The panels are employment per capita, capital, capital per worker, and investment, in that order. All four outcomes are logged; capital is the Penn World Table (PWT) 10.0 `rnna` stock at constant 2017 national prices, capital per worker divides it by same-vintage PWT persons engaged (`emp`), and investment is the investment share of GDP. The effects use the headline specification, which projects logged population, logged human capital index, and standardized inflation out of the outcome before the weights are fitted. The cohorts are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). For each cohort, the dot is its average effect over event times 0 to 20; below the separator is the average cohort effect, the mean of the four cohort effects. Dots are estimates and whiskers are approximate 95 per cent magnitude intervals, the estimate plus or minus 1.96 placebo standard errors. A placebo standard error is the standard deviation of the statistic across its placebos: for a cohort, runs that use every eligible donor once as the placebo; for the average cohort effect, 1,000 joint placebo assignments. The cohorts' whiskers are lighter than the average cohort effect's. Each panel has its own horizontal scale, and the vertical line marks zero. The whiskers show magnitude and are not tests; the placebo p-values for these outcomes are in Table C.1 and Table C.3. The factor outcomes lie outside the headline family (the average cohort effect on GDP per capita and on exports per capita), whose placebo p-values carry a Holm adjustment.

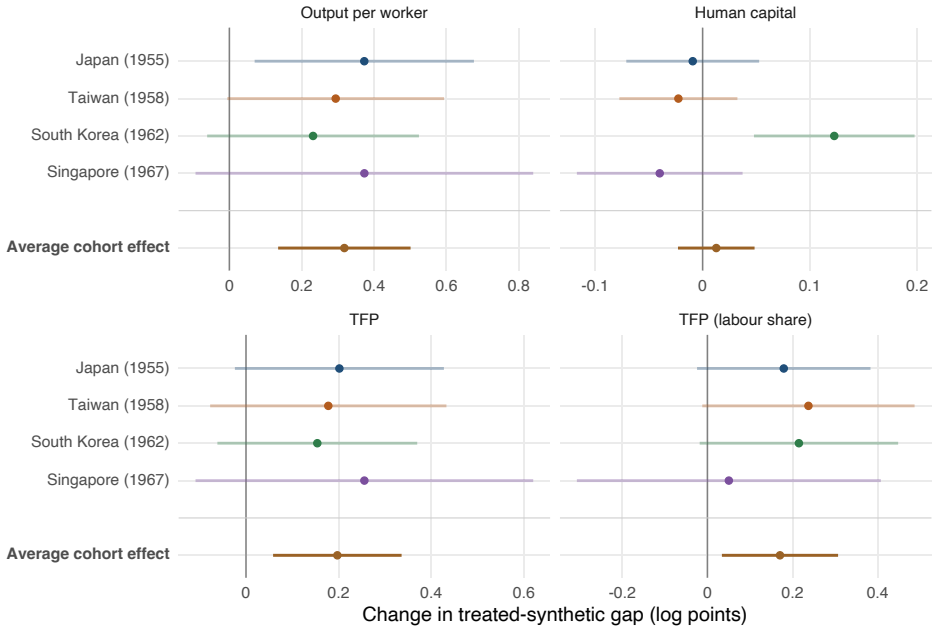


Figure D.2. Cohort effects and their average, with approximate intervals: growth-accounting outcomes

Notes: Each panel reports synthetic difference-in-differences (SDID) effects on one growth-accounting outcome, in log points. The panels are output per worker, human capital, TFP, and TFP (labour share), in that order, and each has its own horizontal scale; the vertical line marks zero. The effects are fitted with no controls: the specification projects no covariates out of the outcome (see Section 8). The effects are separately estimated, each with its own unit weights, time weights, window and donor pool, and are not an additive decomposition of the estimated output effect. The cohorts are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). Their windows begin in 1950, 1952, 1955, and 1960, respectively, and end 20 years after adoption. For each cohort the dot is its average effect over event times 0 to 20, where event time 0 is the adoption year; below the separator is the average cohort effect, the mean of the cohort effects. Dots are estimates and whiskers are approximate 95 per cent magnitude intervals, the estimate plus or minus 1.96 placebo standard errors, where a placebo standard error is the standard deviation of the statistic across its placebos. The cohorts' whiskers are drawn lighter than the average cohort effect's. The figure prints no p-values: the cohorts' rank placebo p-values, over runs that use every eligible donor once as the placebo, are in Table D.1, and the average cohort effect's placebo p-value, over joint placebo assignments, is in Table D.2. The outcomes lie outside the headline family (the average cohort effect on GDP per capita and on exports per capita), whose placebo p-values carry a Holm adjustment. Log output per worker is $\ln(Y/L)$, where Y is real GDP at constant national prices (`rgdpna`) and L is same-vintage PWT persons engaged (`emp`). Log human capital is $\ln H$, where H is the PWT human-capital index (`hc`). Log TFP (labour share) is $\ln[(Y/L)(K/L)^{-\alpha_i}H^{-(1-\alpha_i)}]$, where K is the PWT 10.x `rnna` stock at constant 2017 national prices, Y is real GDP at constant national prices (`rgdpna`), L is same-vintage PWT persons engaged (`emp`), and H is the PWT human-capital index (`hc`). The economy-constant capital share α_i is one minus the mean of the first five admissible labour shares on or after 1950 (strictly between zero and one).

D.2 Accounting for growth: the output response

Figure D.2 and tables D.1 and D.2 show the cohort effects on the growth-accounting outcomes and the placebo inference for the pooled and cohort effects.

Table D.1. Cohort growth-accounting effects with placebo inference

Cohort (adoption year)	Estimate	Placebo SE	Approximate 95% interval	Rank placebo p-value	Rank and pool
Output per worker					
Japan (1955)	0.373	(0.155)	[0.070, 0.676]	0.042	2 of 48
Taiwan (1958)	0.294	(0.153)	[-0.006, 0.593]	0.115	6 of 52
South Korea (1962)	0.231	(0.149)	[-0.062, 0.524]	0.133	8 of 60
Singapore (1967)	0.373	(0.238)	[-0.094, 0.840]	0.123	10 of 81
Human capital					
Japan (1955)	-0.009	(0.032)	[-0.071, 0.053]	0.680	34 of 50
Taiwan (1958)	-0.023	(0.028)	[-0.078, 0.032]	0.333	18 of 54
South Korea (1962)	0.123	(0.038)	[0.048, 0.198]	0.048	3 of 62
Singapore (1967)	-0.040	(0.039)	[-0.117, 0.037]	0.234	22 of 94
TFP					
Japan (1955)	0.202	(0.115)	[-0.024, 0.427]	0.109	5 of 46
Taiwan (1958)	0.178	(0.130)	[-0.077, 0.433]	0.160	8 of 50
South Korea (1962)	0.154	(0.110)	[-0.062, 0.370]	0.155	9 of 58
Singapore (1967)	0.255	(0.186)	[-0.109, 0.620]	0.150	12 of 80
TFP (labour share)					
Japan (1955)	0.179	(0.104)	[-0.024, 0.383]	0.093	4 of 43
Taiwan (1958)	0.237	(0.127)	[-0.012, 0.487]	0.064	3 of 47
South Korea (1962)	0.215	(0.119)	[-0.018, 0.448]	0.094	5 of 53
Singapore (1967)	0.050	(0.182)	[-0.307, 0.407]	0.735	50 of 68

Notes: Each estimate is the cohort's average post-adoption synthetic difference-in-differences (SDID) effect over event times 0 to 20 on one growth-accounting outcome, in log points, with its placebo standard error (SE) in parentheses; the approximate 95 per cent interval is the estimate plus or minus 1.96 SE. Claims about how unusual an effect is rest on the rank placebo p-value, not on the interval, and no cell is marked where the two disagree.

Sample: Each row is one treated economy for one outcome. Donor pools are cohort- and outcome-specific. Cohort windows begin in 1950, 1952, 1955 and 1960 for Japan, Taiwan, South Korea and Singapore, and end 20 years after adoption.

Variables: The outcomes are defined in Table D.2.

Method: The outcomes are fitted with no controls. The rank placebo p-value compares the cohort's effect with placebo effects from runs that use every eligible donor once, as in Table C.3.

Table D.2. Pooled growth-accounting effects with placebo inference

Outcome and statistic	Estimate	Placebo SE	Approximate 95% interval	Placebo p-value	Simulation range	Correlation bound
Output per worker						
Average cohort effect	0.318	(0.093)	[0.135, 0.501]	0.002	[0.000, 0.006]	0.177
h = 10	0.301	(0.105)	[0.096, 0.506]	0.010	[0.004, 0.017]	0.200
h = 15	0.477	(0.139)	[0.205, 0.750]	0.002	[0.000, 0.006]	0.266
h = 20	0.655	(0.163)	[0.335, 0.974]	0.002	[0.000, 0.006]	0.312
Human capital						
Average cohort effect	0.013	(0.018)	[-0.023, 0.048]	0.459	[0.427, 0.489]	0.036
h = 10	0.008	(0.020)	[-0.031, 0.046]	0.648	[0.617, 0.678]	0.039
h = 15	0.013	(0.029)	[-0.045, 0.070]	0.630	[0.599, 0.660]	0.057
h = 20	0.036	(0.038)	[-0.038, 0.109]	0.329	[0.299, 0.358]	0.073
TFP						
Average cohort effect	0.197	(0.071)	[0.058, 0.336]	0.005	[0.001, 0.010]	0.136
h = 10	0.193	(0.081)	[0.035, 0.351]	0.018	[0.010, 0.027]	0.157
h = 15	0.303	(0.103)	[0.102, 0.505]	0.006	[0.002, 0.012]	0.199
h = 20	0.361	(0.121)	[0.123, 0.599]	0.005	[0.001, 0.010]	0.233
TFP (labour share)						
Average cohort effect	0.170	(0.070)	[0.034, 0.307]	0.022	[0.013, 0.032]	0.133
h = 10	0.160	(0.080)	[0.004, 0.316]	0.048	[0.035, 0.062]	0.153
h = 15	0.275	(0.098)	[0.083, 0.467]	0.011	[0.005, 0.018]	0.187
h = 20	0.322	(0.117)	[0.092, 0.552]	0.020	[0.011, 0.030]	0.221

Notes: The table reports pooled synthetic difference-in-differences (SDID) estimates for the four growth-accounting outcomes, in log points, with placebo standard errors (SEs) in parentheses; the approximate 95 per cent interval is the estimate plus or minus 1.96 SE. Claims about how unusual an effect is rest on the placebo p-value, not on the interval. These outcomes lie outside the headline family, so their p-values carry no Holm adjustment.

Variables: Output per worker is $\ln(Y/L)$, human capital is $\ln H$, and total factor productivity (TFP) is $\ln(Y/L) - \frac{1}{3}\ln(K/L) - \frac{2}{3}\ln H$, with Y real GDP at constant national prices, L same-vintage Penn World Table (PWT) persons engaged (`emp`), K the PWT 10.0 `rnna` stock at constant 2017 national prices and H the human capital index, all from the Penn World Table. This stock-based TFP is distinct from PWT's services-based TFP. Labour-share TFP is $\ln(Y/L) - \alpha_i \ln(K/L) - (1-\alpha_i)\ln H$, where the economy-constant capital share α_i is one minus the mean of the first five admissible PWT labour-share observations (`labsh`) from 1950. The treated economies' labour shares are flat for long initial stretches, so labour-share TFP varies the level, not the path, of the capital share. The average cohort effect, the horizon profile, the simulation range and the correlation bound are defined as in Table C.1.

Method: The outcomes are fitted with no controls. The four effects are estimated separately, each with its own unit weights, time weights, window and donor pool, so they are not an additive decomposition of the output-per-worker effect. Placebo inference follows Table C.1.

E. Robustness Appendix

E.1 Leave-one-out

Table E.1 reports the leave-one-out average, the leave-one-out delta, and the delta's placebo p-value for each outcome and each left-out cohort.

E.2 Augmented synthetic control

Table E.2 reports the synthetic difference-in-differences (SDID) and augmented synthetic control (ASCM) cohort effects, average cohort effects, and each method's pre-adoption fit statistic. Figures E.1 and E.2 plot the SDID and ASCM effect paths for GDP per capita and exports per capita, respectively.

E.3 Covariate adjustment

Table E.3 reports the average cohort effect re-estimated under the linear, quadratic, and year-dummy forms of the covariate adjustment. Table E.4 reports the pre-adoption fit of each form.

Table E.1. Leave-one-cohort-out pooled synthetic difference-in-differences (SDID) estimates by outcome

Dropped cohort	Outcome	LOO ATT	Headline-minus-LOO delta	Placebo p
Japan	GDP	0.457	-0.008	0.829
Taiwan	GDP	0.497	-0.049	0.215
South Korea	GDP	0.447	0.001	0.971
Singapore	GDP	0.393	0.056	0.263
Japan	Exports	1.481	-0.125	0.216
Taiwan	Exports	1.383	-0.027	0.803
South Korea	Exports	0.998	0.358	0.001
Singapore	Exports	1.562	-0.206	0.067
Japan	Employment per capita	0.142	-0.010	0.492
Taiwan	Employment per capita	0.165	-0.033	0.032
South Korea	Employment per capita	0.116	0.016	0.401
Singapore	Employment per capita	0.105	0.027	0.183
Japan	Capital	0.498	-0.009	0.824
Taiwan	Capital	0.486	0.003	0.945
South Korea	Capital	0.525	-0.036	0.393
Singapore	Capital	0.446	0.043	0.382
Japan	Capital per worker	0.389	0.009	0.848
Taiwan	Capital per worker	0.402	-0.004	0.949
South Korea	Capital per worker	0.430	-0.032	0.536
Singapore	Capital per worker	0.372	0.026	0.600
Japan	Investment	0.468	0.001	0.987
Taiwan	Investment	0.446	0.022	0.722
South Korea	Investment	0.444	0.024	0.714
Singapore	Investment	0.516	-0.047	0.555

Notes: Each row drops one cohort and recomputes the average cohort effect from the other three. Estimates and deltas are in log points.

Sample: Rows cover each of the four treated cohorts for six outcomes: GDP, Exports, Employment per capita, Capital, Capital per worker, and Investment.

Variables: LOO ATT is the leave-one-out average treatment effect. The delta is the average cohort effect minus the LOO ATT; a positive value means the dropped cohort raised the average cohort effect.

Method: Estimates are for the headline specification, which projects three covariates (logged pre-adoption population, the logged human-capital index and standardised inflation) out of the outcome before the weights are fitted. The Placebo p column is the two-sided placebo p-value of the delta, $(1 + k)/(1 + R)$, where k of the $R = 1{,}000$ seeded joint placebo assignments per outcome give an absolute delta at least as large as the observed absolute delta; the smallest attainable value is 1/1001, printed 0.001.

Table E.2. SDID and ASCM cohort effects and pre-adoption fit

	Effect		SDID	ASCM
	SDID	ASCM	Pre-adoption gap SD	Pre-period RMSPE
GDP per capita				
Japan (1955)	0.424	0.277	0.028	0.034
Taiwan (1958)	0.303	0.327	0.007	0.009
South Korea (1962)	0.453	0.531	0.015	0.081
Singapore (1967)	0.615	0.530	0.021	0.012
Average cohort effect	0.449	0.416	0.018	0.034
Exports per capita				
Japan (1955)	0.982	0.761	0.151	0.329
Taiwan (1958)	1.276	1.289	0.090	0.125
South Korea (1962)	2.429	2.244	0.167	0.208
Singapore (1967)	0.737	0.645	0.258	0.542
Average cohort effect	1.356	1.235	0.167	0.301

Notes: The table compares synthetic difference-in-differences (SDID) (Arkhangelsky et al., 2021) with augmented synthetic control (ASCM) (Ben-Michael, Feller, and Rothstein, 2021) on the headline outcomes, in log points, by cohort and for the average cohort effect, each beside the estimator's own pre-period fit statistic. The comparison is a robustness check across estimation and control procedures, not an estimator-only change.

Sample: The cohorts, with adoption years, are Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). Both estimators use the cohort windows and the regional donor exclusion of the sample calendar. The realised donor pools of SDID and ASCM are identical in every headline cell.

Variables: A cohort effect averages event times 0 to 20, with 0 the adoption year; the average cohort effect is the mean of the four cohort effects, and its fit values are the means of the four cohorts' fit values. The fit statistics are not computed the same way: SDID's pre-adoption gap SD and ASCM's pre-period RMSPE.

Method: Both estimators use the same economic covariates, logged population, the logged human capital index and standardized inflation, as pre-treatment window means, with the same covariate support screen. The control procedures differ (see Section 6.1). SDID effects have placebo inference, with placebo standard errors and p-values in the SDID appendix tables; ASCM cohort effects have pointwise conformal intervals and in-space placebo diagnostics, which come with the ASCM run and are not shown here.

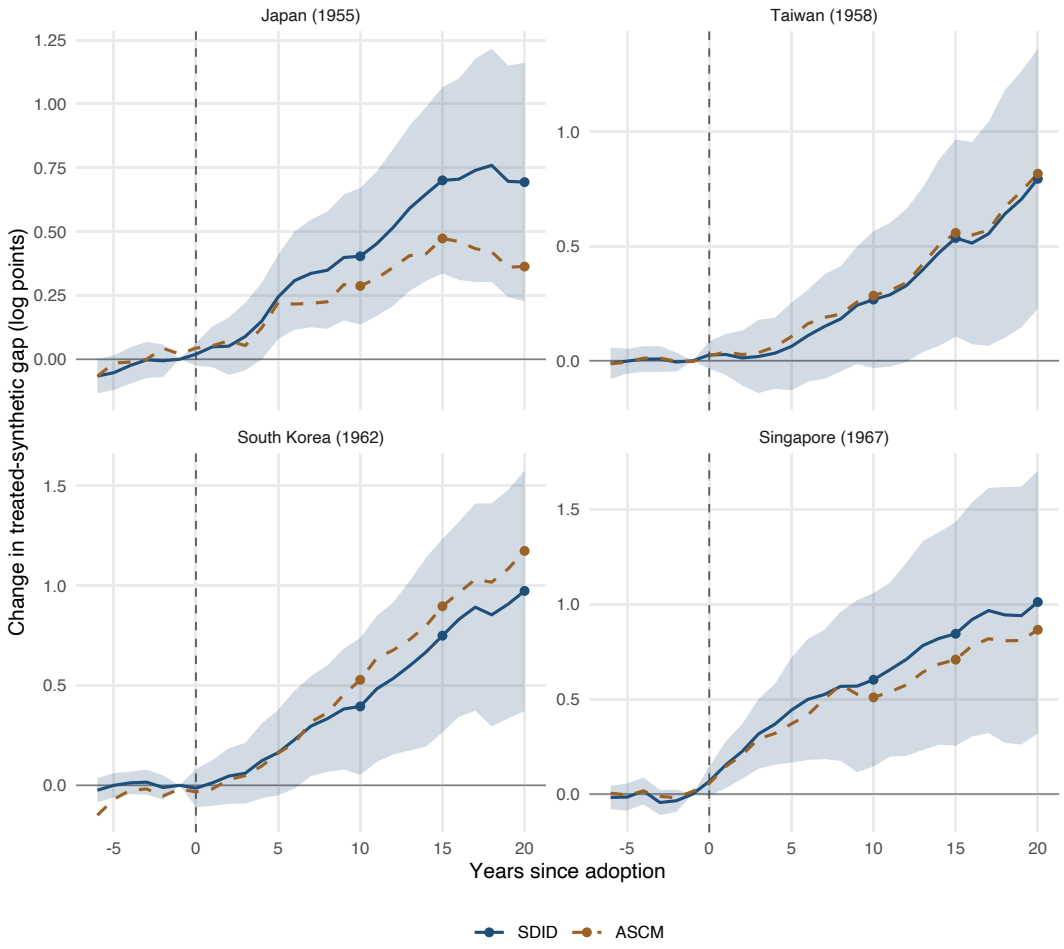


Figure E.1. SDID and ASCM effect paths: GDP per capita

Notes: Each panel compares the yearly effects of synthetic difference-in-differences (SDID) (Arkhangelsky et al., 2021) and augmented synthetic control (ASCM) (Ben-Michael, Feller, and Rothstein, 2021) on GDP per capita, in log points, for one cohort: Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). Event time 0 is the adoption year; markers are at event times 10, 15 and 20. The shaded band is the SDID approximate 95 per cent pointwise band (the estimate plus or minus 1.96 placebo standard errors); ASCM carries no band. Both estimators use the sample calendar's cohort windows and regional donor exclusion. The realised donor pools of SDID and ASCM are identical in every headline cell. The comparison is a robustness check across procedures, not an estimator-only change. Pre-period fit is read by each method's own statistic: for SDID the pre-adoption gap SD and for ASCM the pre-period RMSPE; the two are not computed the same way. ASCM cohort effects have pointwise conformal intervals and in-space placebo diagnostics, which come with the ASCM run and are not shown here. Solver-path precision: ASCM point estimates are pinned down only to the numerical precision of the optimiser's solution path; such differences are numerical sensitivity and not sampling error.

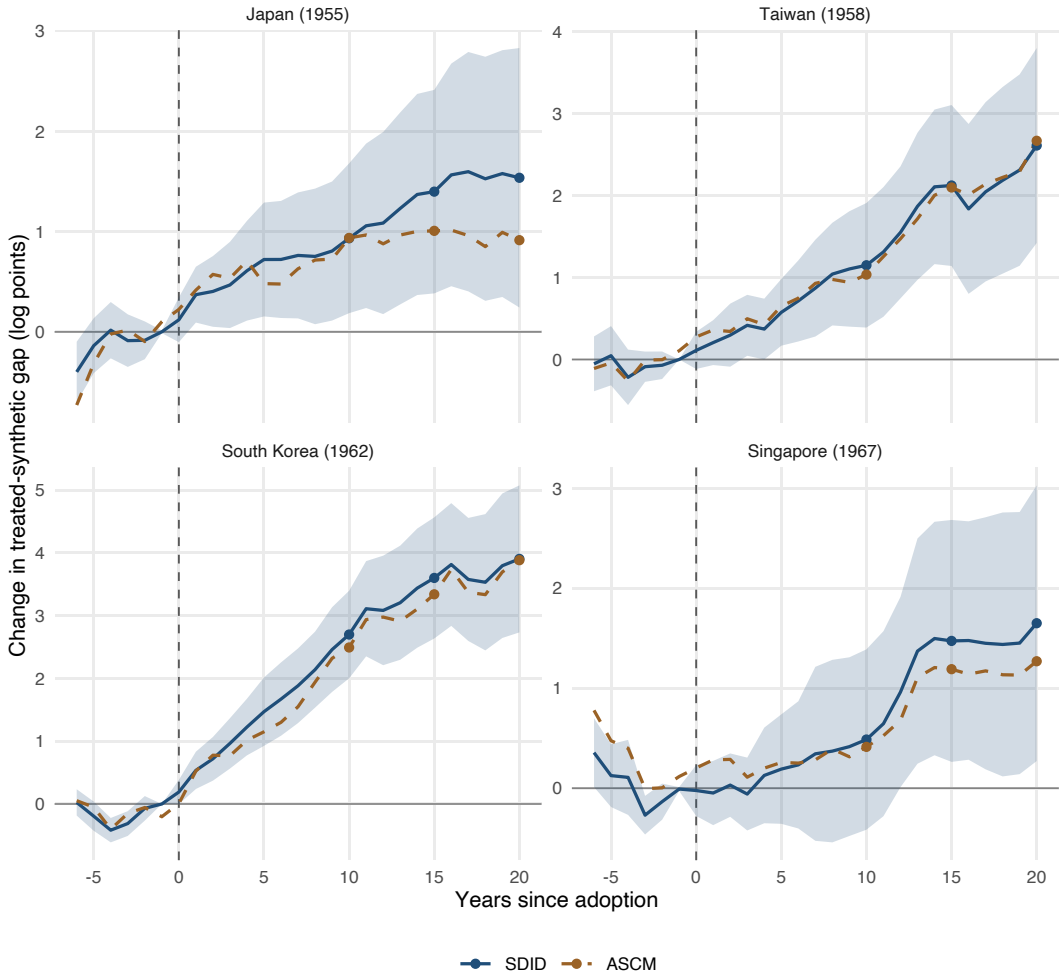


Figure E.2. SDID and ASCM effect paths: exports per capita

Notes: Each panel compares the yearly effects of synthetic difference-in-differences (SDID) (Arkhangelsky et al., 2021) and augmented synthetic control (ASCM) (Ben-Michael, Feller, and Rothstein, 2021) on exports per capita, in log points, for one cohort: Japan (1955), Taiwan (1958), South Korea (1962), and Singapore (1967). Event time 0 is the adoption year; markers are at event times 10, 15 and 20. The shaded band is the SDID approximate 95 per cent pointwise band (the estimate plus or minus 1.96 placebo standard errors); ASCM carries no band. Both estimators use the sample calendar's cohort windows and regional donor exclusion. The realised donor pools of SDID and ASCM are identical in every headline cell. The comparison is a robustness check across procedures, not an estimator-only change. Pre-period fit is read by each method's own statistic: for SDID the pre-adoption gap SD and for ASCM the pre-period RMSPE; the two are not computed the same way. ASCM cohort effects have pointwise conformal intervals and in-space placebo diagnostics, which come with the ASCM run and are not shown here. Solver-path precision: ASCM point estimates are pinned down only to the numerical precision of the optimiser's solution path; such differences are numerical sensitivity and not sampling error.

Table E.3. Average cohort effect by form of the covariate adjustment

Outcome	Linear	Quadratic	Year dummies
Panel A: Core outcomes			
GDP per capita	0.449	0.449	0.446
Exports per capita	1.356	1.347	1.294
Panel B: Additional outcomes			
Employment per capita	0.132	0.130	0.129
Capital	0.489	0.498	0.490
Capital per worker	0.398	0.405	0.376
Investment	0.468	0.455	0.459

Notes: The table compares the average cohort effect of the headline synthetic difference-in-differences (SDID) specification across three forms of the covariate adjustment. For GDP and exports per capita, the Linear column is the average cohort effect in Table 2. Estimates are in log points.

Sample: Each cell is the mean of the four cohort effects for Japan, Taiwan, South Korea and Singapore, each averaged over event times 0 to 20, with adoption at event time 0.

Variables: The covariates are the three headline covariates—logged pre-adoption population, the logged human-capital index and standardised inflation—collected in X_i. Linear interacts them with a linear time trend, $X_i \times t$, which is the headline. Quadratic adds $X_i \times t^2$. Year dummies interact X_i with a dummy for each year of the cohort's window except the first, the reference year.

Method: Each form is estimated on the donor rows with country and year fixed effects, separately for each cohort and window, and the adjusted outcome enters the same SDID estimator.

Table E.4. Pre-adoption fit by form of the covariate adjustment

Outcome	Linear	Quadratic	Year dummies
Panel A: Core outcomes			
GDP per capita	0.021	0.022	0.030
Exports per capita	0.193	0.217	0.259
Panel B: Additional outcomes			
Employment per capita	0.005	0.004	0.005
Capital	0.009	0.012	0.012
Capital per worker	0.020	0.021	0.022
Investment	0.066	0.068	0.075

Notes: The table reports how closely each form of the covariate adjustment fits the pre-adoption years. Lower values mean a closer fit. Values are in log points.

Sample: Each cell averages over the four cohorts: Japan, Taiwan, South Korea and Singapore.

Variables: The columns are the covariate forms of Table E.3. Pre-adoption fit is the root-mean-square of a cohort's pre-adoption effect, measured from the fit's own time-weighted baseline, then averaged over the four cohorts. It differs from the pre-fit SD ratio in Table B.1, which divides the standard deviation of the pre-adoption gap by that of the adjusted treated outcome.