Who Buys Growth Accelerations? Purchaser Routes and Their Persistence

Li Gan*

September 2026

Abstract

Which purchasers support rapid growth, and does foreign-led growth give way to domestic demand? I allocate domestic value added to households, government, investment, and foreign buyers using input–output accounts for 80 economies in 1995–2022. Previous-year prices separate real contributions from relative-price changes across purchasers. Foreign demand leads 19 of 36 independently dated growth accelerations and is more prevalent in smaller economies. Among the 17 foreign-led onsets with ten-year coverage, foreign demand still leads in years eight to ten in 11 and households lead in 4; no cumulative ten-year record flips, because the seven-year selection window builds in a large lead. Domestic leadership emerges in a later, separate window without a cumulative takeover. Applying a common GDP deflator changes 83 of 400 five-year engine assignments and doubles government leadership, but leaves the predominance of household and foreign demand in rapid-growth episodes intact. Historical accounts since 1965 provide supporting evidence with explicit source-quality screens. Purchaser routes describe the composition and persistence of realized growth; they do not rank policies for generating it.

Keywords: growth accelerations; domestic demand; exports; input–output accounts; economic development
JEL codes: O11, O47, F43

*Department of Economics, Texas A&M University, College Station, TX 77843, USA. Email: ganli@tamu.edu.

1 Introduction

Exports are said to start rapid growth and domestic demand to sustain it. That sequence is the development question behind the familiar distinction between growth supported by foreign markets and growth supported by domestic demand. Evaluating it requires identifying the purchasers supporting domestic production during the same dated growth episodes. Neither the largest expenditure share nor the conventional net-export contribution answers that question: the first describes a level, while the second subtracts imports induced by every domestic purchaser as well as those embodied in exports.

This paper measures purchaser routes during growth accelerations and follows them beyond the initial acceleration window. The main evidence combines previous-year-price input–output accounts for 80 economies with independently dated Penn World Table growth episodes. Domestic value added is allocated to household consumption, government and nonprofit consumption, investment, and deliveries to foreign buyers at the national border. Imported content is removed and domestic upstream suppliers are included. The *purchaser state* records each route's share of current domestic value added; the *growth engine* is the largest real contribution over an episode. Germany and Korea illustrate the distinction: both have household-centered recent states, although foreign demand led their growth in every complete five-year episode of 2001–2015.

The central result concerns persistence. Foreign demand leads 19 of 36 seven-year accelerations selected from real GDP per capita, households lead 15, and investment leads 2. Among the 33 onsets with ten subsequent years of purchaser data, 17 are foreign-led in the baseline seven-year window. Selection over that window builds in a lead: at year seven the foreign margin over the best domestic route is 11.5 percentage points in the median episode, so three further years can rarely reverse the cumulative record, and none does. The informative statistic is the engine of years eight to ten taken alone. Foreign demand still leads in 11 of the 17, households in 4, investment in 1, and government in 1. A separate comparison of successive seven-year windows agrees: of 12 initially foreign-led cases with sufficient coverage, 7 remain foreign-led, 4 become household-led and 1 government-led in the subsequent window. A later domestic engine therefore need not replace foreign demand in the cumulative growth record, and a foreign majority persists beyond the selection window. Continuous contribution paths show that annual household leadership can also reflect a fall in the foreign contribution, including during 2009, rather than an increase in the household contribution. These findings do not support a general ten-year export-to-domestic sequence.

Country size helps describe the heterogeneity. Foreign demand leads ten of 12 accelerations in the smallest population tercile and two of 12 in the largest. Conditioning on

initial foreign-demand support halves the population coefficient; the analogous controlled relationship in the full episode panel is not statistically distinguishable from zero. Size and initial exposure therefore describe selected acceleration routes, without establishing a universal growth mechanism. The cross-country pattern does not imply that foreign demand is an available or superior strategy for every small economy.

Measurement affects individual growth histories in a systematic way. Holding nominal purchaser accounts, aggregate real growth, and episode dates fixed, replacing route-specific previous-year prices with a common GDP deflator changes 83 of 400 five-year engine assignments. Government leadership doubles from 25 to 50 episodes under common deflation. Yet the broad conclusion about rapid growth survives: households and foreign demand remain the two most frequent engines. Among 100 period-specific top-quartile growth episodes, their counts are 53 and 38 with route-specific prices and 50 and 37 with a common deflator. The contribution is therefore a correction to particular histories and to the incidence of government leadership, rather than a wholesale reversal of the household–foreign-demand pattern.

The two definitions of rapid growth also warrant separate inference. In the top-quartile sample households are the modal engine; in the acceleration sample foreign demand is modal. A joint country bootstrap gives a 95-percent interval of $[-0.15, 30.62]$ percentage points for the increase in foreign leadership across samples. The change in the foreign-minus-household margin has an interval of $[-4.78, 58.54]$ points. Neither establishes a population ranking reversal. Matching on initial purchaser composition and income further removes the raw household-frequency difference between high-growth and other episodes. These limits are part of the empirical finding.

The paper connects to the literature on growth accelerations and growth spells (Hausmann et al., 2005; Jones and Olken, 2008; Berg et al., 2012) by measuring which purchasers support growth within and after the dated episodes. It also builds on import-adjusted demand accounting (Kranendonk and Verbruggen, 2008; Bussière et al., 2013; Baccaro and Hadziabdic, 2024). The four-way identity and the distinction between export-supported value added and net exports are established results. Relative to that work, this paper isolates the effect of price treatment within common accounts and connects real purchaser contributions to fixed-date acceleration paths. The common-deflator exercise is a controlled comparison within this dataset, not a replication of a prior paper's different countries, periods, or purchaser categories.

Trade-in-value-added research provides the domestic-content interpretation (Johnson and Noguera, 2012; Koopman et al., 2014; Bems et al., 2010). Structural transformation research emphasizes sectoral production and expenditure (Kongsamut et al., 2001; Ngai and

Pissarides, 2007; Herrendorf et al., 2014; Comin et al., 2021); purchaser routes provide a complementary description of the demand supporting that production. A historical extension for 25 economies since 1965 places the recent episodes in a longer perspective. Its route-price anomalies are disclosed and screened, and it does not determine the main acceleration result. The analysis remains an account of realized growth composition. It cannot identify the effect of inducing one purchaser's spending or the income-transmission parameters behind a domestic handoff.

2 Purchaser states, growth engines, and routes

2.1 Domestic value added supported by a purchaser

For country c in year t, let Z_{ct} be the domestic intermediate-input matrix, x_{ct} the vector of gross output, and

$$A^d_{ct} = Z_{ct}\,\mathrm{diag}(x_{ct})^{-1}, \qquad L^d_{ct} = (I - A^d_{ct})^{-1}. \tag{1}$$

Let ν_{ct} be the vector of value added per unit of gross output. For purchaser $k \in \{C, G, I, X\}$, let f^k_{ct} denote domestic deliveries to that purchaser route. Household consumption is C, government and nonprofit consumption is G, gross fixed capital formation plus inventory change is I, and X contains deliveries to all foreign intermediate and final users. Domestic value added supported by purchaser k is

$$V^k_{ct} = \nu'_{ct} L^d_{ct} f^k_{ct}. \tag{2}$$

This is a backward-linkage allocation of domestic value added. An export of a manufactured good supports domestic services and primary inputs upstream; an imported input is excluded because it does not generate domestic value added.

The purchaser state is

$$\alpha^k_{ct} = \frac{V^k_{ct}}{\sum_{j\in\{C,G,I,X\}} V^j_{ct}}, \qquad \sum_k \alpha^k_{ct} = 1. \tag{3}$$

The continuous vector $\boldsymbol{\alpha}_{ct}$ is the primary state variable. No clustering is required. Terms such as "consumption-centered" and "export-oriented" are shorthand for relative positions in this simplex, not structural country types.

2.2 Why this equals import-adjusted demand

Suppose A^m contains imported intermediate coefficients. Gross final demand for purchaser k less the imported intermediates induced along domestic production chains is

$$\mathbf{1}'f^k - \mathbf{1}'A^m(I - A^d)^{-1}f^k. \tag{4}$$

Using the input–output identity that one unit of gross output is exhausted by domestic intermediates, imported intermediates, and value added, this expression equals $\nu'(I-A^d)^{-1}f^k$, up to valuation, tax, and statistical-discrepancy items. It is therefore the same aggregate object as the import-adjusted demand components in Kranendonk and Verbruggen (2008), Bussière et al. (2013), and Baccaro and Hadziabdic (2024). Equation (2) makes the domestic-value-added interpretation explicit and permits sectoral and upstream decompositions, but it does not create a distinct aggregate four-way identity.

2.3 Real growth contributions

Current-price purchaser shares cannot identify a real growth engine because their movements combine quantities and prices. Let $\widetilde{V}^k_{ct}$ be year-t purchaser-supported value added valued at year-$t-1$ prices, from the previous-year-price input–output account. Define the annual contribution of purchaser k to real value-added growth as

$$g^k_{ct} = 100 \times \frac{\widetilde{V}^k_{ct} - V^k_{c,t-1}}{\sum_j V^j_{c,t-1}}. \tag{5}$$

The contributions add to the account's real growth rate exactly, to within 10^{-11}. Reaching that requires care with the source tables: in the historical previous-year-price tables an industry's deliveries summed across destinations need not equal its column gross output, and the domestic input matrix is normalized on gross output, so the accounts are computed on row-balanced tables whose supported value added reproduces the source's own previous-year-price value added to within 10^{-15}. The engine labels are barely sensitive to that choice: scaling the four solved components proportionally onto the same total instead assigns the same purchaser in 267 of 275 episodes, all eight disagreements falling in the historical half where the two conventions can differ, and leaves the leading-engine counts unchanged. Contributions can be negative. In a contraction, the "leading" purchaser is the one with the largest contribution, which may be the least negative contribution.

The foreign-demand contribution g^X_{ct} should not be confused with the conventional net-export contribution. For the acceleration sample I also construct, from the same current-

and previous-year-price ICIO tables,

$$g_{ct}^{NX} = 100 \times \frac{(\widetilde{X}_{ct}^{gross} - X_{c,t-1}^{gross}) - (\widetilde{M}_{ct}^{gross} - M_{c,t-1}^{gross})}{\sum_j V_{c,t-1}^{j}}. \tag{6}$$

The first object asks how much real domestic value added is supported by foreign intermediate and final buyers after removing imported content and tracing domestic suppliers. The second subtracts the real growth contribution of all gross imports from that of gross exports. Imports used by expanding household demand or investment can therefore make g_{ct}^{NX} nonpositive even when g_{ct}^{X} is positive. Conversely, gross exports can rise with little domestic value added when imported content is high.

The conventional object is not a straw man. National-accounts contribution tables define the external component as net exports of goods and services (U.S. Bureau of Economic Analysis, 2026; National Bureau of Statistics of China, 2025), the "traditional method" of Kranendonk and Verbruggen (2008), and the Chinese-language debate over correcting it—input–output allocation of imported intermediates (Shen, 2009) versus direct deduction of imports from exports, or crediting exports with the consumption and investment they induce (Lin and Li, 2003)—turns on exactly the choice of external-demand object. Equation (2) is the input–output allocation, carried out consistently for four purchasers and many countries.

I aggregate annual contributions over nonoverlapping five-year episodes e:

$$G_{ce}^{k} = \sum_{t \in e} g_{ct}^{k}, \qquad E_{ce} = \arg\max_{k} G_{ce}^{k}. \tag{7}$$

The first episode is 1966–1970. The last, 2021–2022, contains only two years and is reported separately. Annual contributions are measured in percentage points; their sum across years records cumulative contributions, not a compound growth rate. A country's purchaser history is the joint sequence $\{\boldsymbol{\alpha}_{ct}, E_{ce}\}$: a slowly moving state together with a more variable marginal engine. A purchaser route is one of the four channels k through which final demand supports domestic value added, so a history records which route holds the state and which route leads each episode.

2.4 Domestic conversion at a fixed network

The accounting framework supplies a fixed-coefficient measure of how one unit of purchaser expenditure reaches domestic production. Let D_{ct}^{k} be total expenditure by purchaser k, including imported final products for the three domestic purchasers. Define the network-

conversion index

$$M_{ct}^{k} = \frac{\nu_{ct}'(I - A_{ct}^{d})^{-1} f_{ct}^{k}}{D_{ct}^{k}}. \tag{8}$$

It is domestic value added supported per unit of purchaser expenditure. It includes indirect domestic suppliers and excludes imported final products and intermediates. To make the marginal interpretation explicit, write $q^k = f^k/D^k$ for the purchaser's final-demand bundle. Holding the base-year value-added coefficients, production network, and bundle fixed,

$$\left.\frac{\partial V^k}{\partial D^k}\right|_{\nu, A^d, q^k} = \nu'(I - A^d)^{-1} q^k = M^k.$$

Thus M^k is a predetermined marginal domestic-conversion index: it answers how much domestic value added the existing network mechanically supports from one additional unit of purchaser-k expenditure. It is not a causal fiscal or trade multiplier.

The direct value added associated with the industries selling to purchaser k is $\nu_{ct}' f_{ct}^k$. The upstream-propagation share is

$$U_{ct}^{k} = 1 - \frac{\nu_{ct}' f_{ct}^{k}}{\nu_{ct}'(I - A_{ct}^{d})^{-1} f_{ct}^{k}}. \tag{9}$$

It records how much purchaser-supported value added arises beyond the final-demand industries. A larger U^k means broader upstream propagation, not necessarily higher productivity or welfare.

2.5 Sources and coverage

LR-WIOD provides annual current-price and previous-year-price world input–output tables for 25 economies and 23 industries from 1965 through 2000 (Woltjer et al., 2021). It supplies the supporting historical extension; the main acceleration evidence uses OECD alone. The sample contains Australia, Austria, Belgium, Brazil, Canada, China, Denmark, Finland, France, Germany, Greece, Hong Kong, India, Ireland, Italy, Japan, Korea, Mexico, the Netherlands, Portugal, Spain, Sweden, Taiwan, the United Kingdom, and the United States.

The OECD inter-country input–output system supplies current-price tables from 1995 through 2022 and previous-year-price tables for real annual changes (OECD, 2025, 2026); the 2023 and 2024 tables in the previous-year-price release are not used, so the panel ends in 2022. The system provides the recent segment for the 25 long-panel economies and a balanced 80-economy panel of purchaser states and real growth contributions from 1995 and 1996, respectively. The 80 economies are those the OECD tables report separately: 38 OECD members and 42 other economies, Hong Kong and Taiwan among them; the rest-of-world

aggregate is not a panel member, and the 25 long-panel economies are a subset. Appendix Table 25 lists them with their 1995 income terciles. OECD industries are first retained at their published detail and then aggregated only after purchaser-supported value added is calculated.

The four categories are defined the same way in both sources, with one documented difference. Household consumption is household final consumption; government includes general-government and nonprofit-institution consumption; investment includes gross fixed capital formation and inventory change; exports include all intermediate and final deliveries to foreign economies. In the OECD tables residents' direct purchases abroad sit in household consumption and non-residents' purchases in the foreign route, whereas LR-WIOD reports both only as unallocated totals outside the industry rows (Woltjer et al., 2021), so before 1995 the foreign route contains cross-border deliveries only and household consumption contains neither component. Using foreign deliveries rather than customs exports is essential: it assigns only domestic value added to X while preserving the full domestic production chain.

2.6 Aligning the current-price state

LR-WIOD and OECD differ in statistical standards, industry detail, valuation, and national source revisions. A raw splice would create mechanical jumps in 1995. I align the composition in additive-log-ratio coordinates. With exports as the reference category, define

$$z_{ct}^{k} = \log\left(\frac{\alpha_{ct}^{k}}{\alpha_{ct}^{X}}\right), \quad k \in \{C, G, I\}. \tag{10}$$

For each country, the correction is the mean OECD-minus-LR-WIOD difference in z^k over 1995–2000. I add this fixed correction to pre-1995 LR-WIOD values and map the four components back to the simplex. OECD values are used from 1995 onward. Before correction, the mean absolute purchaser-share gap between the two sources in the overlap is 2.66 percentage points.

This alignment preserves pre-1995 log-ratio movements in LR-WIOD, not exact percentage-point movements of each share, and prevents the source switch from being read as an economic transition. It is not a claim that the two databases measure identical levels. Appendix Table 27 therefore reports the endpoint results using unadjusted LR-WIOD shares. Alternative overlap windows and level rather than compositional alignment remain additional checks.

No level correction enters equation (5). The long real-growth series uses LR-WIOD contributions through 1995 and OECD contributions beginning in 1996. Every annual change

is therefore calculated from a current-price base and previous-year-price value within the same statistical system.

2.7 Measurement limits

The early LR-WIOD tables are harmonized historical estimates, not 25 annual benchmark censuses. They combine available national benchmark tables with national accounts, trade data, and constrained interpolation and extrapolation. Aggregate agreement with national accounts does not validate the allocation across purchaser routes. Early turning dates and route-specific volume movements require additional scrutiny. This limitation is particularly important for China and India, where early source material is thinner.

The accounts also measure realized demand, not autonomous demand shocks. Exports, investment, government purchases, and consumption are jointly determined with income, prices, policy, and capacity. The labels "engine" and "route" are accounting descriptions. Causal language is reserved for designs that bring external timing and predetermined exposure to the accounts.

The foreign-demand category has a precise boundary. It assigns domestic value added to exports sold to foreign intermediate and final users at the national border; it does not identify whether the final purchaser abroad is a household, government, or investor. Tracing that final foreign purchaser would require a fully linked MRIO destination decomposition. Main-text accounting table headers use "Foreign demand"; in prose, "exports" is shorthand for foreign-demand-supported domestic value added, not gross exports, net exports, or a claim about the identity of the purchaser outside the exporting economy.

The OECD current-price purchaser account expresses support relative to source value added, with a maximum coverage residual of 0.2834 percentage point. I retain that coverage ratio separately and normalize the four purchaser components by their sum before constructing states, bridges, and initial-state controls. Real contributions retain the denominator in equation (5). This distinguishes a compositional normalization from a correction to source data. The historical quality screen and source-specific estimates are reported in Section 7.

3 What price measurement changes

Route-specific deflation is useful only if it changes how growth episodes are interpreted. I isolate its contribution using the same nominal accounts and the same aggregate real growth. Write $a_t^k = V_t^k / \sum_j V_t^j$ and let $\gamma_t = \sum_k g_t^k/100$ be aggregate real growth as a fraction. If all

purchaser-supported values receive a common deflator, their fractional contributions are

$$g_t^{k,\text{common}}/100 = a_t^k(1+\gamma_t) - a_{t-1}^k. \tag{11}$$

They add to γ_t exactly. Thus nominal composition, total real growth, and the independently selected episode lists are identical across the two calculations. Only the allocation of real growth across purchasers changes.

Table 1: How a common deflator changes purchaser-engine assignments

Sample	Prices	N	C	G	I	X	Changed
All episodes	PYP	400	165	25	22	188	83
All episodes	Common	400	152	50	23	175	83
High growth	PYP	100	53	2	7	38	18
High growth	Common	100	50	4	9	37	18
Accelerations	PYP	36	15	0	2	19	5
Accelerations	Common	36	14	0	2	20	5

Notes: PYP denotes the maintained route-specific previous-year-price account. Common denotes equation (11). C, G, I, X denote households, government, investment, and foreign demand. Changed counts compare the same episodes and are repeated for clarity across the two methods. Acceleration windows are the same 36 onsets and $t+1$ through $t+7$.

Table 1 identifies the scope of the correction. Engine assignments change in 20.75 percent of all episodes, 18 percent of high-growth episodes, and 13.89 percent of accelerations. The most systematic aggregate difference concerns government: a common deflator assigns it twice as many episodes. A purchaser whose relative supported-value-added price rises can acquire apparent real leadership when that relative revaluation is not removed. Appendix Tables 14 and 15 show which labels change and their distribution across global periods. Common deflation moves 32 non-government episodes into government leadership and 7 government episodes out, for a net gain of 25. Of that net gain, 13 occur in 2016–2020, when the count rises from 11 to 24. Thus the correction is particularly relevant when assessing government leadership around the pandemic, rather than establishing a new ranking of growth strategies.

The household–foreign-demand comparison is more stable. These two routes jointly lead 91 of 100 high-growth episodes under PYP and 87 under common deflation. Neither price treatment makes investment or government the typical largest contributor in that sample. Reclassification should therefore not be presented as proof that the broad account of rapid growth reverses. It matters for assigning particular country episodes and for assessing the

incidence of government-led growth. Nor does a lower count of government-led episodes estimate the effectiveness of fiscal policy: engine leadership is a relative contribution, often attained when other routes are weak.

4 Growth accelerations and subsequent routes

4.1 Who buys growth accelerations?

The top-quartile definition asks who leads during fast common-calendar windows. To connect the same accounting object to the growth-acceleration literature, I also apply an HPR-style filter to PWT real GDP per capita. An onset must be followed by seven-year log-linear growth of at least 3.5 percent, an acceleration of at least 2 percentage points relative to the prior seven years, and income at the end of the post window above its maximum over the entire prior PWT series, which rules out recoveries to a previous peak. Overlapping candidate onsets within a country are collapsed to the largest growth change. Applied to the whole of PWT the detector returns 208 onsets, of which 104 fall in the 80 panel economies. The 68 before 1996 precede the OECD previous-year-price accounts and are set aside; the remaining 36 onsets, all in 1997–2015, each have a complete seven-year window in the OECD panel, and for each I sum the seven annual purchaser contributions from $t+1$ through $t+7$. This exact alignment is important: assigning an onset to an encompassing calendar block can include pre-onset years and omit later acceleration years.

Table 2: Bridging definitions of rapid growth and purchaser engines

Definition	N	C	G	I	X
Top-quartile five-year growth	100	53.0%	2.0%	7.0%	38.0%
Onsets: calendar block	36	22.2%	0.0%	0.0%	77.8%
Onsets: years 1–5	36	36.1%	2.8%	5.6%	55.6%
Onsets: years 1–7	36	41.7%	0.0%	5.6%	52.8%

Notes: Codes C, G, I, X denote the four purchasers. The first row uses all 100 period-specific top PWT growth quartile episodes. The remaining rows hold fixed the 36 acceleration onsets. "Encompassing block" assigns each onset the engine of the five-year calendar block containing t; the last two rows sum annual contributions from $t+1$ through $t+5$ or $t+7$. The exercise separates sample selection from window alignment. It is an HPR-style filter, not a replication of the original country sample or episode list. "Foreign demand" is domestic value added supported by exports to foreign intermediate and final buyers.

Foreign demand leads 19 of the 36 exact acceleration windows; households lead 15 and investment two. No selected acceleration is government-led. As in the high-growth sample,

19 against 15 is not a margin this many episodes can resolve, and the cross-sample contrast also requires its own interval, reported in Section 6. Table 2 reconciles this with household leadership in 53 of the 100 top-quartile calendar episodes. Sample selection is the main break: the same 36 onsets are 77.8 percent foreign-led when mechanically assigned to their encompassing calendar block. Aligning the window to the onset lowers the foreign share to 55.6 percent over five years and 52.8 percent over seven years. Thus the two 53-percent facts answer different questions; the acceleration result is not a relabeling of the broad high-growth sample. The filter remains descriptive because foreign demand and domestic capacity are jointly selected.

The episode composition is concentrated in the initial globalization and transition period: 21 of 36 onsets occur in 1997–2002, including eleven post-communist economies. That concentration does not explain the modal result by itself. Foreign demand leads 52.4 percent of the 1997–2002 cohort and 53.3 percent of the 15 later onsets. The complete country–onset list, cohort table, and follow-on engines are reported in Appendix Tables 29 and 30.

Two properties of the sample follow from the filter rather than from the accounts. First, 17 of the 36 onsets follow seven years of negative average growth before clearing the prior-peak test. Foreign demand leads 11 of them and 8 of the 19 others, a difference a Fisher exact test does not resolve ($p = 0.20$); the foreign majority is thus more pronounced among rebounds, one more reason the paper does not rest on the within-sample ordering. Second, PWT coverage of the former Soviet and Yugoslav economies begins in 1990, so 1997 is the first onset year the seven-year filter can evaluate for them: 6 of the 7 onsets dated 1997 fall on that first admissible year, with the transition contraction as their entire pre-window. Because the purchaser panel already excludes onsets before 1996, the boundary could only date these onsets earlier or remove them, never later, and the engine assignment uses the fully observed post-onset window.

Country size sharply conditions the route. Table 3 ranks the 36 acceleration episodes two ways. By population, foreign demand leads ten of 12 accelerations in the smallest tercile, seven of 12 in the middle, and two of 12 in the largest—83.3, 58.3 and 16.7 percent. The gradient does not depend on where the tercile cuts fall: across the 36 episodes the rank correlation between log population and foreign leadership is -0.57, with a permutation p-value of <0.001 against 200,000 draws. Ranking instead by output-side real GDP at PPP in the onset year gives the same ordering more weakly—75.0, 50.0 and 33.3 percent, with a rank correlation of -0.40 ($p = 0.016$). Population has the stronger unconditional association in these data. The accounts do not determine which size measure represents the relevant market-access mechanism. Large economies more often mobilize household demand or investment; China is one of the two investment-led exact windows. For small economies

the external market is large relative to domestic demand; this does not make one engine intrinsically superior.

Part of the gradient travels with the inverse relation between country size and trade openness documented by Alesina and Wacziarg (1998), which appears here in the foreign-demand-supported share rather than in gross trade: smaller economies enter the sample with a larger onset-year foreign share (rank correlation between log population and the share -0.55). In a linear probability regression of foreign leadership on log population, the coefficient is -0.17 per log point of population (country-clustered standard error 0.03); conditioning on the onset-year foreign share halves its magnitude to -0.09 (0.04, $p = 0.02$), and adding log income per person leaves it at -0.10 ($p = 0.06$). The share itself enters at 2.32 (0.59). Whether that conditioning removes a confounder or the proposition's own market-access channel is a question the cross-section cannot settle, so the conditional coefficient is reported as a coefficient and not as a decomposition.

Table 3: Country size and the purchaser route during growth accelerations

Size measure	Tercile	Episodes	Household	Government	Investment	Foreign demand
PPP GDP	Small	12	25.0%	0.0%	0.0%	75.0%
PPP GDP	Middle	12	50.0%	0.0%	0.0%	50.0%
PPP GDP	Large	12	50.0%	0.0%	16.7%	33.3%
Population	Small	12	16.7%	0.0%	0.0%	83.3%
Population	Middle	12	41.7%	0.0%	0.0%	58.3%
Population	Large	12	66.7%	0.0%	16.7%	16.7%

Notes: Terciles are formed within the 36 acceleration episodes using PWT output-side real GDP at PPP or population in onset year t, 12 episodes per cell. Each cell reports the share for which that purchaser makes the largest cumulative real domestic-value-added contribution from $t+1$ through $t+7$. For foreign demand, the trend across terciles has $z = -3.27$ ($p = 0.001$) by population and $z = -2.04$ ($p = 0.041$) by PPP GDP. Because the tercile cut is a researcher choice, the text reports rank correlations on the continuous measures instead.

The size gradient is specific to acceleration selection, not a universal law across all growth windows. In the full 400-episode panel, a one-standard-deviation increase in log population lowers the probability of foreign-demand leadership by 11.1 percentage points in the unadjusted period-fixed-effect comparison (standard error 3.5 points). The relationship disappears after controlling for initial income per person and the initial foreign-demand-supported share. Log PPP GDP is not statistically informative in the full panel even before controls. The reading that matches Appendix Proposition 2 is deliberately limited: the selected-acceleration gradient survives the same control at roughly half its size, while the full-panel

gradient does not. The initial foreign-demand share is a suggestive proxy for market access and export capacity—it may be a mediator of the size channel, a confounder, or an outcome of the same latent forces, and this cross-section cannot distinguish those cases. The tempting stronger reading—that conditioning on the share absorbs the proposition's own channel, so the vanishing gradient confirms the mechanism—is exactly what such a cross-section cannot establish, and the controlled null is therefore reported as a null. The proposition also speaks to economic scale without fixing the proxy; that the unconditional gradient appears in population rather than PPP GDP suggests, as interpretation rather than as a test, that the relevant object is the size of the latent domestic market rather than income mass—population here substitutes for the proposition's own Y_0. Country size therefore describes which economies enter the foreign-led acceleration sample; it does not by itself identify the transmission of foreign demand. Appendix Table 31 reports the full coefficients.

4.2 After the onset: the domestic handoff

Does an initially foreign-led acceleration become domestically led at longer horizons? Table 4 keeps the baseline onset dates fixed and restricts all horizons to the same 33 episodes with ten-year coverage. Cumulatively, foreign demand leads 19 at five years, 17 at seven, and 18 at ten, the newcomer at ten being Laos 2006, and 17 of the 17 baseline foreign-led episodes, that is all of them, remain foreign-led at ten years; the same 17 are foreign-led at ten years under the common deflator, holding the original PYP-selected group fixed. The cumulative rows are, however, weak evidence of persistence. Selection on foreign leadership over years one to seven builds in a lead that three further years can rarely reverse: at year seven the foreign margin over the best domestic route is 11.5 percentage points in the median episode, from 0.1 for Estonia 1997 to 42.6 for Ireland 2012. The last row of Table 4 therefore classifies the same episodes by their contributions over years eight to ten alone, after the selection window. Among the 17 baseline foreign-led episodes, foreign demand still leads in 11, households in 4 (Belarus 1997, Kazakhstan 1997, Cambodia 1998, and Hong Kong 2001), investment in 1, and government in 1. The post-window majority is the persistence finding; the cumulative record is arithmetic. Both are statements about the observed ten-year record, not proof that a handoff cannot occur later or under a different policy.

Table 4: Following the same acceleration onsets across horizons

Window	N	C	G	I	X
$t+1$ to $t+5$	33	12	0	2	19
$t+1$ to $t+7$	33	14	0	2	17
$t+1$ to $t+10$	33	14	0	1	18
$t+8$ to $t+10$ only	33	14	2	3	14

Notes: The baseline seven-year detector is run once. All rows retain the same 33 country–onset pairs with complete ten-year coverage. The first three rows sum contributions from $t+1$; of the 17 episodes foreign-led at seven years, all remain foreign-led at ten. The last row classifies the same episodes by the sum over years $t+8$ through $t+10$ only; among the 17 baseline foreign-led episodes, 11 are foreign-led in that window. This is distinct from rerunning the detector with longer pre- and post-windows.

A subsequent nonoverlapping window asks a different question. A second seven-year purchaser window is observed for 27 earlier accelerations, of which 15 retain the same leading route. Among the 12 initially foreign-led cases, 7 remain foreign-led, 4 switch to households, and 1 to government in $t+8$ through $t+14$ (Table 5). Their median follow-on per-capita growth is 2.5 percent and 66.7 percent remain above 2 percent. This comparison agrees with the post-window row of Table 4 and complements the growth-slowdown literature (Eichengreen et al., 2012). It does not establish a domestic takeover of growth accumulated since the onset.

The ten-year row in Appendix Table 32 instead reruns the onset detector. Its household majority refers to 24 differently selected onsets. It cannot be used as evidence of a handoff within the baseline 36. Selection, cumulative horizon, and subsequent-window leadership are kept separate throughout.

Winner labels can, however, flip for two different reasons: because household contributions accelerate, or because foreign contributions collapse. Figure 1 separates them by tracing the continuous path. For the 19 foreign-led onsets, no average cumulative crossing occurs at any observed horizon: in the balanced ten-year subset of 17 episodes, the cumulative totals at year ten are 33.6 versus 12.6 percentage points. The annual gap turns positive only at year 8 and again in years 11–12. Each flicker is carried by the episodes that reach that horizon in 2009, when foreign margins collapsed, rather than by a rise in the mean household margin: at year eight these are the 6 onsets of 2001 (mean foreign margin -3.5 points, household 0.9), and without them the year-eight gap is -0.7 points; at year 12 they are the 4 onsets of 1997, and at year 11 1 of the 13 episodes at that horizon, a 1998 onset. The mean household margin never exceeds 1.7 points at any horizon, and the year-eight gap is indistinguishable from zero (bootstrap interval -0.9 to 2.7 points). The annual flickers should therefore be

read as decaying foreign impulses meeting a stable domestic margin, not as domestic demand taking over the cumulative record. The paths do not identify income transmission, capacity, or displacement. Selection on foreign leadership over the first seven years already imposes the cumulative ordering over that full selection window; the post-window row of Table 4 and the subsequent seven-year window supply the persistence evidence. Appendix A illustrates possible channels without estimating their parameters.

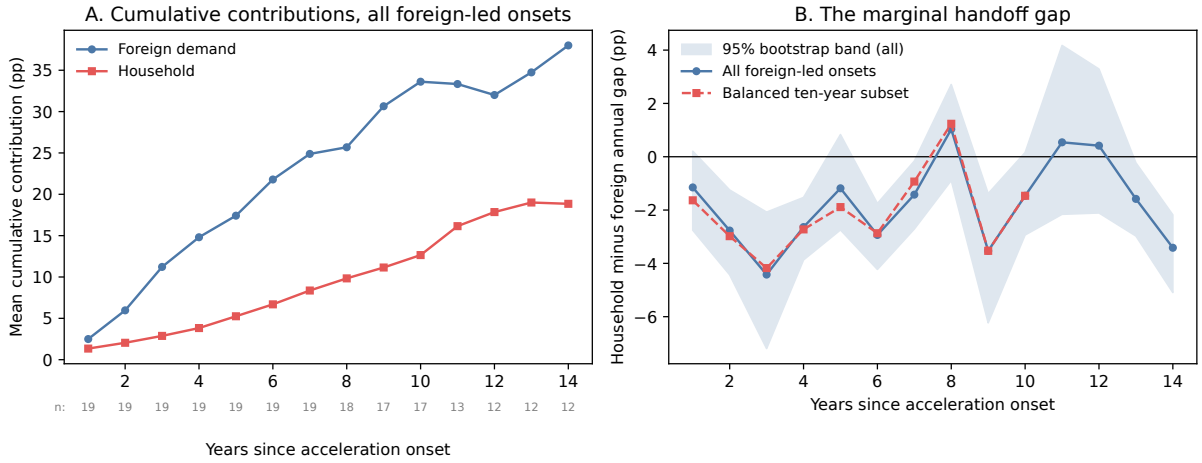


Figure 1: Cumulative and annual contributions after foreign-led accelerations

Notes: Panel A shows mean cumulative previous-year-price contributions of foreign and household demand, aligned to the acceleration onset, for the 19 foreign-led onsets; the gray row reports the number of episodes observed at each horizon. Panel B shows the mean annual household-minus-foreign gap with a 95 percent episode-bootstrap band; the dashed line restricts to the 17 episodes with complete ten-year windows, so composition changes do not drive the path.

Table 5: The purchaser engine in the window after an acceleration

Initial engine	N	C	G	I	X	Median growth	$\geq 2\%$
Household	13	7	0	1	5	3.1%	61.5%
Investment	2	1	0	1	0	5.9%	100.0%
Foreign demand	12	4	1	0	7	2.5%	66.7%

Notes: Codes C, G, I, X denote the four purchasers. The 27 accelerations for which a second complete seven-year purchaser window is observed. Columns count episodes by the engine of that next window; median growth is the PWT log-linear per-capita growth rate from $t+7$ through $t+14$.

Inventory change does not generate the short-window investment findings. When investment is restricted to gross fixed capital formation, both investment-led accelerations and all

seven investment-led broad-panel high-growth episodes remain investment-led. 9 of the 10 investment-led long-panel episodes remain so; 1 changes label. Appendix Table 33 reports the counts.

5 Foreign demand is not net exports

The four purchaser routes measured here treat foreign demand as the domestic value added that exports support. The conventional national-accounts engine is net exports instead, and the two are different objects. Appendix Proposition 4 settles the direction of the difference in advance: the foreign-demand contribution weakly exceeds the net-export contribution whenever imports induced by the three domestic purchasers do not fall. A sign disagreement is therefore available by construction, and what the sample has to establish is how often it occurs and how large it is. The condition is close to automatic but not vacuous: it holds in 35 of the 36 accelerations. The exception is Nigeria from 2002, where imports induced by the three domestic purchasers fall in total by 0.21 percentage point a year (household-induced imports fall by 0.29 point a year, or 2.1 points over the seven-year window, while government- and investment-induced imports rise slightly) and the measured wedge is -0.25 point a year (-1.7 points cumulatively)—a reminder that Appendix Proposition 4 governs the accounting benchmark and not every realization of it. Across the 36 episodes the two contributions—the second as defined in equation (6)—correlate at only 0.35. Foreign-demand-supported domestic value added rises in every one of the 36; the cumulative net-export contribution is nonpositive in 22 of them and in 11 of the 19 foreign-led accelerations. Within the foreign-led group, foreign demand contributes 3.55 percentage points per year on average, whereas the mean net-export contribution is -0.08 point. Foreign purchases can therefore support domestic value added even as imports rise with production and domestic absorption. This is precisely why "export-led" should not be read as "net-export-led." The acceleration sample shows that the wedge of Appendix Proposition 4, appendix equation (19)—contested in the Chinese three-engines debate of Section 2.3—is not specific to China. The import decomposition makes the mechanism observable. Table 6 sets the contribution comparison beside it, and Figure 2 plots the two contributions against each other across the 36 windows. In the 22 sign-disagreement episodes, gross imports add 5.77 percentage points per year; 2.01 points are induced by household demand, 1.76 by investment, 1.75 by the imported inputs embodied in foreign-demand-supported production, and 0.25 by government purchases. Imports rise on both sides of the economy rather than serving as a mechanical negation of the foreign-demand route.

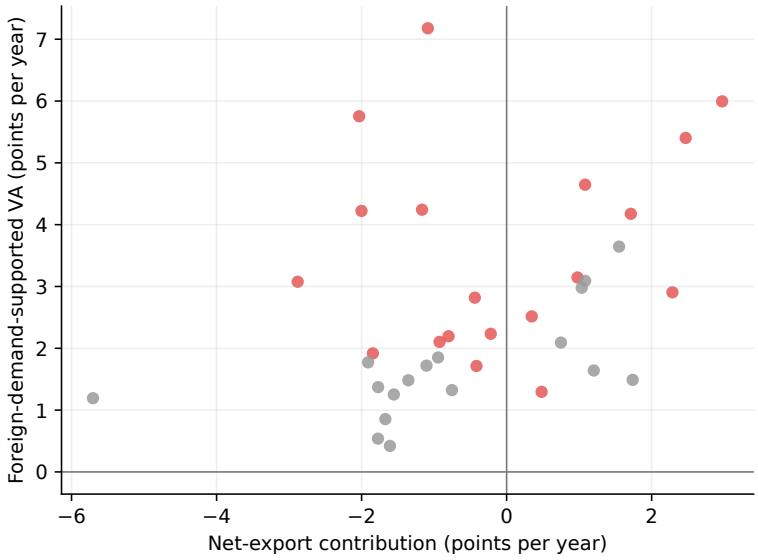


Figure 2: Foreign demand against net exports in the 36 acceleration windows

Notes: Each point is one of the 36 exact acceleration windows; red points are foreign-led. Both axes are average annual percentage points of domestic value added. Points in the upper-left quadrant have rising foreign-demand-supported value added and a nonpositive net-export contribution. The corresponding gradient across all 400 five-year episodes is reported in Appendix Figure 7.

Table 6: Foreign-demand-supported value added against the net-export contribution

Panel A. Average annual contributions during the acceleration window

Engine	N	VA-X	Exports	Imports	Net exports	Opposite signs
Household	15	1.61	2.11	2.96	-0.85	66.7%
Investment	2	2.25	3.18	3.24	-0.06	50.0%
Foreign demand	19	3.55	6.87	6.95	-0.08	57.9%

Panel B. Imports induced along the domestic chain, by purchaser

Sample	N	C	G	I	X	Imports
All accelerations	36	1.57	0.21	1.41	1.88	5.08
Foreign-demand-led	19	1.89	0.27	1.66	3.13	6.95
Opposite signs	22	2.01	0.25	1.76	1.75	5.77

Notes: Panel A entries are average annual percentage-point contributions over $t+1$ through $t+7$. Gross exports, gross imports, and net exports use the OECD current- and previous-year-price ICIO tables. The last column is the share with a positive cumulative foreign-demand-supported value-added contribution and a nonpositive cumulative net-export contribution. Panel B assigns imported final products and imported intermediates induced along the domestic production chain to the same four purchaser routes; entries are average annual percentage-point contributions.

6 Fast growth, initial states, and sample selection

6.1 High-growth episodes in the broad post-1995 panel

The long panel identifies historical transitions, but it omits many economies that grew rapidly after 1995. I therefore calculate the same previous-year-price purchaser contributions for all 80 economies in the balanced OECD panel. There are five complete five-year episodes from 1996–2000 through 2016–2020, or 400 country–episodes. High growth is defined independently of the purchaser accounts: for each period, I select the top quartile of compound real GDP growth from Penn World Table 11.0, measured from the year before the episode through its final year (Feenstra et al., 2015). This produces 100 high-growth episodes. A separate absolute definition retains the 153 episodes with at least 4 percent annual real GDP growth.

PWT and summed I–O growth correlate at 0.864 across the 400 episodes, with a median absolute annual-growth gap of 0.28 percentage point. Excluding the largest 5 percent of discrepancies leaves 90 high-growth episodes with similar median contributions. This checks aggregate growth, not the purchaser allocation.

Table 7 reports median annual contributions. The top-quartile episodes grow by 6.2 percent annually in both PWT and the I–O accounts. Household purchases contribute a median 2.4 percentage points, exports 2.1 points, investment 1.1 points, and government 0.6 point. These componentwise medians need not add to the median total, but they show that high growth is not simply high investment support. The absolute 4-percent rule gives the same ordering.

Table 7: Purchaser contributions in the broad panel's high-growth episodes

Episode group	N	PWT	I–O	C	G	I	X
Top growth quartile	100	6.2	6.2	2.4	0.6	1.1	2.1
Top quartile, validation sample	90	6.2	6.2	2.3	0.6	1.1	2.1
At least 4 percent growth	153	5.4	5.6	2.1	0.6	1.0	2.0
Other episodes	300	2.5	2.4	0.7	0.4	0.3	0.9

Notes: The sample contains 400 complete five-year country–episodes for 80 economies in 1996–2020. PWT growth is compound annual growth in constant-national-price real GDP. I–O growth and the four purchaser entries are median annual previous-year-price contributions, in percentage points. "Top growth quartile" is defined separately within each global period; the validation row excludes the largest 5 percent of absolute PWT–I–O growth gaps in the full sample; the 4-percent group overlaps the top quartile. Componentwise medians need not add to the median total. Across the 400 episodes, average annual PWT growth and the summed purchaser contributions correlate at 0.864 with a median absolute difference of 0.28 percentage point; PWT rather than I–O growth determines the high-growth sample, and excluding the largest 5 percent of absolute discrepancies leaves 90 high-growth episodes with virtually unchanged median contributions.

The identity of the largest contributor is more revealing. Households lead in 53 percent of high-growth episodes, exports in 38 percent, investment in 7 percent, and government in 2 percent. The difference between the first two is not statistically resolved; intervals follow below. Figure 3 divides countries into period-specific terciles using real GDP per capita in the year before the episode: households lead most lower-tercile high-growth episodes, while exports lead a majority of middle-tercile and most upper-tercile ones. The upper tercile contains only ten observations, so the gradient is descriptive rather than a stable development law.

The broad panel therefore changes the interpretation of familiar country labels. China is investment-centered in its recent purchaser state, and investment leads its growth in 2006–2010 and 2011–2015. But rapid growth more generally is not predominantly investment-led in domestic-value-added terms. Among high-growth episodes outside the long sample, Vietnam and Cambodia supply export-led examples; Myanmar, Côte d'Ivoire, and Bangladesh supply household-led examples; and Laos in 2011–2015 supplies an investment-led example. This heterogeneity is the result. It does not show that household demand causes growth: consumption can respond endogenously to income, credit, transfers, and earlier export or investment expansion.

Table 8 makes the state–engine distinction concrete for six familiar policy shorthands. A

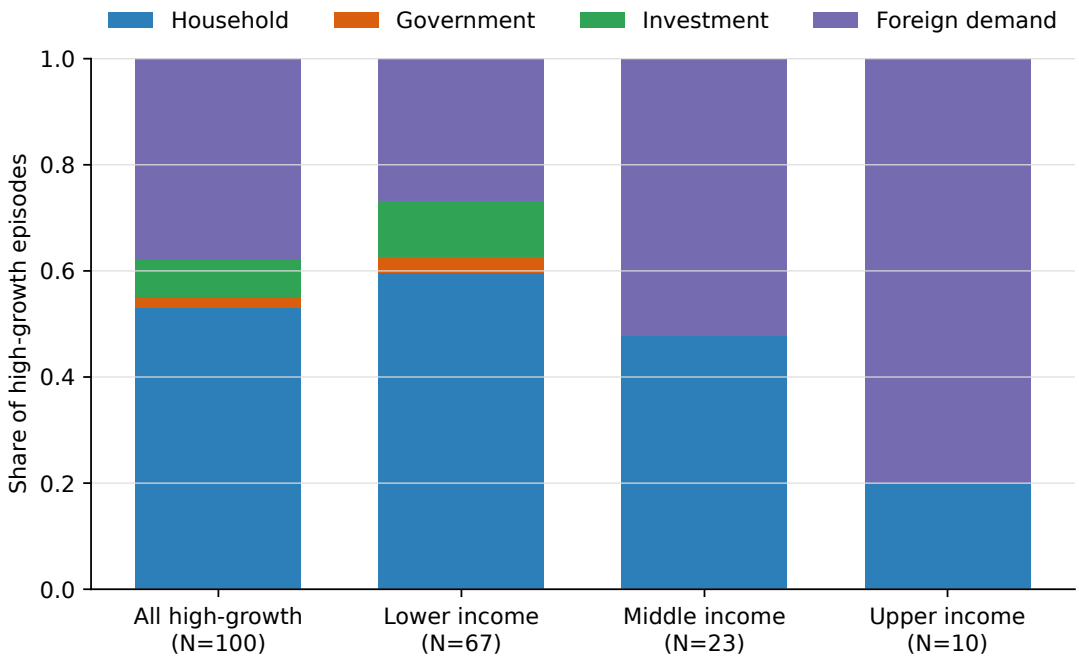


Figure 3: Leading purchaser in broad-panel high-growth episodes

Notes: High growth is the top period-specific quartile of PWT real GDP growth. Initial-income terciles use PWT output-side real GDP per capita in the year before each episode. Bars show the fraction for which each purchaser makes the largest cumulative real contribution. Counts are country–episodes.

label can describe a current state reasonably well and still conceal a changing marginal engine. Germany and Korea, for example, were led by foreign demand in every complete episode of 2001–2015 but hold household-led 2020–2022 states in the pandemic-affected endpoint window.

Table 8: Familiar labels versus measured purchaser facts

Economy	Conventional shorthand	2020–2022 state	2001–2020 engines
China	Investment/export led	I (35.7%)	X–I–I–C
United States	Consumption led	C (56.6%)	C–C–C–C
Germany	Export led	C (32.0%)	X–X–X–G
Korea	Export led	C (33.2%)	X–X–X–G
India	Consumption/services led	C (52.5%)	C–C–C–C
Sweden	Public/open economy	X (32.0%)	X–C–X–X

Notes: Codes C, G, I, X denote households, government, investment, and foreign demand. The state is the largest average purchaser-supported share in 2020–2022. Engine sequences list complete five-year episodes beginning with 2001–2005 and ending with 2016–2020. Conventional labels are policy shorthands, not estimated country types; the endpoint state includes the pandemic years.

6.2 How precisely are these frequencies measured?

The shares above are counts of episodes, and episodes repeat within countries, so the sampling unit is the country. Table 9 attaches a country-clustered bootstrap interval to each leading share and to the margin between the two leading routes, over 20,000 draws that resample countries with replacement.

The household frequency in high-growth episodes has a 95-percent interval of 39–67 percent across 47 countries; the foreign frequency in accelerations has an interval of 37–69 percent across 32 countries. The respective lead-over-runner-up intervals are $[-11, +41]$ and $[-21, +42]$ percentage points. Neither ordering is statistically resolved. An equal-share benchmark is uninformative here: households already lead 41.25 percent and foreign demand 47 percent of all 400 episodes, so four equally probable categories are not a relevant null.

Table 9: Country-clustered intervals for the leading-route frequencies

Sample	Route	Share / margin	95% interval
High-growth episodes	Households	53.0	[39.4, 67.0]
	Foreign demand	38.0	[25.0, 51.9]
	Difference	+15.0	[-11.2, +40.9]
Growth accelerations	Foreign demand	52.8	[36.8, 69.4]
	Households	41.7	[25.7, 58.3]
	Difference	+11.1	[-20.6, +42.4]

Notes: Bootstrap over 20,000 draws resampling countries with replacement and pooling their episodes, so a country contributing several episodes moves the statistic as one unit. Shares are of all episodes in each selected sample; difference rows are percentage-point margins. The intervals condition on the observed episode lists and do not account for common global shocks, redetection of onsets, or uncertainty in the source input–output tables.

Table 10: The cross-sample contrast has its own uncertainty

Acceleration minus high growth	Estimate	95% interval
Foreign frequency	14.78	[-0.15, 30.62]
Foreign minus household margin	26.11	[-4.78, 58.54]

Notes: Joint bootstrap with 100,000 draws and seed 20260905. The 54 countries in the union of the two selected samples are resampled with common multiplicities in both samples. Episode lists remain fixed. Entries are percentage points.

The modal route differs across the two observed lists, but that is not an established

population ranking reversal. Table 10 shows that both cross-sample intervals include zero. Global-period sensitivity is reported separately in Appendix Table 16: it addresses sample composition across time without pretending that a small number of global periods supplies precise time-cluster inference.

6.3 Does the engine merely reproduce the initial state?

The first condition is mechanical capacity: a purchaser that initially supports a large share of production has more scope to make the largest growth contribution. Table 11 makes this benchmark explicit. The initial largest purchaser and the realized engine differ in 33 of the 100 high-growth episodes. Engine–state coincidence is therefore important but incomplete, because levels and real changes are distinct objects.

Table 11: Initial-state benchmark for broad-panel growth engines

Sample or measure	N	C	G	I	X	Initial match
High growth: initial leader	100	66.0%	0.0%	2.0%	32.0%	—
High growth: engine	100	53.0%	2.0%	7.0%	38.0%	67.0%
Other episodes: engine	300	37.3%	7.7%	5.0%	50.0%	52.0%
Matched high growth: engine	52	55.8%	3.8%	7.7%	32.7%	61.5%
Matched controls: engine	52	55.8%	5.8%	1.9%	36.5%	67.3%

Notes: Codes C, G, I, X denote the four purchasers; initial match is engine–state coincidence. The first row classifies each high-growth episode by the largest purchaser share in the year before the episode. The remaining rows classify by the realized five-year engine. Matching is within the same global period, with replacement, on standardized additive-log-ratio purchaser shares and log real GDP per capita. A one-unit Euclidean caliper retains 52 high-growth episodes. "Same as initial" is the fraction for which the engine equals the initial largest purchaser. Among matched pairs, household leadership is 55.8 percent in both groups; export leadership is 32.7 versus 36.5 percent; investment leadership is 7.7 versus 1.9 percent, a difference of four versus one episodes. Standardized mean differences on the matched variables range from -0.13 to 0.22 for the four shares and are -0.10 for initial income.

The unmatched comparison makes household leadership look distinctive: it occurs in 53.0 percent of high-growth episodes but 37.3 percent of other episodes. Among the 52 matched pairs, however, household leadership is identical and export leadership similar across the two groups, and the small investment contrast rests on a handful of episodes (notes to Table 11). Neither the raw household frequency nor the investment contrast should be read as a causal return.

The remaining diagnostics reinforce the limits of reading the engine as a strategy. With genuinely earlier information, a recursive impulse forecast identifies 60.4 percent of later en-

gines, against 60.0 percent for persistence; adding predicted transmission does not improve accuracy. When the realized future impulse is supplied instead, accuracy rises to 86.3 percent, and base-year network conversion adds 3.3 percentage points. The latter is an oracle accounting exercise, not a forecast. Appendix H reports both. Country-held-out high-growth classification *conditional on the realized engine* also gains no accuracy from the network indices. Initial conditions and network conversion describe the accounts without providing a reliable independent forecast of the next growth engine.

7 Historical perspective and source quality

The historical extension asks whether purchaser reorientation and repeated engine changes predate the recent acceleration sample. In the 25-economy panel, the median within-country current-price household share falls by 11.0 percentage points between 1965–1969 and 2018–2022, while foreign support rises by 12.6 points. Raw and aligned endpoints have the same broad directions (Appendix Table 27). These are current-price compositions, not estimates of a secular volume reallocation.

Annual previous-year-price effects avoid chaining across decades. With $g_{ct} = \sum_j g_{ct}^j$ expressed in percentage points, the share that would obtain after the year's volume movement is

$$\widetilde{\alpha}_{ct}^k = \frac{\alpha_{c,t-1}^k + g_{ct}^k/100}{1 + g_{ct}/100}. \tag{12}$$

The annual volume effect is $\widetilde{\alpha}_{ct}^k - \alpha_{c,t-1}^k$. This calculation uses raw within-source states, not the log-ratio bridge. Table 12 reports country medians of within-country average annual effects. Household effects are negative and foreign effects positive in both sources. Government's pooled median is -0.033 percentage point per year, with signed-rank $p = 0.096$: a negative point estimate and no statistically clear rise, rather than proof of no increase. The OECD-only median is -0.089 point, negative in 19 of 25 economies ($p < 0.001$).

Table 12: Annual volume effects by source and quality screen

Source/sample	N	C	G	I	X
Pooled	25	-0.279	-0.033	-0.184	0.510
LR-WIOD	25	-0.345	0.018	-0.274	0.627
OECD	25	-0.257	-0.089	-0.028	0.383
Screened pooled	21	-0.279	-0.042	-0.133	0.494

Notes: Entries are percentage points per year, taking the median across countries of their mean annual effects. LR-WIOD covers annual changes in 1966–1995; OECD covers 1996–2022 for the same 25 economies. The screen is defined below and removes flagged economies from the pooled diagnostic. Componentwise medians need not sum to zero.

Aggregate reconciliation is insufficient validation. Historical solved value added closes to the source total within 10^{-15}, and the 150 pre-1996 episode growth rates correlate with PWT at 0.975 (median gap 0.32 point). China accumulates 234.5 points in the I–O accounts against 241.1 in PWT. Nevertheless, China's sum of separately chained route indices is 40.15 times the corresponding aggregate index in 1965–1975. Median annual ratios of current to previous-year-price export-supported value added imply revaluation of -30.64 percent in China and -19.96 percent in India in 1966–1995. These constructed ratios are not independently observed export prices.

I therefore flag a historical economy if any route value is nonpositive, if a within-source decade's ratio of summed route chains to the aggregate chain lies outside $[0.5, 2]$, or if the absolute median annual implied revaluation of any route exceeds 15 percent. The rule uses accounting properties rather than engine labels. Its thresholds were chosen during validation, not preregistered; tighter and looser screens are disclosed in Appendix Table 18. It is a sensitivity diagnostic, not certification of the remaining source estimates. Main acceleration results use OECD throughout and are unaffected by this historical screen. Appendix Table 17 lists the flagged economies and their diagnostics. Decade-chain compositions are retained only as qualified supplemental indices; they do not support the main volume inference.

Table 13: Engine changes across sources and the historical quality screen

Sample	Pairs	Changes	Switch (%)	C–X share (%)
All sources	250	104	41.6	66.3
Historical pairs	125	45	36.0	75.6
OECD pairs	100	46	46.0	52.2
Screened all sources	210	80	38.1	68.8

Notes: Pairs contain two consecutive complete five-year episodes. Historical and OECD rows exclude pairs crossing the source boundary. The screened row removes flagged historical economies throughout the long panel for a transparent country-exclusion diagnostic. The final column is the household–foreign share of all changes.

Across all 250 complete adjacent pairs, 104 change engine, a rate of 41.6 percent (country-bootstrap interval 32.8–50.4 percent). Household–foreign exchanges account for 66.3 percent of changes (interval 54.1–78.9 percent). Table 13 shows the source-specific and screened results. Longer episode windows and a 0.5-point winning margin reduce relabeling but leave substantial switching (Appendix Table 28). Historical data thus support the need to date purchaser engines, while their quality limits prevent early country sequences from carrying the main development argument.

8 Conclusion

Growth accelerations can be supported by domestic or foreign purchasers, and the route supporting the change need not be the largest purchaser in the initial state. The main post-1995 evidence does not reveal a general ten-year transition from foreign-led to domestically led growth. Selection builds a lead into the cumulative record, so the informative statistic is the engine after the selection window: among the 17 baseline foreign-led onsets with ten-year coverage, foreign demand still leads years eight to ten in 11 and households in 4. Some later seven-year windows become domestically led, which is compatible with persistent foreign leadership in the cumulative record. Distinguishing these objects changes the interpretation of a domestic handoff.

Price measurement matters for individual histories and for government leadership. Common deflation changes one-fifth of five-year assignments and doubles the number classified as government-led, while leaving household and foreign demand as the prevalent routes in rapid-growth episodes. Neither the separate frequency rankings nor their cross-sample reversal is precisely established. Initial purchaser composition and country size account for part of the descriptive differences; network indices add little independent forecasting value.

The policy implication is to distinguish a purchaser's domestic production exposure from the return to expanding its spending. At fixed production coefficients and purchaser bundles, a specified expenditure change maps to domestic value added through M^k. Financing, induced changes in other demand, capacity, productivity, and welfare require further evidence. The accounts can locate the purchasers supporting an observed acceleration and document whether that support persists. They cannot determine which policy would create the next one.

Data availability

The replication package contains the analysis code, the frozen derived inputs, source manifests, and a runner that reproduces every generated table and number from the frozen inputs. It will be deposited in a public repository on acceptance and is available from the author before then. The original data are the OECD inter-country input–output tables, LR-WIOD, and Penn World Table 11.0, all publicly available from their producers under their own terms; the manifest records their locations and access conditions, and fresh extraction from them is documented separately from reproduction from the supplied derived inputs.

A An illustrative transmission framework

This appendix records a conceptual interpretation of the accounting evidence. Its behavioral parameters are not estimated; the fixed-onset results do not identify income transmission or test a calibrated transition model.

To move from the accounting identity to a growth mechanism, take the base-year production structure as predetermined, suppress country and time subscripts, and write the purchaser index as a subscript. Let Y_0 be initial real value added and let $d_k = \Delta D_k / Y_0$ denote an externally or internally generated purchaser impulse. Let $\chi_k(z_k) \geq 0$ summarize scalable capacity—supplier availability, finance, labor, and infrastructure relevant to route k—and let $\Omega_k(d_k, z) \geq 0$ be displacement, crowding out, and other general-equilibrium leakage, expressed as a share of Y_0, with $\Omega_k(0, z) = 0$ so that a route with no impulse contributes nothing. Displacement is expressed against the route's own contribution; cross-route displacement would make the engine ranking non-separable and is deliberately left outside this local framework. The framework is stated for positive impulses ($d_k > 0$), which suffices for the high-growth and ignition applications. Contractions would require a separate sign convention, because released resources attenuate rather than amplify a decline. The route's real contribution is

$$r_k = M_k \chi_k(z_k) d_k - \Omega_k(d_k, z) \equiv B_k(d_k, z) d_k, \tag{13}$$

where $B_k = M_k \chi_k - \Omega_k / d_k$. The observed purchaser contribution in equation (5) is the empirical counterpart of r_k after the impulse and equilibrium responses have occurred. When $d_k > 0$ and expenditure and value added are measured in compatible real terms, the accounts permit the ex post calculation $B_k = r_k / d_k$. This realized domestic transmission coefficient is not thereby a causal multiplier: causal interpretation requires an externally generated purchaser impulse. Nor does B_k separately identify its capacity and displacement components. The single equation identifies only their composite; separating χ_k from Ω_k requires additional capacity measures, exclusion restrictions, or a structural model. Equation (13) is therefore a local transmission model, not a structural welfare model. Input–output accounts identify M_k and U_k. The distinction follows the network insight that the aggregate effect of a shock depends on both its location and equilibrium adjustment (Acemoglu et al., 2012; Baqaee and Farhi, 2019). The index U_k describes where conversion occurs and can proxy supplier reach entering χ_k, but it does not create additional aggregate value added once M_k is fixed.

This separation creates two different exercises. A forecast asks which purchaser impulse will arrive and which engine will lead. A conditional scenario instead takes a candidate impulse as given and asks how the existing domestic network converts it. For the forecast

exercise, B_k cannot simply be moved backward one period. Given information $\mathcal{I}_{t-1}$, define

$$\widehat{e}_{t|t-1} = \arg\max_k \left\{ \widehat{B}_{k,t|t-1} \widehat{d}_{k,t|t-1} \right\}. \tag{14}$$

For equal-sized purchaser scenarios, $\widehat{B}_k$ ranks effective domestic transmission, while predetermined M_k isolates its accounting component. Forecasting the actual next engine additionally requires the route-specific impulse $\widehat{d}_k$. Section E and Appendix Table 35 evaluate both the forecast and the conditional mapping.

Proposition 1 (A purchaser state need not be its growth engine)**.** *Assume unique maxima. Let* $s = \arg\max_k \alpha_k$ *be the largest component of the initial purchaser state and* $e = \arg\max_k r_k$ *the marginal engine. Then* $e \neq s$ *if and only if there is some* $j \neq s$ *such that*

$$r_j > r_s. \tag{15}$$

Consequently, a large initial purchaser share is neither necessary nor sufficient for that route to lead subsequent growth.

The result follows directly from the two definitions: s selects the largest level share, e the largest real contribution, and the two selections need not agree. Its substantive content is nevertheless important: α_k is a level, whereas d_k is an impulse and B_k is its domestic transmission. Divergence between state and engine is therefore the footprint of route-specific impulses and transmission, not an anomaly. Uniqueness is an assumption about the data as much as the framework; empirically, near-ties are treated as indeterminate under the 0.5-point winning margin of Section D.2. The 33 percent mismatch in the broad high-growth sample shows that the distinction is empirically consequential.

Proposition 2 (A country-size gradient for foreign demand)**.** *Suppose an accessible foreign-demand opportunity is* $\Delta D_X = \min\{\varepsilon_X Y_0, \overline{D}_X\}$*, while a domestic route* $q \in \{C, G, I\}$ *receives* $\Delta D_q = \varepsilon_q Y_0$*. Assume* $\Omega_X = \omega_X(z) d_X$*, so that* B_X *is invariant to the impulse, and hold* $B_X, B_q, \varepsilon_X, \varepsilon_q$*, and* $\overline{D}_X$ *fixed and positive. Once* $Y_0 > \overline{D}_X/\varepsilon_X$*,*

$$\frac{\partial (r_X - r_q)}{\partial Y_0} = -\frac{B_X \overline{D}_X}{Y_0^2} < 0. \tag{16}$$

The foreign-demand advantage therefore declines as economic size rises, conditional on market access and transmission capacity; with a size-invariant distribution of other margins, foreign-demand leadership becomes less likely.

The proposition is a scale result, not a universal population law. A large economy can remain export-led if its accessible market opportunity or export capacity rises with its size.

It nevertheless explains why the same absolute foreign opportunity has greater aggregate leverage in a small economy and motivates the observed PPP-GDP and population gradients.

Proposition 3 (Domestic handoff is conditional)**.** *Let a temporary foreign impulse follow* $d_{Xh} = \rho^h d_{X0}$ *with* $0 < \rho < 1$ *and* $d_{X0} > 0$*. Let export-supported income feed a domestic-demand stock according to* $w_{h+1} = \delta w_h + \lambda r_{Xh}$*, where* $w_0 = 0$*,* $0 \leq \delta \leq 1$*, and* $\lambda \geq 0$*. If* $r_{Xh} = B_X d_{Xh}$ *and* $r_{Ch} = B_C w_h$*, with* $B_X, B_C > 0$*, household demand overtakes foreign demand at horizon* h *if and only if*

$$B_C \lambda \sum_{\ell=0}^{h-1} \delta^{h-1-\ell} \rho^\ell > \rho^h. \tag{17}$$

A handoff is more likely when export income passes strongly into domestic demand (λ), domestic demand persists (δ), and domestic capacity converts it into production (B_C). It is not guaranteed at any finite horizon. The empirical horizon comparisons describe timing; without identified parameters they do not test this behavioral specification.

Proposition 4 (Accounting corollary: foreign-demand growth is not net-export growth)**.** *Let* $m_k \in [0, 1]$ *be the total imported content of an additional unit of purchaser demand, including imported intermediates, and suppose the imports induced by the three domestic purchasers do not fall in total,* $\sum_{k \in \{C,G,I\}} m_k d_k \geq 0$*, which holds in particular when the three domestic impulses are nonnegative. In the fixed-coefficient benchmark, with* $\chi_k = 1$ *and* $\Omega_k = 0$*, the normalized foreign-demand-supported contribution and conventional net-export contribution are*

$$r_X = (1 - m_X) d_X, \qquad n_X = (1 - m_X) d_X - \sum_{k \in \{C,G,I\}} m_k d_k. \tag{18}$$

Therefore

$$r_X - n_X = \sum_{k \in \{C,G,I\}} m_k d_k \geq 0. \tag{19}$$

Positive foreign-demand-supported value added can coexist with a nonpositive net-export contribution when import growth induced by domestic absorption is sufficiently large. (Because the benchmark sets $\chi_k = 1$ *and* $\Omega_k = 0$*, zero impulses are innocuous, and because the condition is on the total, one purchaser's induced imports may fall provided the others' rise by more.)*

This benchmark abstracts from the valuation and statistical-discrepancy terms handled in the empirical ICIO accounts. Its economic content is exact: net exports subtract imports

associated with all purchasers, whereas the foreign route subtracts only the imported content embodied in foreign sales.

Equation (13) also shows why network conversion alone cannot rank engines. Even if $M_a > M_b$, route b can contribute more when its demand impulse or scalable capacity is larger, or its displacement is smaller. The ranking by U_k is still less determinate because upstream breadth affects aggregate growth only through capacity and adjustment. With capacity and displacement unrestricted, the framework implies no ranking of high-growth incidence by M_k or U_k; that non-implication is a property of the framework, not a falsifiable prediction.

The four results have different epistemic status, and the empirical sections are read accordingly: Proposition 1 is definitional (the evidence shows the distinction is quantitatively important, not that a behavioral model is confirmed); Proposition 2 is a genuine comparative static (the acceleration evidence is consistent with it, and the broad panel does not provide a robust test); Proposition 3 is a dynamic hypothesis whose structural parameters are not econometrically identified here ; and Proposition 4 is an accounting corollary whose empirical wedge is a direct demonstration rather than a test.

For completeness, if a policy instrument p_j changes both demand and z, the differentiable version of equation (13) implies

$$\frac{\partial \sum_k r_k}{\partial p_j} = \sum_k \left(B_k + d_k \frac{\partial B_k}{\partial d_k} \right) \frac{\partial d_k}{\partial p_j} + \sum_k d_k \nabla_z B_k \frac{\partial z}{\partial p_j}.$$

Only a locally constant-B restriction eliminates the derivative terms. The paper's accounting scenario instead fixes ν, A^d, q^k and uses $\Delta V^k = M^k \Delta D^k$; it does not treat average realized transmission as an unrestricted policy derivative.

B Measurement sensitivity

Table 14: Engine reassignment under common deflation

PYP / common	C	G	I	X
Households	134	13	2	16
Government	1	18	1	5
Investment	1	6	13	2
Foreign demand	16	13	7	152

Notes: Rows are maintained PYP labels and columns common-deflator labels in the same 400 five-year episodes. Diagonal cells retain their label.

Table 15: Price sensitivity across global periods

Period	N	Changed	PYP G	Common G
1996–2000	80	12	2	3
2001–2005	80	23	2	5
2006–2010	80	14	9	15
2011–2015	80	8	1	3
2016–2020	80	26	11	24

Notes: Each global period contains 80 economies. Government counts use the two alternative price treatments in the same nominal accounts.

Table 16: Global-period and onset-cohort sensitivity

Sample	N	C (%)	G (%)	I (%)	X (%)
High growth: omit 1996–2000	80	55.0	2.5	8.8	33.8
High growth: omit 2001–2005	80	52.5	2.5	7.5	37.5
High growth: omit 2006–2010	80	51.2	1.2	7.5	40.0
High growth: omit 2011–2015	80	52.5	2.5	5.0	40.0
High growth: omit 2016–2020	80	53.8	1.2	6.2	38.8
Accelerations: 1997–2002	21	38.1	0.0	9.5	52.4
Accelerations: 2003–2015	15	46.7	0.0	0.0	53.3

Notes: High-growth rows omit one entire global period from the fixed high-growth list. Acceleration rows split the fixed onset list. Entries are episode frequencies; these checks do not estimate a time-clustered sampling distribution.

Table 17: Historical route-quality flags

Economy	Max. chain ratio	Median X revaluation	Nonpositive route
CHN	40.15	-30.64%	No
GRC	2.87	-7.49%	No
IND	1.91	-19.96%	No
PRT	1.01	6.95%	Yes

Notes: The diagnostic rule is stated in Section 7. The maximum chain ratio uses the three pre-1995 decades; revaluation is the country median of annual current/PYP export-supported-value-added ratios minus one. Nonpositive denotes any nonpositive supported route value. These are diagnostics of constructed accounts, not external price validation.

Table 18: Sensitivity to historical quality thresholds

Chain band	Revaluation bound	Retained economies	Switch (%)
[0.67, 1.5]	10%	18	39.4
[0.50, 2]	15%	21	38.1
[0.33, 3]	20%	23	40.4

Notes: In every row a nonpositive route value also triggers exclusion. Screens are diagnostic sensitivity rules selected during validation, not alternative estimates of a causal parameter.

C Supporting historical states and engine sequences

Table 19 compares five-year windows at the endpoints of the long sample. The first two columns are cross-country medians of purchaser shares. The third is the median within-country endpoint change.

Table 19: Purchaser states at the endpoints of the long panel

Purchaser	1965–1969	2018–2022	Median country change
Household	0.493	0.379	-0.110
Government	0.158	0.210	+0.063
Investment	0.190	0.139	-0.067
Foreign demand	0.131	0.255	+0.126

Notes: The sample contains 25 economies. Shares allocate domestic value added to the four purchaser routes. Endpoint shares are country medians in the indicated windows. The final column is the median country-specific change; componentwise medians need not sum to zero.

The main long-run direction does not depend on the bridge. Appendix Table 27 replaces the aligned pre-1995 shares with raw LR-WIOD shares while leaving the OECD endpoint unchanged. Household support falls by 13.8 rather than 11.0 percentage points in the median country, and export support rises by 12.0 rather than 12.6 points. Across countries, raw and aligned endpoint changes have correlations of 0.81–0.98; their signs agree in 92–100 percent of cases. The bridge matters for the exact magnitude of some country changes, especially household and government support, but not for the paper's directional long-run fact.

The household share falls by 11.0 percentage points within the median country, and the export share rises by 12.6 points. Government gains 6.3 points, while investment loses 6.7. These changes are economically large. They also clarify why final-expenditure-to-GDP ratios are insufficient: the purchaser accounts remove imported content and allocate domestic upstream value added to the purchaser route that draws it. The state is a current-price composition. Decade-chain indices in Table 20 are supplemental and can exhibit severe chain drift, even within ten-year windows. Their bounded shares do not validate the route-price allocation. Households lose chained volume share in 6 of 6 decade medians (range -6.21 to -1.32 points); foreign demand gains +2.98 to +7.42 points. Government current-price medians rise in 4 windows, while its chain-share median falls in 4 (range -1.68 to +0.47 points). These index summaries are not the main evidence. The unchained annual statistics use 1,425 country-years in 25 economies. Median effects are -0.279 points for households (22 negative, p <0.001), -0.033 for government (15 negative, p 0.096), -0.184 for investment (25 negative, p <0.001), and +0.510 for foreign demand (0 negative, p <0.001). Source-specific and screened estimates in Table 12 govern the interpretation.

Table 20: Supplemental current-price and chained-index compositions

	Households		Government		Investment		Foreign demand	
Decade	Price	Volume	Price	Volume	Price	Volume	Price	Volume
1965–1975	-4.17	-6.21	+3.23	+0.47	-0.87	-3.49	+2.56	+7.42
1975–1985	-1.77	-2.43	+0.55	+0.12	-2.81	-3.88	+4.31	+6.14
1985–1995	-2.01	-3.50	+0.25	-1.06	-0.33	-2.07	+1.94	+6.80
1995–2005	-2.87	-3.34	-0.41	-1.68	+0.14	+0.09	+2.89	+5.35
2005–2015	-0.77	-1.32	+0.62	-0.32	-0.84	-1.29	+2.09	+2.98
2015–2022	-3.24	-2.20	-0.45	-0.95	+0.53	-0.02	+2.81	+3.42

Notes: Median across economies of the within-country change over each decade, in percentage points. All 25 economies enter every window except 1965–1975, where Portugal's chained volume index is undefined because its 1975 investment contribution is −77.9 points. "Price" is the change in the current-price purchaser share. "Volume" is the change in a volume share: each route's real growth is chained from the base year of the window and expressed as a share of the four, which is bounded by construction. Windows never straddle the LR-WIOD/OECD boundary, since a change of source is not a price movement. Cells are separate medians, so they do not add across routes. Decade windows do not eliminate chain drift: China's first-decade ratio to the aggregate chain is 40.15. These indices are reported for transparency and are not treated as validated volume shares.

Using 2015–2019 instead of the pandemic-affected 2018–2022 endpoint gives the same direction: household support falls 9.6 points, foreign-demand support rises 12.1 points, government gains 5.8 points, and investment falls 6.2 points. Appendix Table 34 reports both terminal windows.

Figure 4 shows annual cross-country medians. The household share trends downward over most of the sample. Export support rises, with acceleration around the expansion of global production networks. Government support rises more gradually. Investment is cyclical but does not display a general secular rise. The vertical line marks the point at which the OECD series becomes the maintained source.

The broad OECD panel confirms that the recent part of this result is not an artifact of selecting countries with long historical data. Across 80 economies observed in 1995–2022, the median within-country change between three-year endpoint windows is −6.14 percentage points for households and +6.23 points for exports. Government rises by 0.93 points and investment falls by 1.08. The larger sample is useful for recent cross-sectional comparisons; the 25-economy sample is necessary for identifying multi-decade routes.

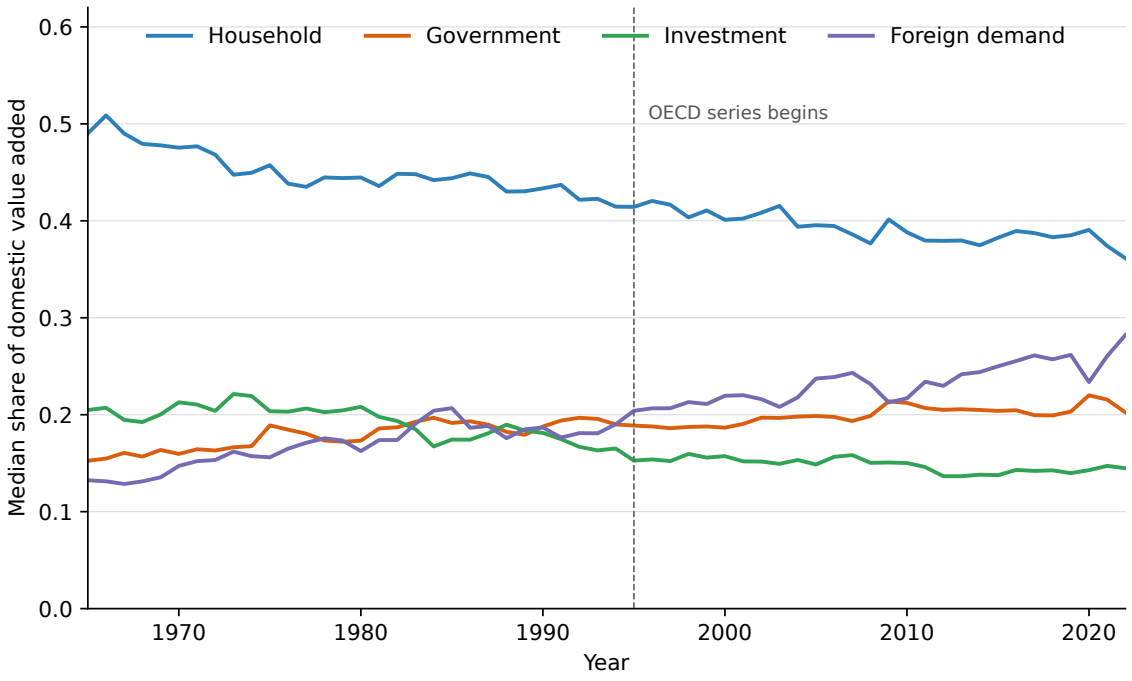


Figure 4: Median purchaser-supported value-added shares, 1965–2022

Notes: Annual medians across the balanced 25-economy sample. Pre-1995 LR-WIOD shares are aligned to OECD in additive-log-ratio coordinates using the 1995–2000 overlap. OECD values are used from 1995. The four shares sum to one within each country-year.

Current-price household and foreign shares move in opposite directions in the typical economy. Annual volume effects support the same directions, while government's pooled volume effect has an unresolved sign. None of these comparisons implies a universal sequence of development stages.

D Growth engines and transitions

D.1 Which purchaser made the largest real contribution?

Figure 5 reports the share of countries for which each purchaser is the leading five-year growth engine, and Appendix Table 26 gives the underlying counts for selected episodes. Household demand dominates the first episode, leading in 17 of 25 economies in 1966–1970. By 1996–2000, exports lead in 16. In 2016–2020, exports lead in 11, households in 6, government in 5, and investment in 3. The short 2021–2022 block is unusually export-led and should not be compared mechanically with a full five-year episode. The five government-led cases in 2016–2020 include the 2020 pandemic fiscal response and should not be read as a normal-times regime change.

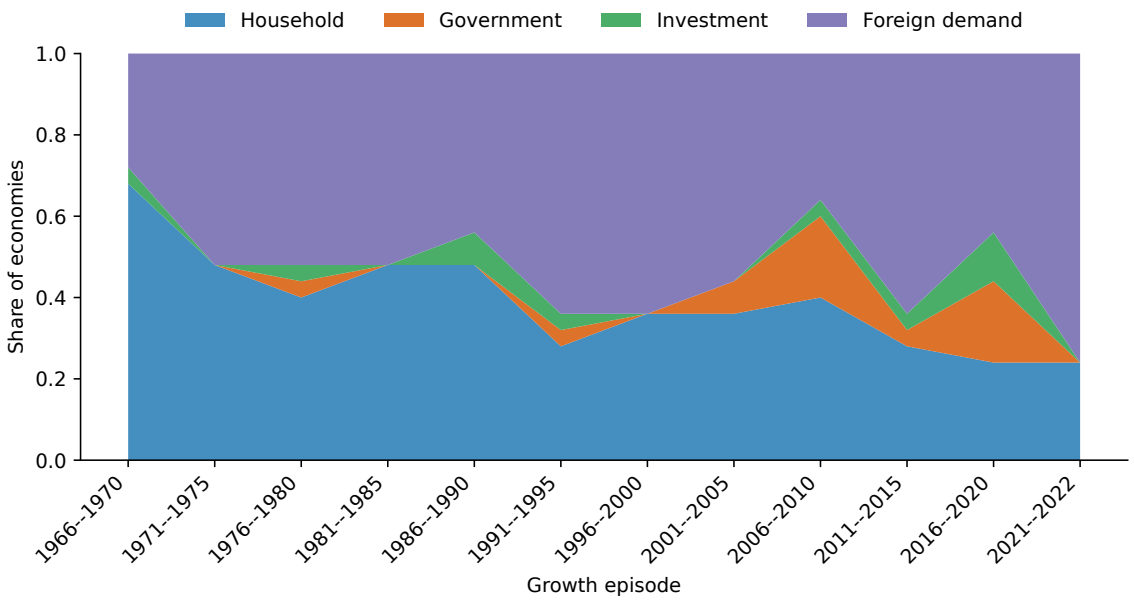


Figure 5: Share of economies by leading real growth engine

Notes: Fraction of the 25 economies for which the indicated purchaser makes the largest cumulative contribution to real value-added growth. Episodes contain five annual changes except 2021–2022. Contributions use previous-year-price input–output accounts.

Across the 275 complete country–episodes, exports lead in 139 and households in 111. Government leads in 15 and investment in 10. Counting the two-year 2021–2022 block as well would put exports at 158 against 117, but 19 of those 25 short episodes are export-led, so the wider margin is an artifact of the block rather than a feature of the panel. The small counts for government and investment do not mean that these purchasers are unimportant. They can make large positive contributions without exceeding household or export demand, and they can be decisive during particular recessions and recoveries. The result does mean that permanent four-type country labels would exaggerate how often government or investment is the single largest real engine.

D.2 Routes within countries

Figure 6 displays every country sequence. Some routes are persistent: of the 11 episodes each country contributes through the last complete window in 2020, household demand leads all 11 in the United States, exports lead all 11 in Belgium and 11 in Ireland, household demand leads 8 in India, and exports lead 10 in Austria and 8 in Sweden. Many others switch repeatedly. The historical sequences of flagged economies, including China and India, are displayed for transparency rather than treated as validated early turning dates. The screened aggregate results are in Table 13.

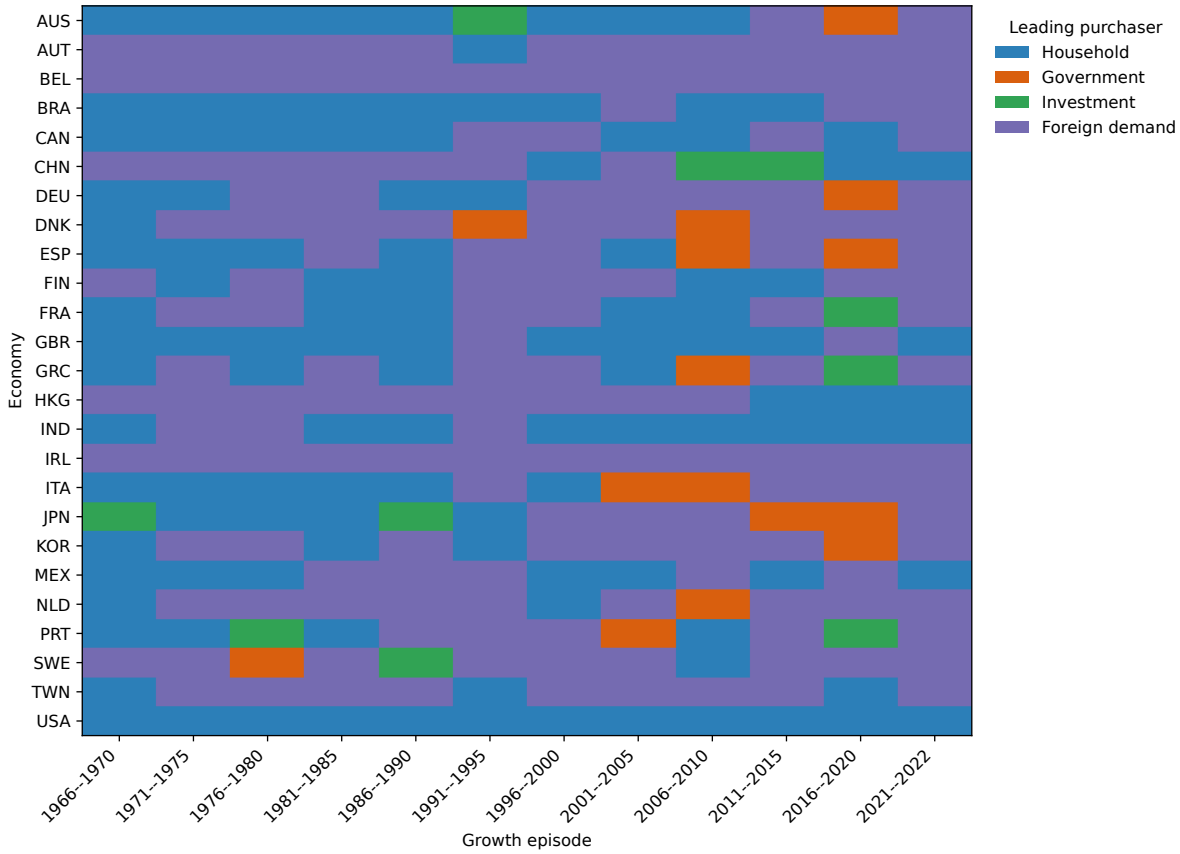


Figure 6: Country sequences of leading growth engines

Notes: Each cell records the purchaser with the largest cumulative real contribution to value-added growth in a country–episode. The last column is a two-year episode.

Table 21 counts adjacent episodes, using the 250 pairs in which each episode spans a full five years, so the short 2021–2022 block does not enter. The diagonal contains 146 persistent pairs; the remaining 104 change engine. Household-to-export transitions occur 40 times and export-to-household transitions 29 times. Together they represent 66.3 percent of all engine changes. Government leadership is left in 8 pairs, of which 7 go to exports. That last cell is thin, and deliberately so: pooling the two-year block would add five further departures from government leadership, but all five are the same transition—government over 2016–2020 followed by the export-led 2021–2022 stub—so counting them as evidence that public purchase hands off to exports would be close to circular.

Table 21: Transitions between adjacent growth-engine episodes

Previous \ next	Household	Government	Investment	Foreign demand	Total
Household	59	3	3	40	105
Government	1	2	0	7	10
Investment	5	0	1	1	7
Foreign demand	29	10	5	84	128
Total	94	15	9	132	250

Notes: Rows are the previous episode and columns the next episode. Diagonal cells indicate persistence. A pair enters only when both episodes span a full five years, which excludes every pair involving the two-year 2021–2022 block and leaves 250 pairs across 25 economies. Pooling that block instead gives 275 pairs and a switch share of 42.2 percent against 41.6 percent here. The household–export margin dominates on either basis, at 66.3 percent of changes here against 62.9 percent pooled; the one count pooling moves materially is the number of exits from government leadership, discussed in the text.

Observation 1 (The main margin is household versus export demand)**.** *Growth engines persist in 58.4 percent of complete adjacent episode pairs. When the engine changes, the dominant movement is between domestic household demand and foreign demand, not progression through four ordered stages.*

The switch rate is frequency-dependent, but the substantive result is not. Panel A of Appendix Table 28 applies the same complete-pair rule at other window lengths. The switch share is 48.0 percent with three-year windows, 41.6 percent with five-year windows, 40.8 percent with seven-year windows, and 35.0 percent with ten-year windows. Panel B addresses near ties. Requiring the winning purchaser to exceed the runner-up by 0.5 percentage point of cumulative real value-added growth classifies 11.6 percent of complete five-year episodes as indeterminate. Among adjacent pairs in which both engines remain determinate, 37.4 percent still switch. Some five-year relabeling is therefore noise or a near tie, but a substantial medium-run route-switching margin remains.

This finding changes the paper's empirical strategy. A four-cluster map is easy to display but fragile: a continuous composition is forced into mutually exclusive labels, and clusters can be driven by small differences in a category whose share is not large. The continuous state and real engine are interpretable without an arbitrary distance metric. Clustering can remain a robustness or exploratory device, but it should not define the contribution.

E Initial conditions and network diagnostics

E.1 Illustrative conditions associated with high-growth configurations

For this section only, I call an engine–episode pair a "high-growth configuration" when the engine is the realized engine in an episode that also belongs to the period-specific top quartile of PWT growth. This is an ex post joint outcome, not a welfare criterion and not an intervention. Within each realized engine, Table 22 compares predetermined conditions in high-growth and other episodes. Every condition is standardized within the global period, so an entry is the high-growth minus other-episode mean in standard-deviation units.

Table 22: Predetermined profiles of high-growth episodes, by realized engine

Condition	C (53/112)	G (2/23)	I (7/15)	X (38/150)
High-growth rate	32.1%	8.0%	31.8%	20.2%
Initial support of the realized engine	+0.21	-1.06	+1.56	+0.79
Initial log income per person	-0.97	-2.03	-1.69	-0.77
Initial household support	+0.21	+1.31	+0.96	-0.14
Initial government support	-0.78	-1.06	-1.26	-0.94
Initial investment support	+0.41	+0.31	+1.56	-0.51
Initial export support	+0.00	-0.95	-0.96	+0.79
Initial service share	-0.79	-0.42	-0.81	-0.93
PWT investment expenditure share	+0.03	-0.38	+0.79	+0.11
PWT gross trade openness	-0.33	-0.63	-0.76	-0.15
PWT human capital	-0.71	-2.14	-1.44	-0.94
Initial log capital per person	-0.85	-2.30	-1.82	-0.80
Initial log population	+0.55	-1.22	+1.25	-0.01

Notes: Columns C, G, I, X denote households, government, investment, and foreign demand; parentheses give high-growth/other episode counts. The high-growth rate is the fraction of episodes with the indicated engine that are in the top period-specific PWT growth quartile. All remaining entries are standardized high-growth minus other-episode differences within the same realized engine; variables are standardized within period before the comparison. Conditions are measured in the year before the episode. Human capital is unavailable for ten observations. The investment column rests on seven high-growth episodes, including China in 2006–2015 and several cases with weaker underlying accounts; the government column has two high-growth cases out of 25 government-led episodes. These are descriptive profiles, not estimated treatment-effect moderators.

Three provisional patterns are economically coherent. Household-led high growth appears in initially poorer, less-service-intensive, more populous economies—consistent with a large latent domestic market and broad income transmission, and equally consistent with consumption responding to growth generated elsewhere. Investment-led high growth is concentrated in poorer, capital-scarce economies with unusually high initial investment support,

a capacity-building configuration resting on very few episodes. Export-led high growth likewise occurs in poorer, less-service-intensive economies, and it is accompanied by an already large foreign-demand-supported domestic production base—a sharper marker than gross trade openness. The government column is too thin to support a general account of fiscal growth transmission (notes to Table 22).

The first row of the table needs a caveat that the others do not. A route's contribution is its initial support multiplied by its own growth rate, so a larger initial share raises that route's contribution mechanically, and the engine label is assigned on exactly that contribution. The row therefore mixes an accounting factor with whatever structural content it has, and it cannot be read as absorptive capacity on its own. Decomposing it does not resolve the problem either: the second factor, the route's own growth rate, is close to the outcome that defines a high-growth episode. The row is retained because it is the natural first thing a reader asks about, and it is duplicated exactly by the four route-specific rows below it, where the same quantities can be compared across engines rather than only along the diagonal.

The common idea is absorptive structure. Purchaser demand is more likely to coincide with rapid domestic growth when the economy already has the domestic market, capital-formation capacity, or supplier network that accompanies the translation of that demand into domestic value added. This is a descriptive profile of selected outcomes, not an estimate of whether a purchaser impulse works or why. That stronger question requires an externally generated purchaser shock, a first stage showing that the shock changes the intended purchaser contribution, and an interaction with a predetermined capacity measure. The interaction's outcome must be real output, employment, or productivity rather than the purchaser share alone. That externally identified exercise belongs in a separate causal paper.

E.2 Does the leading conversion route change across vintages?

The transmission framework treats M^k as predetermined. That is a statement about timing, not about persistence, and the panel can be asked the stronger question directly: over two decades of network change, does the purchaser with the highest marginal conversion stay the same?

It usually does. The five base years are 1995, 2000, 2005, 2010, and 2015. The leading route is the one the accounting definition makes likely. Government consumption is mostly non-market output with little imported content, so government purchases have the highest M^k in between 72 and 76 of 80 economies at every one of the 5 base years, and foreign demand leads between 4 and 7 economies. South Africa has household leadership in 2000,

Table 23: Marginal conversion falls for every purchaser, and least for government

Route	Median M^k 1995	Median M^k 2015	Mean change (p.p.)	Falls in	p	Rank corr.
Households	0.726	0.688	-4.79	63/80	<0.001	0.79
Government	0.893	0.870	-1.53	58/80	0.009	0.71
Investment	0.604	0.586	-2.59	52/80	0.009	0.77
Foreign	0.795	0.739	-4.17	59/80	<0.001	0.80

Notes: M^k is the network-conversion index of equation (8), measured at the 1995 and 2015 ICIO base years for the 80 economies present in both. "Falls in" counts economies whose conversion is lower in 2015. The p-value is for a paired test that the mean change is zero. "Rank corr." is the Spearman correlation of the cross-country ordering between the two years. These are accounting measurements, not estimated elasticities.

2005, and 2010; the non-government leaders are therefore not exclusively foreign demand. What the panel adds is that the leader almost never changes and that the exceptions are systematic. The leading purchaser is unchanged across all 5 vintages in 70 of 80 economies. Of the 10 that do change, 9 trade only between government and foreign demand, and they are predominantly commodity exporters whose export conversion is both high and volatile. The lead is not marginal: the median economy's leading route converts between 11 and 17 percent more of a unit of expenditure than its closest rival, in every vintage.

Levels are a different matter, and here the accounts show a clear common movement. Table 23 reports the change between the 1995 and 2015 base years. Marginal conversion falls for every purchaser: by 4.79 percentage points for households, 4.17 for foreign demand, 2.59 for investment, and 1.53 for government. A unit of final demand of any kind reached less domestic production in 2015 than in 1995, which is what two decades of lengthening international supply chains should do to an import-adjusted accounting measure.

The erosion is uneven, and its pattern is informative rather than surprising. Government conversion falls by 3.26 percentage points less than household conversion ($p < 0.001$) and 2.64 less than foreign conversion ($p < 0.001$); against investment the difference is 1.06 points and is not distinguishable from zero ($p = 0.15$). Government purchases are concentrated on public administration, education, and health, which are close to non-tradable and carry the lowest import content of the four expenditure components in the 18 OECD economies of Bussière et al. (2013), so the route whose bundle is hardest to import is also the route whose domestic content held up best. Relative positions are meanwhile persistent: the cross-country rank correlation between 1995 and 2015 is 0.79 for households, 0.71 for government, 0.77 for investment, and 0.80 for foreign demand.

The same panel answers a further question: which margin carries more of the variation in M^k, the economy or the purchaser? The design is balanced — 80 economies by four routes

— so the two main effects and the interaction are orthogonal. In 2015 the purchaser main effect accounts for 47.4 percent of the variance and the economy main effect for 40.8 percent, leaving 11.8 percent in the interaction; in 1995 the corresponding shares are 50.0 and 38.2 percent. The same ordering appears in the dispersions: the median within-economy standard deviation across the four routes, 0.122, exceeds the across-economy standard deviation once purchaser means are removed, 0.107. For the domestic content of a marginal demand impulse, then, which purchaser receives it matters at least as much as which economy does. That is a statement about the conversion map alone and carries no behavioural content: it says nothing about propensities to spend, and therefore nothing directly about the size of any multiplier.

Two implications follow for the rest of the paper. First, the identity of the highest-conversion purchaser is common to almost every economy and does not move, while the level of each route's conversion is a persistent country characteristic—rank correlations of 0.71 to 0.80 between 1995 and 2015—that drifts down by between 1.53 and 4.79 points over the two decades. Nothing here tests what moves it, and the index remains a base-year accounting measure rather than a structural parameter. Second, the stability is a property of the conversion map alone. It says nothing about which purchaser leads realized growth, which is governed by where demand actually goes: Section D.2 shows that the realized engine changes far more often than the leading conversion route does.

E.3 Network conversion, propagation, and conditional lift

For the descriptive episode evidence, I use a high-growth lift,

$$L_k = \frac{\Pr(H_{ce} = 1 \mid E_{ce} = k)}{0.25}, \tag{20}$$

where H_{ce} denotes the period-specific top PWT growth quartile. A lift above one is an ex post signal of compatibility between an engine and observed conditions; it is not part of the transmission framework and not the return to choosing that engine. Table 24 puts the three indices on a common scale. Network conversion and upstream propagation are measured in the year before each episode. The raw lift is equation (20). The conditional lift is the mean country-held-out predicted probability among episodes with the indicated realized engine, divided by 0.25. Because that model carries free realized-engine intercepts, the conditional lift largely reproduces the raw lift mechanically and should not be read as an adjusted comparison. The last column reassigns each episode to the indicated route: the episode's own characteristics are held fixed while the engine indicator and the three route-indexed features—initial support of that route, its network conversion, and its upstream

share—are set to the candidate route's values. It reports a model score after changing those inputs, not an identified counterfactual for making that route lead.

Table 24: Accounting indices and high-growth incidence

Purchaser	Conversion	Upstream	Raw lift	Conditional lift	Readiness	Episodes
Household	0.692	0.356	1.28	1.29	1.09	53/165
Government	0.874	0.231	0.32	0.41	0.63	2/25
Investment	0.594	0.440	1.27	1.25	0.94	7/22
Foreign demand	0.754	0.426	0.81	0.81	0.82	38/188

Notes: Conversion is network conversion, domestic value added per unit of purchaser expenditure. Upstream is the fraction of purchaser-supported domestic value added arising beyond the final-demand industries. Both entries are medians over 400 country–base-year cells per purchaser in 1995, 2000, 2005, 2010, and 2015. Raw lift is the observed high-growth rate divided by 0.25. Conditional lift uses nested country-held-out ridge-logit predictions for episodes with the realized engine; because that model includes realized-engine indicators, it tracks the raw lift closely. Readiness reassigns every episode to the indicated route—setting the engine indicator and the route-indexed features (initial support, network conversion, upstream share) to that route's values, with the episode's other characteristics unchanged—and averages the predicted probability. It varies the route rather than the sample. None of the three is a causal multiplier.

The accounting indices reveal why domestic retention and high-growth incidence must not be conflated. Government purchases have the highest median network conversion, 0.874, because much public consumption is domestically produced, and the lowest raw high-growth lift, 0.32. The two numbers should not be read against each other. Under the largest-contribution rule, government leads when the other routes are weak: in the broad panel, 20 of the 25 government-led episodes fall in 2006–2010 or 2016–2020, 16 have at least one other route contracting, and on the PWT measure that defines the lift 21 grow below their period median and 6 contract. Government leadership therefore marks episodes in which other demand was weak; its lift is a selection statistic, not a property of public purchases, and the cell is too thin to carry weight either way. The informative comparison is among the other three purchasers. Investment has the lowest conversion, 0.594, in part because investment goods are import-intensive, but the highest upstream share, 0.440, and a raw lift of 1.27. Household demand combines a conversion rate of 0.692 with a lift of 1.28. Exports have relatively high conversion and propagation, 0.754 and 0.426, but a raw lift of 0.81. High domestic content is therefore neither necessary nor sufficient for an engine to coincide with rapid aggregate growth. The transmission framework explains why no ranking is implied: conversion and upstream reach affect growth only through capacity and displacement, which

the accounting tables do not observe. The indices are insufficient reduced-form predictors in this sample and remain candidate moderators for an external-shock design.

The conditional signal is estimated with nested country-held-out folds. The baseline includes period and engine indicators; initial support of the realized engine; initial income, service share, investment expenditure, trade openness, human capital, capital per person, and population; and engine-specific slopes. Ridge penalties are selected inside each training fold. The network specification adds M^k and U^k and their engine interactions.

The result is a useful discipline. Classification conditional on the realized engine has substantial country-held-out signal: its AUC is 0.783. Because the realized engine and engine-indexed features enter the baseline, this is not a forecast using only information available before the episode. Adding network conversion and upstream propagation does not improve it; the AUC falls to 0.779 and both loss measures worsen slightly (Appendix Table 36). Country-held-out permutation importance shows that initial income carries the baseline prediction most strongly (AUC drop 0.085), followed by the service share (0.052) and initial support of the realized engine (0.035). In the network specification, conversion and upstream propagation reduce AUC by only 0.011 and 0.003 when permuted. The indices are therefore mostly redundant with income, sector structure, and initial purchaser support rather than wholly uninformative; Appendix Table 37 gives the full listing. The conditional lifts remain close to the raw lifts for the three non-government purchasers, which is what free engine intercepts imply and is therefore not evidence. The readiness column narrows the differences. Across the three purchasers with enough episodes to compare, the raw lifts run from 0.81 to 1.28 and readiness from 0.82 to 1.09, so the range narrows from 0.47 to 0.27; the government cell, on 25 episodes, sits below both ranges and is too thin to carry weight either way. Two cautions belong with this column rather than after it. Its spread is produced by the same two network indices that the next subsection finds add nothing to out-of-sample accuracy, and by initial support, which carries the accounting component noted above; and the underlying model is regularized, so the exact compression moves with the penalty the inner cross-validation selects. The column is a scenario device for reading the raw rates with more care, not an estimate of what a route delivers. The network indices are interpretable measures of retention and reach, but in this sample they are not stand-alone predictors of which realized engine will accompany high growth. Their stronger use is as predetermined moderators of externally generated purchaser shocks in equation (13).

E.4 Can the network predict or condition the next engine?

Two recursive exercises discipline the switching framework; their design, the realized-transmission construction, and the full accuracy grid are in Appendix H and Table 35. First, a pseudo-out-of-time implementation of equation (14) shows no incremental forecast value from the transmission coefficient: among determinate later episodes, the predicted-B signal differs from the impulse forecast alone by -1.4 percentage points, with a country-cluster bootstrap 95 percent interval of $[-5.6, 2.9]$. Second, an oracle diagnostic separates the demand and network margins: the realized future impulse alone identifies 86.3 percent of engines (89.6 percent of determinate episodes), and multiplying it by predetermined base-year network conversion raises the rates to 89.6 and 92.4 percent—gains of 3.3 and 2.8 percentage points, with bootstrap intervals $[0.8, 6.3]$ and $[0.0, 5.7]$, and only 12 of 240 assignments changing. Because the realized impulse uses future information from the same previous-year-price accounts, neither row is a forecast or a policy evaluation. The reading is nonetheless sharp: knowing where demand moves does most of the work, the predetermined network adds a modest positive conditional mapping gain, and the network is therefore better read as a scenario map once an impulse is observed or posited than as a standalone early-warning device.

F Panel construction and engine robustness

Table 25: The 80 economies of the broad panel

Code	Economy	Income	Code	Economy	Income	Code	Economy	Income
AGO	Angola	L	FIN*†	Finland	U	NGA	Nigeria	L
ARE	United Arab Emirates	U	FRA*†	France	U	NLD*†	Netherlands	U
ARG	Argentina	M	GBR*†	United Kingdom	U	NOR†	Norway	U
AUS*†	Australia	U	GRC*†	Greece	M	NZL†	New Zealand	U
AUT*†	Austria	U	HKG*	Hong Kong	U	PAK	Pakistan	L
BEL*†	Belgium	U	HRV	Croatia	M	PER	Peru	L
BGD	Bangladesh	L	HUN†	Hungary	M	PHL	Philippines	L
BGR	Bulgaria	M	IDN	Indonesia	L	POL†	Poland	M
BLR	Belarus	L	IND*	India	L	PRT*†	Portugal	M
BRA*	Brazil	M	IRL*†	Ireland	U	ROU	Romania	L
BRN	Brunei	U	ISL†	Iceland	U	RUS	Russia	M
CAN*†	Canada	U	ISR†	Israel	U	SAU	Saudi Arabia	U
CHE†	Switzerland	U	ITA*†	Italy	U	SEN	Senegal	L
CHL†	Chile	M	JOR	Jordan	L	SGP	Singapore	M
CHN*	China	L	JPN*†	Japan	U	STP	Sao Tome and Principe	L
CIV	Côte d'Ivoire	L	KAZ	Kazakhstan	L	SVK†	Slovakia	M
CMR	Cameroon	L	KHM	Cambodia	L	SVN†	Slovenia	M
COD	D.R. Congo	L	KOR*†	Korea	M	SWE*†	Sweden	U
COL†	Colombia	L	LAO	Laos	L	THA	Thailand	M
CRI†	Costa Rica	M	LTU†	Lithuania	M	TUN	Tunisia	L
CYP	Cyprus	U	LUX†	Luxembourg	U	TUR†	Türkiye	M
CZE†	Czechia	M	LVA†	Latvia	M	TWN*	Taiwan	U
DEU*†	Germany	U	MAR	Morocco	L	UKR	Ukraine	L
DNK*†	Denmark	U	MEX*†	Mexico	M	USA*†	United States	U
EGY	Egypt	L	MLT	Malta	M	VNM	Vietnam	L
ESP*†	Spain	M	MMR	Myanmar	L	ZAF	South Africa	M
EST†	Estonia	M	MYS	Malaysia	M			

Notes: Every economy reported separately in the OECD inter-country tables. A dagger marks the 38 OECD members as of 2026; an asterisk marks the 25 long-panel economies with LR-WIOD coverage from 1965. Income is the tercile (L, M, U) of PWT output-side real GDP per capita in 1995, cut within the 80.

Table 26: Leading growth engines in selected episodes

Episode	Household	Government	Investment	Foreign demand
1966–1970	17	0	1	7
1986–1990	12	0	2	11
1996–2000	9	0	0	16
2006–2010	10	5	1	9
2016–2020	6	5	3	11
2021–2022	6	0	0	19

Notes: Entries count countries. The four entries in each row sum to 25. The leading engine maximizes the purchaser's cumulative real growth contribution within the episode.

Table 27: Long-run endpoint changes with and without database alignment

Purchaser	Aligned change	Raw change	Correlation	Same sign	Median gap
Household	-11.0	-13.8	0.87	92%	2.8
Government	+6.3	+8.2	0.98	100%	2.3
Investment	-6.7	-5.9	0.92	96%	0.7
Foreign demand	+12.6	+12.0	0.81	92%	0.7

Notes: Changes and median gaps are percentage points between 1965–1969 and 2018–2022. "Aligned" applies the country-specific additive-log-ratio bridge; "raw" uses unadjusted LR-WIOD shares before 1995. Correlations and sign agreement compare the 25 country-specific endpoint changes. The median gap is the median absolute difference between raw and aligned changes.

Table 28: Growth-engine robustness to episode length and near ties

Specification	Episodes	Indeterminate	Adjacent pairs	Switch share
A. Complete episodes by window length				
3-year windows	475	—	450	48.0%
5-year windows	275	—	250	41.6%
7-year windows	175	—	125	40.8%
10-year windows	125	—	100	35.0%
B. Complete five-year episodes with a minimum winning margin				
0.00 pp margin	275	0 (0.0%)	250	41.6%
0.25 pp margin	275	17 (6.2%)	219	39.7%
0.50 pp margin	275	32 (11.6%)	195	37.4%
1.00 pp margin	275	60 (21.8%)	156	30.1%

Notes: Panel A retains only episodes with the stated number of annual changes. Panel B uses complete five-year episodes and labels an episode indeterminate when the largest cumulative purchaser contribution is less than the stated margin above the second largest. Its switch share uses adjacent pairs for which both engines are determinate. Margins are percentage points of real value-added growth; the pipeline stores them as fractions, so the 0.5-point margin appears in the generated files and in `growth_engine_summary.json` as 0.005.

G Acceleration episodes and robustness

Table 29: Growth-acceleration onsets and purchaser engines

Economy	Onset	Exact-window engine	Follow-on engine	Cohort
Belarus	1997	Foreign demand	Household	1997–2002
Croatia	1997	Foreign demand	Government	1997–2002
Estonia	1997	Foreign demand	Foreign demand	1997–2002
Greece	1997	Household	Foreign demand	1997–2002
Kazakhstan	1997	Foreign demand	Household	1997–2002
Latvia	1997	Household	Foreign demand	1997–2002
Lithuania	1997	Household	Foreign demand	1997–2002
Cambodia	1998	Foreign demand	Household	1997–2002
Russian Federation	1999	Household	Household	1997–2002
Bulgaria	2001	Foreign demand	Foreign demand	1997–2002
China, Hong Kong SAR	2001	Foreign demand	Household	1997–2002
Czechia	2001	Foreign demand	Foreign demand	1997–2002
Romania	2001	Household	Foreign demand	1997–2002
Singapore	2001	Foreign demand	Foreign demand	1997–2002
Slovakia	2001	Foreign demand	Foreign demand	1997–2002
Thailand	2001	Foreign demand	Foreign demand	1997–2002
Türkiye	2001	Investment	Household	1997–2002
China	2002	Investment	Investment	1997–2002

Table 29: Growth-acceleration onsets and purchaser engines (continued)

Economy	Onset	Exact-window engine	Follow-on engine	Cohort
Colombia	2002	Household	Household	1997–2002
Indonesia	2002	Household	Household	1997–2002
Nigeria	2002	Household	Household	1997–2002
Argentina	2003	Household	Household	2003–2015
Morocco	2003	Household	Foreign demand	2003–2015
Peru	2003	Household	Household	2003–2015
Viet Nam	2003	Foreign demand	Foreign demand	2003–2015
Bangladesh	2005	Household	Household	2003–2015
Lao People's DR	2006	Household	Investment	2003–2015
Malta	2011	Foreign demand	–	2003–2015
Hungary	2012	Foreign demand	–	2003–2015
Ireland	2012	Foreign demand	–	2003–2015
Latvia	2012	Foreign demand	–	2003–2015
Lithuania	2012	Foreign demand	–	2003–2015
Philippines	2012	Household	–	2003–2015
Cyprus	2014	Foreign demand	–	2003–2015
Romania	2014	Household	–	2003–2015
Croatia	2015	Foreign demand	–	2003–2015

Table 30: Growth-acceleration engines by onset cohort

Onset cohort	Episodes	Household	Government	Investment	Foreign demand
1997–2002	21	38.1%	0.0%	9.5%	52.4%
2003–2015	15	46.7%	0.0%	0.0%	53.3%

Notes: The split isolates the 1997–2002 transition and early-globalization cohort from onsets in 2003–2015. Entries are shares within each cohort.

Table 31: Country size and foreign-demand leadership in all 400 episodes

Size measure	Outcome	Unadjusted	S.E.	With controls	S.E.
Log PPP GDP	Foreign-led indicator	-0.035	0.032	-0.013	0.031
Log PPP GDP	Foreign advantage (points)	0.134	0.234	0.132	0.181
Log population	Foreign-led indicator	-0.111	0.035	-0.013	0.032
Log population	Foreign advantage (points)	-0.337	0.229	0.132	0.181

Notes: Coefficients are per one standard deviation of the indicated size measure. All regressions include period indicators; controlled specifications also include initial income per person and the initial foreign-demand-supported share. Standard errors are clustered by country. The foreign advantage is the foreign-demand contribution minus the largest of the other three annual contributions. Because log PPP GDP equals log population plus log income per person, the two controlled size specifications are reparameterizations of one another and therefore produce nearly identical coefficients.

Table 32: Acceleration engines under alternative thresholds and horizons

Specification	N	Household	Government	Investment	Foreign demand
Baseline	36	41.7%	0.0%	5.6%	52.8%
Post growth at least 3.0 percent	42	42.9%	0.0%	4.8%	52.4%
Post growth at least 4.0 percent	34	41.2%	0.0%	5.9%	52.9%
Acceleration at least 1.5 points	40	42.5%	0.0%	7.5%	50.0%
Acceleration at least 2.5 points	31	38.7%	0.0%	6.5%	54.8%
Five-year pre/post windows	53	35.8%	0.0%	3.8%	60.4%
Ten-year pre/post windows	24	58.3%	0.0%	0.0%	41.7%

Notes: Each row re-runs the onset detector and requires a complete post-onset purchaser window of the stated horizon. The baseline uses seven-year pre- and post-windows, post growth of at least 3.5 percent, and an acceleration of at least 2 percentage points. Candidate onsets before the OECD purchaser panel are excluded.

Foreign demand remains the modal route under the nearby five- and seven-year specifications and threshold changes. It leads exactly half under the looser 1.5-point acceleration rule. With ten-year windows, households lead 58.3 percent and foreign demand 41.7 percent, in a newly

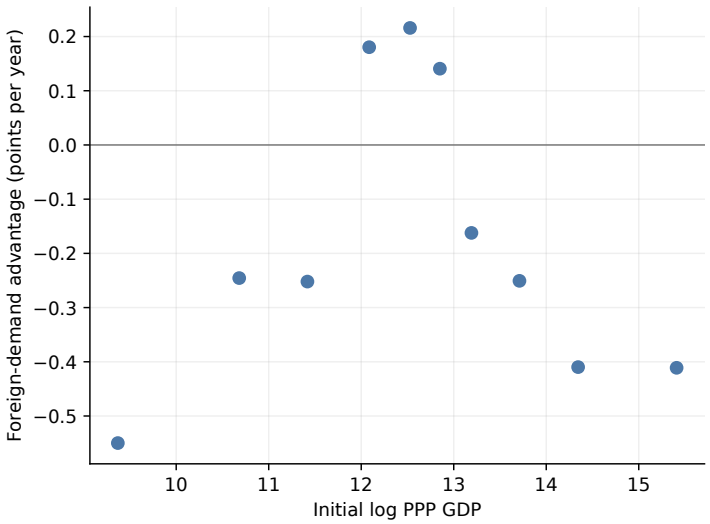


Figure 7: The size gradient across all five-year episodes

Notes: Decile means from all 400 complete five-year episodes. The vertical outcome is the foreign-demand contribution minus the largest other purchaser contribution. The binned pattern is downward-sloping, but the corresponding observation-level slope is not statistically significant, and the population gradient in Table 31 does not survive controls. The figure is reported here rather than in the main text because the relation it displays is one the text declines to claim.

selected set of 24 onsets. This row changes selection and cannot measure a handoff within the baseline 36; the fixed-onset comparison is Table 4.

Table 33: Investment-engine robustness excluding inventories

Sample	Original I-led	GFCF-only I-led
Long-panel I-led	10	9
Broad-panel high-growth I-led	7	7
Exact-acceleration I-led	2	2

Notes: "GFCF-only" replaces gross capital formation plus inventory change with gross fixed capital formation when ranking the four cumulative purchaser contributions. LR-WIOD and OECD source tables preserve the split.

H Endpoint and prediction diagnostics

Table 34: Long-run purchaser states with a pre-pandemic terminal window

Purchaser	1965–1969	2015–2019	Δ to 2015–2019	2018–2022	Δ to 2018–2022
Household	0.493	0.386	-0.096	0.379	-0.110
Government	0.158	0.203	+0.058	0.210	+0.063
Investment	0.190	0.142	-0.062	0.139	-0.067
Foreign demand	0.131	0.260	+0.121	0.255	+0.126

Notes: Entries are cross-country medians of five-year average purchaser shares for the balanced 25-economy long panel. Change columns report the median within-country terminal-minus-1965–1969 change, not the change between the two cross-country medians. The 2015–2019 window excludes the pandemic; 2018–2022 is the maintained endpoint.

Two constructions underlie the forecasting and oracle diagnostics summarized in the main text and reported in Table 35. First, the purchaser accounts permit an ex post check on the composite transmission coefficient in equation (13). I construct the real expenditure impulse from current- and previous-year-price OECD tables and normalize each annual change by prior-year current value added before summing both the impulse and the domestic-value-added contribution over each five-year episode. When the resulting cumulative positive impulse is at least 0.02, the median realized $B_k = r_k/d_k$ is 0.581 for households, 0.814 for government, 0.495 for investment, and 0.681 for foreign demand; respectively, 84, 61, 67, and 80 percent of route–episodes meet this threshold. These are descriptive transmission ratios—ratios of five-year sums rather than bounded fractions—not causal multipliers. The medians sit below the corresponding conversion medians of Table 24 for every route, with the same route ranking; this is consistent with $\chi_k \leq 1$ and positive displacement in $B_k = M_k\chi_k - \Omega_k/d_k$, although the samples and normalizations differ, so the comparison is indicative rather than a test.

Second, the recursive signal implements equation (14) for the 2006–2010, 2011–2015, and 2016–2020 episodes. For each route, ridge models forecast d_k and B_k using only earlier episodes, lagged route outcomes, base-year purchaser support, the network indices, income, sector structure, investment, trade openness, human capital, capital per person, and population; penalties are chosen within the expanding training sample using held-out countries. Across 240 later episodes, the impulse forecast alone identifies the next engine in 60.4 percent of cases; predicted $B_k\widehat{d}_k$ and lagged $B_k\widehat{d}_k$ each reach 59.2 percent, while persistence reaches 60.0 percent. Restricting attention to 211 episodes in which the realized winner exceeds the runner-up by at least 0.5 percentage point raises the rates to 65.9 percent for the impulse

forecast, 64.5 percent for both B signals, and 63.0 percent for persistence.

Table 35: Forecasting and oracle diagnostics for the next purchaser engine

Signal	Engine hit	Switch hit	Switch precision	Switch recall	Switch destination
Previous engine persists	60.0	60.0	–	0.0	0.0
Largest initial purchaser	57.9	68.3	63.5	49.0	22.9
Forecast impulse only	60.4	67.5	62.9	45.8	28.1
Base-year $M_k \times \widehat{d}_k$	58.3	65.8	60.0	43.8	25.0
Lagged $B_k \times \widehat{d}_k$	59.2	66.3	61.5	41.7	24.0
Predicted $B_k \times \widehat{d}_k$	59.2	66.3	61.2	42.7	25.0
Oracle: realized d_k only	86.3	91.7	89.6	89.6	76.0
Oracle: base-year $M_k \times d_k$	89.6	92.5	89.8	91.7	84.4

Notes: Entries are percentages across 240 country–episodes in 2006–2010, 2011–2015, and 2016–2020. Route-specific ridge models use only earlier episodes; penalties are chosen in country-held-out folds within each expanding training sample. A switch is defined relative to the preceding five-year engine. Switch-destination accuracy is conditional on an actual switch. M_k is the base-year domestic network-conversion index. Realized $B_k = r_k/d_k$ is defined for episodes whose cumulative impulse—annual impulses normalized by lagged value added and summed over the episode—is at least 0.02; its forecast uses only information observed before the episode. The final two rows instead use the realized future impulse; they are oracle accounting diagnostics, not forecasts or policy evaluations. None of the rows is a causal estimate.

Table 36: Country-held-out classification given the realized engine

Country-held-out specification	Brier score	Log loss	AUC
Given realized engine	0.144	0.455	0.783
Add network indices	0.148	0.468	0.779

Notes: The outcome is membership in the period-specific top PWT growth quartile. Countries, rather than observations, are held out in ten outer folds; the ridge penalty is chosen using five inner country folds. Lower Brier score and log loss and higher area under the receiver-operating-characteristic curve (AUC) indicate better classification. Both specifications condition on realized-engine indicators and route-indexed features; neither is a strictly pre-episode forecast.

Table 37: Permutation importance in the held-out prediction model

Predetermined feature	Baseline AUC drop	Network-model AUC drop
Initial support of realized engine	0.035	0.016
Initial income per person	0.085	0.041
Initial service share	0.052	0.034
PWT investment share	0.012	0.009
Gross trade openness	-0.008	-0.007
Human capital	-0.006	0.000
Capital per person	0.028	0.024
Population	0.002	-0.002
Network conversion	–	0.011
Upstream propagation	–	0.003

Notes: Entries are mean declines in the pooled country-held-out AUC when the indicated feature is permuted in held-out observations, leaving each fold's fitted ridge-logit model unchanged. Means use 30 permutations. Negative entries indicate that permuting the feature slightly improves prediction.

I Selected continuous country paths

Figure 8 displays six economies chosen to illustrate different routes, not to estimate types. The figure emphasizes the continuous state and makes visible both slow reallocations and temporary movements. Early annual variation reflects the harmonized LR-WIOD construction and should be interpreted at medium-run frequencies.

J Reproducibility

The broad 1995–2022 accounts, long current-price state, long growth engines, and broad high-growth episodes are generated separately:

```
python code/build_routes.py
python code/build_long_run_routes.py --cached
python code/build_growth_engines.py
python code/build_broad_engine_conditions.py
python code/build_acceleration_dynamics.py
python code/build_handoff_path.py
python code/build_engine_effectiveness_indices.py --cached
```

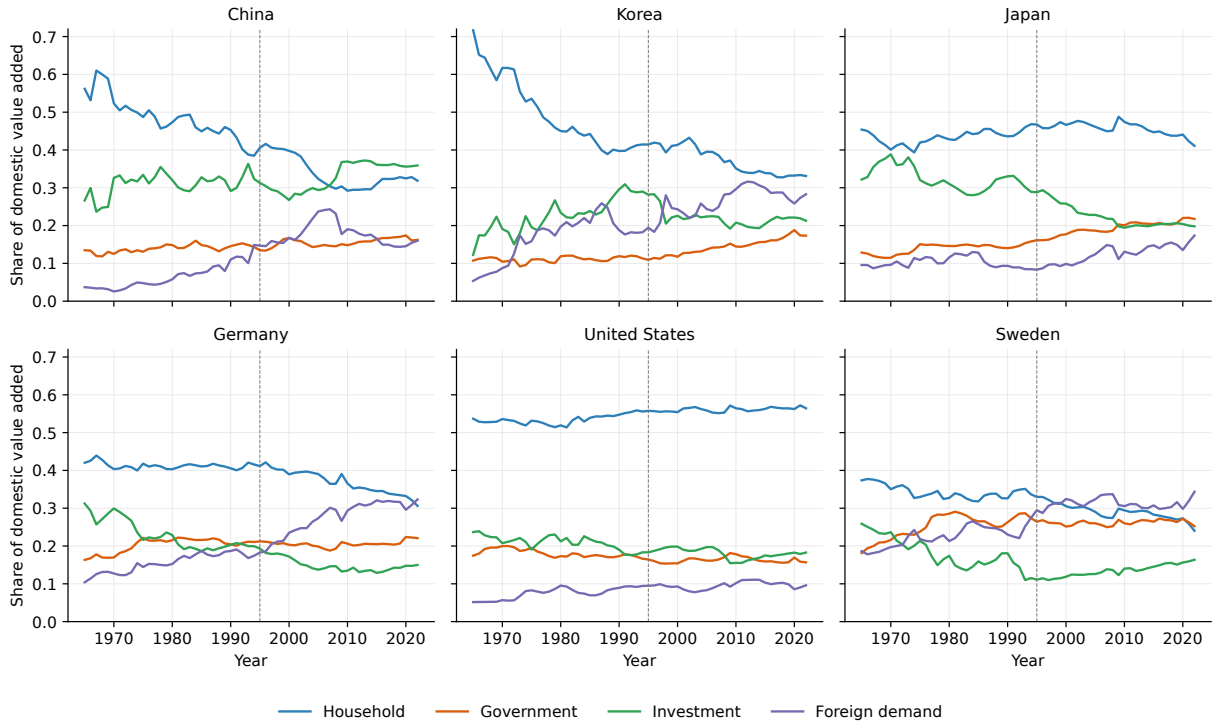


Figure 8: Purchaser-state paths in selected economies

Notes: Shares of domestic value added supported by each purchaser. The vertical line marks the switch to the OECD current-price series in 1995. Countries are illustrative and were not selected by an estimated cluster.

```
python code/build_conversion_stability.py
python code/build_engine_switch_signal.py
python code/build_headline_intervals.py
python code/build_jde_revision.py
python code/beamer_figures_cn.py
```

The commands above reuse the frozen historical and network caches. Removing `--cached` reruns those two source extractions. The growth and switch-signal scripts also reuse existing price-account caches; the replication README distinguishes a cached analysis build from a fresh source build. The JDE revision script produces the common-deflator comparison, fixed-onset horizons with post-window engines, the sample disclosures, joint intervals, quality screens, the economy list, the macro file, and the deck's tables.

The route script builds the purchaser states from the OECD current-price purchaser account; the long-run script extracts and aligns the LR-WIOD states; the growth script computes the annual previous-year-price contributions for both panels, validates episode growth against PWT, and classifies the high-growth episodes; the conditions script builds the initial-state benchmark and the matched comparison; the acceleration and handoff scripts build the

onset sample, its follow-on windows, the net-export comparison, and the continuous paths; the effectiveness-index, stability, and switch-signal scripts construct the network indices, their stability across base years, and the recursive signal; and the interval script attaches the country-clustered intervals. The stability script runs after the effectiveness-index script, and the interval script after the growth and acceleration scripts. Generated tables and numbers are stored in `output/` and figures in `figures/`. The OECD current-price purchaser account is the one constructed in Gan (2026); its annual panels are read without modification.

References

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

Alesina, A. and Wacziarg, R. (1998). Openness, country size and government. *Journal of Public Economics*, 69(3):305–321.

Baccaro, L. and Hadziabdic, S. (2024). Operationalizing growth models. *Quality & Quantity*, 58:1325–1360.

Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem. *Econometrica*, 87(4):1155–1203.

Bems, R., Johnson, R. C., and Yi, K.-M. (2010). Demand spillovers and the collapse of trade in the global recession. *IMF Economic Review*, 58(2):295–326.

Berg, A., Ostry, J. D., and Zettelmeyer, J. (2012). What makes growth sustained? *Journal of Development Economics*, 98(2):149–166.

Bussière, M., Callegari, G., Ghironi, F., Sestieri, G., and Yamano, N. (2013). Estimating trade elasticities: Demand composition and the trade collapse of 2008–09. *American Economic Journal: Macroeconomics*, 5(3):118–151.

Comin, D., Lashkari, D., and Mestieri, M. (2021). Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374.

Eichengreen, B., Park, D., and Shin, K. (2012). When fast-growing economies slow down: International evidence and implications for China. *Asian Economic Papers*, 11(1):42–87.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The next generation of the Penn World Table. *American Economic Review*, 105(10):3150–3182.

Gan, L. (2026). Who buys structural transformation? Purchaser exposure and sectoral responses. Working paper, Texas A&M University.

Hausmann, R., Pritchett, L., and Rodrik, D. (2005). Growth accelerations. *Journal of Economic Growth*, 10(4):303–329.

Herrendorf, B., Rogerson, R., and Valentinyi, Á. (2014). Growth and structural transformation. In Aghion, P. and Durlauf, S. N., editors, *Handbook of Economic Growth*, volume 2, pages 855–941. Elsevier.

Johnson, R. C. and Noguera, G. (2012). Accounting for intermediates: Production sharing and trade in value added. *Journal of International Economics*, 86(2):224–236.

Jones, B. F. and Olken, B. A. (2008). The anatomy of start-stop growth. *Review of Economics and Statistics*, 90(3):582–587.

Kongsamut, P., Rebelo, S., and Xie, D. (2001). Beyond balanced growth. *Review of Economic Studies*, 68(4):869–882.

Koopman, R., Wang, Z., and Wei, S.-J. (2014). Tracing value-added and double counting in gross exports. *American Economic Review*, 104(2):459–494.

Kranendonk, H. and Verbruggen, J. (2008). Decomposition of GDP growth in some European countries and the United States. *De Economist*, 156(3):295–306.

Lin, J. Y. and Li, Y. (2003). Exports and economic growth in China: A demand-oriented analysis (in Chinese). *China Economic Quarterly*, 2(4).

National Bureau of Statistics of China (2025). What are the "three engines" pulling economic growth? (in Chinese). Statistical knowledge series, stats.gov.cn; accessed August 2026.

Ngai, L. R. and Pissarides, C. A. (2007). Structural change in a multisector model of growth. *American Economic Review*, 97(1):429–443.

OECD (2025). OECD inter-country input-output (ICIO) tables, 2025 edition. Inter-Country Input-Output database, `https://oe.cd/icio`; accessed August 2026.

OECD (2026). Inter-country input-output tables at previous year prices, 2026 edition. Inter-Country Input-Output database. Paired current-price and previous-year-price tables; accessed August 2026.

Shen, L. (2009). Evaluating the pulling effects of the "three engines" (in Chinese). *The Journal of Quantitative & Technical Economics*, 26(4):139–151.

U.S. Bureau of Economic Analysis (2026). Table 1.1.2. Contributions to percent change in real gross domestic product. National Income and Product Accounts, apps.bea.gov/iTable; accessed September 2026.

Woltjer, P., Gouma, R., and Timmer, M. P. (2021). Long-run world input-output database: Version 1.1 sources and methods. GGDC Research Memorandum 190, Groningen Growth and Development Centre. Dataset DOI: 10.34894/A7AXDN.