Online Appendix

Exchange-Rate Transmission through Multinational Firms: Evidence from Japan

Ryan Kim[*] Bin Ni[†] Hyunseung Oh[‡] Choongryul Yang[§]

[*]Seoul National University. E-mail: rkim59@snu.ac.kr.

[†]Hosei University. E-mail: bin@hosei.ac.jp.

[‡]Federal Reserve Board. E-mail: hyunseung.oh@frb.gov.

[§]Federal Reserve Bank of Atlanta. E-mail: choongryul.yang@atl.frb.org.

A Firm-Level Data and Empirics

A.1. Summary Statistics

Table A.1 reports summary statistics for the matched parent-affiliate data on Japanese multinational firms in 2020. The sample consists of parent firms that own at least one foreign affiliate, appear in the parent-firm panel with a branch employment decomposition, and have a constructed exchange-rate shock.

Table A.1: Summary Statistics: Affiliate Networks and Parent Workforce

	Mean	SD	p10	p25	p50	p75	p90
Panel A: Foreign affiliate networks							
Number of affiliates per parent	4.9	10.1	1.0	1.0	2.0	5.0	10.0
Number of host countries per parent	3.4	4.5	1.0	1.0	2.0	4.0	8.0
Affiliate employment	270	1,244	2	8	40	164	522
Affiliate sales (million yen)	9,888	57,729	72	305	1,279	5,079	18,043
Ownership share of affiliates (%)	90.2	20.0	51.0	97.9	100.0	100.0	100.0
Share of current income paid to parent (%)	74.2	160.4	0.0	5.1	31.5	75.0	153.3
Revenue share: local host market (%)	67.4	38.8	0.0	30.8	90.1	100.0	100.0
Revenue share: Japan (%)	17.0	31.2	0.0	0.0	0.1	15.6	80.8
Revenue share: third countries (%)	15.6	28.8	0.0	0.0	0.0	15.3	67.3
Panel B: Parent workforce							
Parent employment	1,099	3,267	88	148	316	793	2,055
Employment composition shares (% of own parent employment)							
Head-office share	43.2	28.4	9.4	18.0	37.0	67.2	87.1
Non-production support share	22.5	23.4	1.8	5.0	13.5	32.7	61.2
Other employment share	34.3	30.4	0.0	4.0	28.6	61.7	79.2
of which production worker	25.1	28.8	0.0	0.0	10.7	48.8	71.3

Notes: All statistics are for 2020, the pre-depreciation baseline year. The common parent sample consists of firms that own at least one foreign affiliate, appear in the parent-firm panel with a branch employment decomposition, and have a constructed exchange-rate shock. Panels A and B describe the same 2,210 parent firms, covering 10,776 affiliates in 102 host countries. The first two rows of Panel A, affiliates per parent and host countries per parent, are computed at the parent-firm level, with one observation per parent ($N = 2,210$). All other Panel A statistics are computed at the affiliate level, with one observation per affiliate ($N = 10,776$). Ownership share is the parent's equity stake in the affiliate. Share of current income paid to parent is affiliate payments to the parent as a fraction of affiliate current profit. Because this ratio is unstable for affiliates with near-zero profits, it is winsorized at the 99th percentile; the median is unaffected. Revenue shares from the local host market, Japan, and third countries sum to 100 percent. Panel B describes the same parent firms on the workforce side; 62 percent are in manufacturing, defined as JSIC 091–329. Composition rows report, across firms, the distribution of each category's share of its own parent's employment.

A.2. Data Construction

Survey coverage and linkage. The parent survey (BSJBSA) covers firms with at least 50 employees and at least 30 million yen in capital, spanning mining, manufacturing, electricity and gas, wholesale, retail, and selected service industries; roughly 25,000–30,000 firms respond each year, with a response rate above 80 percent. The affiliate survey (BSOBA) covers all Japanese firms operating at least one foreign affiliate, excluding parents in finance, insurance, and real estate.[5] A foreign affiliate is a foreign entity in which Japanese investors hold at least 10 percent of equity; majority-owned affiliates of majority-owned subsidiaries are also covered. When several Japanese firms invest in the same affiliate, the questionnaire is filed by the largest Japanese investor. The two surveys share a permanent firm identifier, which we use to link each affiliate to its parent.

Cleaning. We drop affiliate-year observations in which total sales and total purchases are both unreported or both zero; this effectively removes affiliates that are not in operation. The geographic decomposition of revenues and purchases is built only from affiliates that report the corresponding components: when a total is unreported, its components are treated as unreported, and we verify that recovered totals track reported totals closely (correlations above 0.95 for every component). For parent firms that switch industries over time, we fix the industry at its final-year classification, since the survey classifies firms by their largest-revenue product in the previous year. Parent total employment is the sum of head-office employment, employment at establishments outside the head office, and employees dispatched to other firms; parent wages include bonuses and exclude retirement benefits. Bilateral exchange rates are from the IMF International Financial Statistics, measured as yen per unit of foreign currency. Transaction shares are measured over 2016–2020 and exposure weights are fixed at their 2020 values; outcomes are 2020–2022 log changes. Table A.2 traces the path from the cleaned affiliate panel to the baseline estimation samples.

Where the exposure comes from. Table A.3 reports the thirteen largest host countries by affiliate revenue, their currencies, and the geographic decomposition of revenues and purchases.

Persistence of exposure shares. Table A.4 reports, for each exposure share, the R^2 from regressing the share on affiliate fixed effects only and on year fixed effects only, in the affiliate-year panel for 2015–2020.

A.3. Additional Evidence on the Reallocation Mechanism

In the model, parent non-production support labor is an input into foreign affiliate production, and a depreciation that raises the value of affiliate operations pulls this labor toward affiliate-related activity. The evidence matches this reallocation: total parent employment is roughly flat

[5] The sector composition table nonetheless shows seven parents in the broad finance and real-estate grouping, 0.4 percent of the estimation sample. The industry code there is the parent's own classification in the parent survey, which is assigned by main line of business and need not coincide with the frame the affiliate survey screens on, and the broad grouping also covers adjacent industries such as leasing. Dropping these parents leaves the parent-level estimates unchanged, for example the average-wage coefficient moves from 0.169 to 0.168 and the support-employment coefficient from 0.384 to 0.387.

Table A.2: Sample Construction

	Count
Affiliate side	
Affiliates in cleaned panel, 2016–2020	34,094
with non-missing, non-zero sales or purchases	28,904
with constructed exchange-rate shocks (2020 base)	13,939
in the baseline employment regression	9,240
Parent side	
Parents matched to the parent survey	38,747
with parent-level shocks	3,416
in the baseline wage regression	1,811

Notes: Counts of affiliates (affiliate side) and parent firms (parent side) at each step of the sample construction. The affiliate rows count distinct affiliates in the cleaned 2016–2020 panel, those surviving the reporting screen, those with constructed 2020-base exchange-rate shocks, and those in the baseline affiliate employment regression of Table 1. The parent rows count firms matched to the parent survey, those with constructed parent-level shocks, and those in the baseline wage regression of Table 2. The summary-statistics sample of Table A.1 lies between the last two parent rows; it additionally requires the branch employment decomposition.

Table A.3: Composition of Affiliate Exposure by Host Country

		Revenue				Purchases			
Host country	Currency	Share	Local	Japan	Third	Share	Local	Japan	Third
USA	USD	0.30	0.64	0.03	0.33	0.29	0.54	0.37	0.09
CHN	CNY	0.17	0.64	0.14	0.23	0.17	0.68	0.24	0.08
THA	THB	0.08	0.56	0.13	0.31	0.08	0.72	0.18	0.11
SGP	SGD	0.05	0.32	0.21	0.47	0.05	0.35	0.25	0.40
IDN	IDR	0.03	0.56	0.12	0.32	0.03	0.66	0.19	0.15
GBR	GBP	0.03	0.46	0.08	0.46	0.03	0.30	0.23	0.46
HKG	HKD	0.03	0.42	0.16	0.42	0.04	0.28	0.45	0.28
BEL	EUR	0.02	0.27	0.02	0.71	0.02	0.07	0.85	0.08
DEU	EUR	0.02	0.61	0.03	0.36	0.02	0.23	0.51	0.25
AUS	AUD	0.02	0.72	0.11	0.17	0.02	0.45	0.48	0.06
CAN	CAD	0.02	0.45	0.02	0.54	0.02	0.33	0.34	0.33
KOR	KRW	0.02	0.68	0.09	0.23	0.02	0.57	0.35	0.08
TWN	TWD	0.02	0.68	0.09	0.23	0.02	0.41	0.44	0.15

Notes: Thirteen largest host countries by aggregate affiliate revenue, averaged over 2016–2020, which together span the twelve currencies in Figure 2. Share is the country's share of aggregate affiliate revenue (purchases). Local, Japan, and Third decompose the country's affiliate revenue (purchases) by counterpart market and sum to one within each block.

while non-production support employment rises, and affiliate employment, dispatched workers, and investment expand.

This appendix adds two further tests of the mechanism and examines two alternative channels. First, the reallocation in the model sends parent labor toward affiliate operations, and

Table A.4: Persistence of Exposure Shares

	R^2: Affiliate FE	R^2: Year FE
Network weights		
Affiliate share of parent revenue	0.975	0.000
Affiliate share of parent purchases	0.969	0.000
Transaction-geography shares		
Revenue share: local market	0.907	0.000
Revenue share: Japan	0.924	0.000
Revenue share: third countries	0.875	0.000
Purchase share: local market	0.878	0.000
Purchase share: Japan	0.887	0.000
Purchase share: third countries	0.845	0.000
Third-country regional shares		
Third-country revenue: North America	0.832	0.000
Third-country revenue: Asia	0.893	0.000
Third-country revenue: Europe	0.886	0.000
Third-country revenue: other regions	0.801	0.000
Third-country purchases: North America	0.847	0.000
Third-country purchases: Asia	0.870	0.000
Third-country purchases: Europe	0.877	0.000
Third-country purchases: other regions	0.716	0.000

Notes: Each row reports the R^2 from regressing the indicated share on affiliate fixed effects only (first column) and on year fixed effects only (second column), in the affiliate-year panel for 2015–2020. Network weights are each affiliate's share of its parent's total affiliate revenue and purchases, the outer weights used in the nested shift-share. Transaction-geography shares decompose each affiliate's revenue and purchases across the local host market, Japan, and third countries, and the third-country regional shares split third-country revenue and purchases across North America, Asia, Europe, and other regions.

dispatched workers, parent employees posted to foreign affiliates, are where this flow is directly recorded. Table A.5 shows that the dispatched-worker response to the revenue-side shock is roughly three times larger for sales- and services-oriented affiliates (wholesale and retail trade, and business services), the same commercial functions counted as non-production support employment at the parent. The differential is specific to parent-provided labor: the dispatched-worker increase is much smaller for manufacturing affiliates, and affiliate employment shows no differential for either affiliate type, whether measured as the survey total, which includes dispatched workers, or net of them. What moves is parent-supplied support component.

Second, if the rise in non-production support employment serves foreign operations, it should be concentrated among parents whose foreign operations are largest. Table A.6 splits parents at the median of their total foreign affiliate employment in 2020. Columns 1 and 2 estimate the parent regression separately on each group, and column 3 pools them and interacts both shocks with an above-median indicator. Parents below the median show no support response, while their counterparts above the median do. The pooled differential is significant

Table A.5: Affiliate-Side Response by Affiliate Type: Sales/Services vs. All Other Affiliates

	Disp Emp	Aff Emp	Local Emp
	(1)	(2)	(3)
$\Delta \ln E_a^S$	0.113*	0.207***	0.210***
	(0.064)	(0.071)	(0.072)
$\Delta \ln E_a^S \times$ Sales/Serv	0.254**	-0.095	-0.123
	(0.107)	(0.129)	(0.126)
Implied effect, Sales/Serv	0.367	0.112	0.087
Host-country FE	✓	✓	✓
N	6,542	9,448	9,326

Notes: Affiliate-level long-difference (2020–2022) regressions of the indicated outcome on the revenue-side shock $\Delta \ln E_a^S$ and its interaction with an indicator for *sales/services* affiliates (wholesale and retail trade and business services; affiliate industry codes 230, 280, 290), with manufacturing and other affiliates as the omitted group. The first row reports the effect for manufacturing and other affiliates, the second the additional effect for sales/services affiliates, and the "implied effect" row their sum. Column (1), Disp Emp, is the number of workers dispatched from the parent to the affiliate (parent-provided labor); column (2), Aff Emp, is total affiliate employment as reported in the survey, which includes dispatched workers. Column (3), Local Emp, nets dispatched workers out of the total; we treat the survey total as the primary measure since the net measure differences two reported series. All specifications include host-country fixed effects and the purchase-side shock $\Delta \ln E_a^P$, also interacted with the type indicator. Standard errors, two-way clustered by host country and parent firm, are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

at the ten percent level, and the purchase-side interaction is not significant, so the pattern is specific to the revenue side. We read the split as suggestive.

An alternative channel for the parent-side response is financial rather than organizational: more exposed parents repatriate affiliate income and spend the windfall on domestic hiring. If so, the response should be stronger for initially more financially constrained firms. Table A.7 interacts the revenue-side shock with the parent's standardized initial leverage (total liabilities over total assets), a standard firm-level measure of financial constraint (Giroud and Mueller, 2017). The interaction is insignificant for every outcome: more leveraged parents do not respond more strongly.

A second alternative is a vertical intermediate-input channel. A depreciation lowers the foreign-currency cost of Japanese inputs, so affiliates that source more from the parent country could expand for a standard trade-cost reason rather than through parent-side support. The purchase-side shock absorbs average cost-side exposure as a level control, but the saving on Japanese inputs grows with the same local depreciation that drives the revenue-side shock, so the two could move together across affiliates. To test this channel, Table A.8 interacts the revenue-side shock with the affiliate's Japan-sourcing intensity, its purchases from Japan relative to total purchases. The vertical channel predicts a positive interaction, and we find none: across the same seven outcomes as the affiliate baseline, the interaction is insignificant everywhere

Table A.6: Parent Support Employment and Foreign Affiliate Employment

	NP support employment		
	(1)	(2)	(3)
	Below median	Above median	All parents
$\Delta \ln E_f^S$	0.060	0.574**	0.088
	(0.241)	(0.253)	(0.223)
$\Delta \ln E_f^S \times$ Above median			0.629*
			(0.318)
$\Delta \ln E_f^P$	0.017	0.199	-0.022
	(0.233)	(0.309)	(0.215)
$\Delta \ln E_f^P \times$ Above median			0.101
			(0.388)
Industry FE	✓	✓	✓
N	449	659	1,123

Notes: The dependent variable is the long-difference (2020–2022) in log non-production support employment at the Japanese parent. Parents are split at the median of total foreign affiliate employment in 2020, the summed employment of their affiliates. Columns (1) and (2) estimate the regression separately on the below- and above-median groups, so the controls and industry fixed effects are free to differ across groups; column (3) pools the two and interacts both the revenue- and purchase-side shocks with an above-median indicator. All columns match the specification of the main parent regressions in Table 2, with the baseline parent controls (purchase-side shock, affiliate revenue and purchase coverage, export/import shares, R&D, revenue, employment) and industry fixed effects; standard errors clustered by industry are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Parent-Side Spillback and Initial Financial Position

	(1)	(2)	(3)	(4)	(5)	(6)
	Avg Wage	Total Emp	Head Emp	NP Support	Other Emp	Fix Asset
$\Delta \ln E_f^S$	0.139	-0.041	-0.162	0.501*	-0.313	-0.229***
	(0.124)	(0.071)	(0.140)	(0.291)	(0.214)	(0.074)
$\Delta \ln E_f^S \times Z$	-0.134	-0.122	-0.277	0.460	0.077	-0.051
	(0.230)	(0.155)	(0.442)	(0.628)	(0.375)	(0.162)
N	1,776	1,776	1,776	1,104	1,425	1,776

Notes: Long-difference (2020–2022) parent-level regressions of the indicated outcome on the revenue-side shock $\Delta \ln E_f^S$ and its interaction with the parent's standardized 2020 leverage Z, total liabilities over total assets, with mean 0 and s.d. 1. β_S is the shock effect at the mean of Z and $\beta_{S\times Z}$ the differential per one-standard-deviation increase in Z. The main effect of Z is included throughout the specifications. All specifications include the baseline parent controls (purchase-side shock, affiliate revenue and purchase coverage, export/import shares, R&D, revenue, employment) and industry fixed effects; standard errors clustered by industry are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

except profit, where it is negative, so the baseline estimates do not reflect cheaper Japanese inputs.

Table A.8: Affiliate Response and Intra-firm Import Intensity from Japan

	(1) Value Added	(2) Profit	(3) Total Pay	(4) Ret. Earn	(5) Emp	(6) Disp. Emp	(7) Cap. Invest.
$\Delta \ln E_a^S$	0.491**	0.760***	0.632***	0.847***	0.145**	0.185**	0.242
	(0.219)	(0.241)	(0.165)	(0.298)	(0.058)	(0.091)	(0.319)
$\Delta \ln E_a^S \times$ Japan-input share	-0.181	-0.324***	-0.076	-0.151	-0.068	0.029	-0.379
	(0.141)	(0.121)	(0.201)	(0.213)	(0.046)	(0.047)	(0.256)
N	8,384	5,601	4,194	4,096	9,240	6,379	4,642

Notes: Affiliate-level long-difference (2020–2022) regressions of the indicated outcome on the revenue-side shock $\Delta \ln E_a^S$ and its interaction with the affiliate's Japan-input share, purchases from Japan over total purchases averaged over 2016–2020 and standardized to mean 0 and s.d. 1. All specifications include the purchase-side shock and host-country×sector fixed effects. Standard errors two-way clustered by host country and parent firm are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.4. Robustness

At the affiliate level we vary the fixed-effect structure. Table A.9 removes the fixed effects entirely, and Table A.10 replaces the host-country-by-sector fixed effects with host-country fixed effects, which absorb host-country-level shocks such as local pandemic restrictions while retaining cross-sector variation within countries. The baseline pattern is stable in both, as value added, profit, payments to the parent, and employment respond to the revenue-side shock, and royalties and rental costs remain null.

Varying the fixed effects also changes where the identifying variation comes from, and the two sources face different confounders. Under host-country-by-sector fixed effects, the bilateral exchange-rate change is common within a host country, so exposure varies only through pre-period transaction geography, and identification requires those shares to be unrelated to other 2020–22 shocks. Without fixed effects, the currency movements themselves contribute variation, and identification instead requires host-country conditions to be unrelated to how far each currency moved. The estimates are similar across the two, so unobserved characteristics common to affiliates in the same host country and sector are not what generated them, including the pandemic-era demand and supply disruptions that dominate 2020–22.

At the parent level, the estimation sample is concentrated in manufacturing, and within manufacturing in metal, machinery, and electronics, as Table A.11 shows. Restricting to manufacturing parents (Table A.12) gives the same picture as the pooled baseline, somewhat more strongly: wages rise, total employment holds steady, non-production support employment grows while other employment shrinks, and fixed assets decline.

Removing the industry fixed effects (Table A.13) changes little: wages rise, non-production support employment increases, fixed assets fall, and total and head-office employment stay flat. The decline in other employment is similar in size.

Table A.9: Affiliate Outcomes: No Fixed Effects

Panel A: Profitability and Financial Flows	Value Added	Profit	Total Pay	Royalty	Ret. Earn
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln E_a^S$	0.465**	0.554**	0.816***	0.229	0.515**
	(0.230)	(0.237)	(0.197)	(0.188)	(0.235)
$\Delta \ln E_a^P$	-0.404**	-0.570**	-0.392*	0.023	-0.299
	(0.166)	(0.228)	(0.220)	(0.207)	(0.223)
Sector×Country FE					
N	8613	5787	4367	2509	4261
R-squared	0.005	0.004	0.005	0.001	0.002

Panel B: Real Activity	Emp	Disp. Emp	Rent Cost	R&D Cost	Cap. Invest.
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln E_a^S$	0.155*	0.183**	-0.004	0.150	0.399
	(0.083)	(0.081)	(0.127)	(0.277)	(0.355)
$\Delta \ln E_a^P$	-0.220***	-0.202***	-0.160	0.088	-0.213
	(0.073)	(0.064)	(0.125)	(0.377)	(0.288)
Sector×Country FE					
N	9466	6556	6497	933	4830
R-squared	0.004	0.004	0.000	0.000	0.001

Notes: Long-difference estimates of the baseline affiliate specification with the host-country-by-sector fixed effects removed. Outcomes and variable definitions are as in Table 1. Standard errors are two-way clustered by host country and parent firm. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Table A.14 measures R&D expenditure, the total wage bill, and total employment in 2018 rather than 2020, matching the event study of Figure 4. The baseline empirical pattern remains unchanged: support employment rises, other employment falls, fixed assets decline, and total and head-office employment do not respond.

Constructing the exposure. The baseline exposure averages the affiliate transaction shares over 2016–2020 and assigns third-country transactions with European destinations to the euro and the rest to the U.S. dollar. Tables A.15 and A.16 vary these two choices. Shifting the share window to 2015–2019 leaves the parent and affiliate results strong and, if anything, strengthens the support and wage responses. Assigning all third-country transactions to the U.S. dollar, or instead assigning third-country Europe to the euro and third-country Asia to the Renminbi, moves the estimates little, because third-country sales are a minority of affiliate revenue and the host-market and Japan components are unaffected. The last row of Table A.15 instead fixes the parent-network weights at their 2016–2020 average rather than their 2020 values, and the

Table A.10: Affiliate Outcomes: Host-Country Fixed Effects Only

Panel A: Profitability and Financial Flows	Value Added	Profit	Total Pay	Royalty	Ret. Earn
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln E_a^S$	0.699**	0.767***	0.766***	0.198	0.869***
	(0.282)	(0.241)	(0.185)	(0.304)	(0.231)
$\Delta \ln E_a^P$	-0.281**	-0.310*	-0.446**	-0.018	0.036
	(0.125)	(0.169)	(0.194)	(0.184)	(0.174)
Country FE	✓	✓	✓	✓	✓
Sector×Country FE					
N	8597	5773	4347	2494	4247
R-squared	0.041	0.032	0.035	0.055	0.033

Panel B: Real Activity	Emp	Disp. Emp	Rent Cost	R&D Cost	Cap. Invest.
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln E_a^S$	0.176***	0.209**	-0.056	0.411	0.189
	(0.059)	(0.082)	(0.128)	(0.304)	(0.199)
$\Delta \ln E_a^P$	-0.208***	-0.196***	-0.191	-0.232	-0.439
	(0.056)	(0.053)	(0.121)	(0.402)	(0.326)
Country FE	✓	✓	✓	✓	✓
Sector×Country FE					
N	9448	6542	6482	926	4820
R-squared	0.023	0.019	0.016	0.031	0.022

Notes: Long-difference estimates of the baseline affiliate specification with host-country fixed effects in place of the host-country-by-sector fixed effects. Outcomes and variable definitions are as in Table 1. Standard errors are two-way clustered by host country and parent firm. *, **, *** denote significance at the 10%, 5%, and 1% levels.

parent estimates are essentially unchanged, consistent with the near-perfect persistence of the within-parent weights documented above. The affiliate operating, payment, and headcount responses are significant under every variant, and the parent wage, support-employment, and fixed-asset responses are stable throughout.

Table A.11: Sector Composition of the Parent-level Sample

Sector	Firms	Share (%)
Manufacturing		
Food, Textile, Chemical, Oil, Plastics	464	25.6
Metal, Machinery, Electronics	770	42.5
Transport Equipment, Precision, Other Mfg.	270	14.9
Manufacturing, total	**1,504**	**83.0**
Non-manufacturing		
Agri/Mining/Construction	21	1.2
Wholesale and Retail	212	11.7
Finance and Real Estate	7	0.4
Utilities/Info/Services/Other	67	3.7
Non-manufacturing, total	**307**	**17.0**
All sectors	**1,811**	**100.0**

Notes: Sector composition of the parent-level estimation sample. Manufacturing follows JSIC Division E (codes 091–329). Shares are computed over the set of firms entering the parent-level first-difference specification.

Table A.12: Parent-level Outcomes: Manufacturing Firms Only

	Avg Wage	Total Emp	Head Emp	NP Support Emp	Other Emp	Fix Asset
$\Delta \ln E_f^S$	0.175*	0.002	-0.241	0.451**	-0.349**	-0.173**
	(0.115)	(0.061)	(0.201)	(0.240)	(0.186)	(0.065)
	[0.098]	[0.040]	[0.157]	[0.176]	[0.155]	[0.069]
$\Delta \ln E_f^P$	-0.017	-0.079	-0.050	0.106	-0.021	-0.050
	(0.093)	(0.063)	(0.146)	(0.207)	(0.205)	(0.045)
	[0.081]	[0.049]	[0.118]	[0.226]	[0.161]	[0.062]
p90–p10 effect (E^S)	0.041	0.000	-0.056	0.103	-0.077	-0.040
Industry FE	✓	✓	✓	✓	✓	✓
N	1,504	1,504	1,504	927	1,232	1,504
R^2	0.118	0.052	0.054	0.083	0.051	0.059

Notes: First-difference estimates for the manufacturing subsample (JSIC 091–329). Outcomes are as in Table 2. Standard errors clustered by industry are in parentheses; exposure-robust standard errors are in brackets. The p90–p10 effect scales the $\Delta \ln E^S$ coefficient by the p90–p10 range of the shock. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Table A.13: Parent-level Outcomes: Robustness without Industry Fixed Effects

	Avg Wage	Total Emp	Head Emp	NP Support Emp	Other Emp	Fix Asset
$\Delta \ln E_f^S$	0.146*	-0.011	-0.129	0.554***	-0.270	-0.197***
	(0.090)	(0.054)	(0.174)	(0.194)	(0.198)	(0.063)
	[0.080]	[0.038]	[0.127]	[0.152]	[0.173]	[0.063]
$\Delta \ln E_f^P$	-0.022	-0.025	-0.058	0.075	0.059	-0.036
	(0.083)	(0.059)	(0.114)	(0.185)	(0.194)	(0.050)
	[0.070]	[0.047]	[0.100]	[0.198]	[0.164]	[0.059]
Industry FE						
N	1,834	1,834	1,834	1,144	1,466	1,834
R^2	0.082	0.016	0.013	0.014	0.007	0.016

Notes: First-difference estimates of the baseline parent-level specification with the industry fixed effects removed; controls are identical to the baseline. Outcomes are as in Table 2. Standard errors clustered by industry are in parentheses; exposure-robust standard errors are in brackets. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Table A.14: Parent-level Outcomes: Controls Measured in 2018

	Avg Wage	Total Emp	Head Emp	NP Support Emp	Other Emp	Fix Asset
$\Delta \ln E_f^S$	0.153*	-0.012	-0.133	0.347**	-0.394**	-0.150**
	(0.108)	(0.052)	(0.183)	(0.188)	(0.227)	(0.060)
	[0.084]	[0.039]	[0.129]	[0.147]	[0.169]	[0.066]
$\Delta \ln E_f^P$	-0.030	-0.036	-0.000	0.165	-0.141	-0.074
	(0.112)	(0.055)	(0.121)	(0.190)	(0.231)	(0.051)
	[0.078]	[0.051]	[0.104]	[0.197]	[0.179]	[0.063]
p90–p10 effect (E^S)	0.037	-0.003	-0.033	0.083	-0.091	-0.037
Industry FE	✓	✓	✓	✓	✓	✓
N	1,726	1,726	1,726	1,086	1,373	1,726
R^2	0.046	0.068	0.055	0.085	0.076	0.067

Notes: First-difference estimates of the baseline parent-level specification, including industry fixed effects, with R&D expenditure, the total wage bill, and total parent employment measured in 2018 instead of 2020. The affiliate coverage and trade-share controls remain measured in 2020, as in the baseline. Outcomes and sample construction are as in Table 2. Standard errors clustered by industry are in parentheses; exposure-robust standard errors are in brackets. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Table A.15: Parent-level Outcomes: Robustness to Exposure Construction

	Ave Wage	Total Emp	Head Emp	NP Support	Other	Fix Asset
Baseline (2016–2020)	0.169**	-0.014	-0.160	0.384**	-0.297*	-0.201***
	[0.080]	[0.038]	[0.129]	[0.152]	[0.173]	[0.064]
2015–2019 window	0.215***	-0.024	-0.163	0.478***	-0.283	-0.190***
	[0.081]	[0.040]	[0.134]	[0.151]	[0.175]	[0.065]
Third-country → USD	0.189**	-0.047	-0.183	0.362**	-0.283	-0.223***
	[0.078]	[0.038]	[0.126]	[0.148]	[0.174]	[0.063]
Third-country EU→EUR, Asia→CNY	0.167**	-0.013	-0.149	0.412***	-0.321*	-0.196***
	[0.082]	[0.040]	[0.133]	[0.160]	[0.171]	[0.064]
2016–2020 average weights	0.166**	-0.015	-0.146	0.368***	-0.287*	-0.140**
	[0.071]	[0.035]	[0.110]	[0.141]	[0.147]	[0.056]

Notes: Long-difference (2020–2022) parent-level estimates under alternative constructions of the exposure measure; outcomes run across the columns and specifications down the rows. The rows vary the pre-period share window (baseline 2016–2020 versus 2015–2019), the third-country currency rule (baseline: Europe to the euro, rest to the U.S. dollar; "Third-country → USD": all third-country to the U.S. dollar; "Third-country EU→EUR, Asia→CNY": Europe to the euro and Asia to the renminbi), and, in the last row, the parent-network weights (baseline: fixed at 2020; "2016–2020 average weights": averaged over the pre-period window). Host-market and Japan transactions are assigned to the host and yen currencies throughout. All specifications include the baseline controls and industry fixed effects. Entries are the revenue-side coefficient with shift-share robust standard errors in brackets, as in Table 2. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Table A.16: Affiliate Outcomes: Robustness to Exposure Construction

	Value Added	Profit	Total Pay	Ret. Earn	Emp	Disp. Emp	Cap. Invest.
Baseline (2016–2020)	0.497**	0.754***	0.622***	0.832***	0.149***	0.184**	0.211
	(0.213)	(0.230)	(0.160)	(0.289)	(0.057)	(0.088)	(0.329)
2015–2019 window	0.501**	0.828***	0.496***	0.926***	0.102**	0.174**	0.256
	(0.254)	(0.254)	(0.156)	(0.351)	(0.050)	(0.083)	(0.220)
Third-country → USD	0.470**	0.686***	0.647***	0.819***	0.132**	0.177**	0.384
	(0.202)	(0.219)	(0.165)	(0.303)	(0.057)	(0.086)	(0.298)
Third-country EU→EUR, Asia→CNY	0.569**	0.754***	0.643***	0.763**	0.178***	0.216**	0.317
	(0.264)	(0.276)	(0.164)	(0.348)	(0.060)	(0.092)	(0.355)

Notes: Long-difference (2020–2022) affiliate-level estimates under alternative constructions of the exposure measure; outcomes run across the columns and specifications down the rows, defined as in Table A.15 (the parent-network weight row does not apply at the affiliate level). All specifications include host-country×sector fixed effects. Entries are the revenue-side coefficient with standard errors two-way clustered by host country and parent firm in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% levels.

B Cross-Country Evidence

This appendix describes the data and specification behind Figure 8 and reports robustness.

B.1. Data, Exposure, and Specification

Bilateral outward FDI positions are from the IMF Coordinated Direct Investment Survey, which reports positions in U.S. dollars and is available from 2009. Exchange rates are period-average rates against the U.S. dollar from the IMF International Financial Statistics, from which any bilateral rate follows as a cross rate; a rise is a depreciation. Because the exchange-rate file does not carry euro-area members as individual reporters, they are not in the sample. Direct investment income is from the IMF Balance of Payments Statistics (BPM6), real GDP growth is from the World Economic Outlook, and nominal GDP from the same source scales the income outcome. Conduit and special-purpose-entity hubs (Bermuda, the British Virgin Islands, the Cayman Islands, Curaçao, Cyprus, Guernsey, Hong Kong SAR, Ireland, the Isle of Man, Jersey, Liberia, Luxembourg, Malta, Mauritius, the Netherlands, Panama, and Singapore) and three small offshore financial centers of the same character (Barbados, Macao SAR, and Samoa) are excluded because their reported positions reflect where investment is booked rather than where affiliate activity takes place. Hyperinflation episodes, defined as annual depreciation above roughly 120 log points, are dropped, and outcomes are winsorized at the 2.5 and 97.5 percentiles. The estimation sample covers 70 economies with shock years 2014–2024.

For horizons $h = 0, 1, 2$, we estimate

$$\Delta_h y_{c,t} = \alpha_c + \lambda_t + \beta_h X_{c,t}^{FDI} + \gamma_h X_{c,t}^{Trade} + \delta_h Z_{c,t} + \rho_h y_{c,t-1} + \varepsilon_{c,t+h},$$

with country and year fixed effects and standard errors clustered by country. The exposure is the FDI-weighted network depreciation, $X_{c,t}^{FDI} = \sum_d (\mathrm{FDI}_{cd,0}/\mathrm{GDP}_{c,0})\, \Delta s_{cdt}$, where Δs_{cdt} is the bilateral depreciation of country c against destination d, measured in percent, and $\mathrm{FDI}_{cd,0}$ is the predetermined outward position, averaged together with GDP over $[t-5,\ t-3]$. Because bilateral positions exist from 2009, 2014 is the earliest shock year with a full window. Because the weights sum to the outward-FDI position relative to GDP rather than to one, $Z_{c,t}$ controls for that sum, following the treatment of incomplete shares in Borusyak et al. (2022). The control $X_{c,t}^{Trade}$ is the trade-weighted effective depreciation, constructed by applying pre-period bilateral goods-trade shares to the same bilateral exchange-rate movements because the published IMF effective exchange rate covers only about half of the estimation sample; it enters nominal for the income outcome and CPI-deflated for real GDP. For the income outcome, the dependent variable is the change in the ratio of direct investment income credits to GDP from $t-1$ to $t+h$, in percentage points. Because numerator and denominator are converted at the same year's exchange rate, this ratio equals the local-currency ratio and is invariant to the reporting currency, so it retains

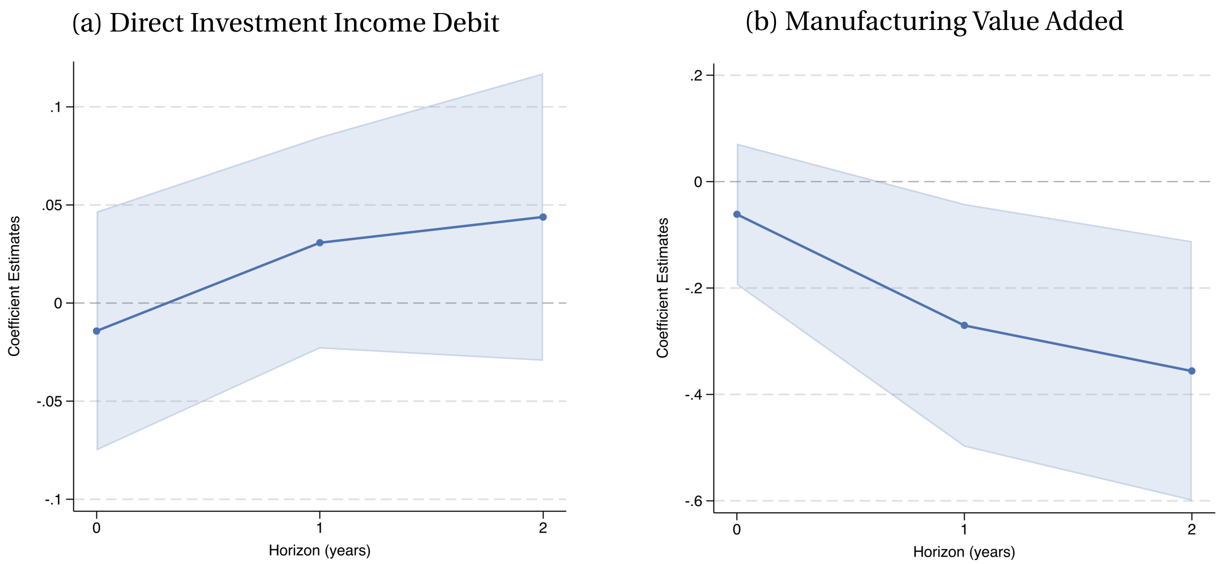

Figure B.1: Cross-Country Panel Local Projections: Additional Outcomes

Notes: This figure reports the coefficients β_h on the FDI-network depreciation exposure from the same panel local projections as Figure 8, with 90% confidence intervals based on standard errors clustered by country. In panel (a), the dependent variable is the change in the ratio of direct investment income debits to GDP from $t-1$ to $t+h$, in percentage points, with the nominal trade-weighted control; debit is income paid to nonresident direct investors. In panel (b), the dependent variable is cumulative manufacturing value added growth (World Bank WDI, constant prices) from $t-1$ to $t+h$, in percent, with the CPI-deflated trade-weighted control. All specifications include the sum of the FDI weights, the lagged dependent variable, and country and year fixed effects.

the home-currency valuation of foreign-currency income. For real GDP, the dependent variable is cumulative growth from $t-1$ to $t+h$, in percent.

A unit increase in exposure corresponds to one percentage point of GDP-share-weighted network depreciation. The standard deviation of exposure is 1.6 percentage points in the credit sample and 1.9 percentage points in the real GDP sample. Using these outcome-specific standard deviations, the estimates imply direct investment income credits 0.13 percentage points of GDP higher and cumulative real GDP growth 0.8 percentage points lower at the two-year horizon.

B.2. Additional Outcomes and Robustness

Figure B.1 reports the same projection for two additional outcomes. Panel (a) shows the gross debit, direct investment income paid to nonresident direct investors, with the same currency-invariant ratio construction as the credit. The response referenced in the main text is small and imprecisely estimated at every horizon; payments to inward direct investors are home-currency profits with no translation component, so a flat debit is what the mechanism implies. Panel (b) shows cumulative manufacturing value added growth, the tradable production the mechanism reallocates across borders; it falls at one to two years, in line with the real GDP response in Figure 8.

Table B.1 reports robustness of the two responses plotted in Figure 8. Controlling for one lag of both shocks leaves the responses intact. Converting income and GDP to local currency with the bilateral U.S. dollar rate before differencing, the other currency-invariant measurement, gives the same or stronger credit response. Both responses are as strong within shock years 2017–2024. The remaining rows adapt the local-projection difference-in-differences ingredients of Dube et al. (2023) to a continuous treatment. Controlling for the shocks dated between $t+1$ and $t+h$ isolates the response to a one-off shock and is close to the baseline. Additionally restricting the sample to observations whose two lagged absolute exposures lie below the sample median, their clean-control condition adapted to a continuous treatment, cuts the panel to roughly a third; the responses keep their sign, with magnitudes no longer well disciplined. Dropping the country fixed effects, which the long-differenced outcome makes largely redundant, gives smaller responses that remain significant at one to two years.

Table B.2 instead varies how the exposure itself is built, addressing three concerns: that it is driven entirely by the dollar, that it mixes the geography of a country's outward-FDI network with the size of that network, and that it reflects demand conditions in the destination markets. Dropping the United States as a destination, and then also dropping every currency that is pegged to the dollar, leaves both responses intact and, if anything, somewhat larger, so neither depends on the dollar leg of the exposure. The next row separates geography from size. Because the weights are bilateral positions over GDP, the exposure splits exactly into two additive terms: the outward-FDI-to-GDP ratio times the country's own depreciation against the dollar, and the position-weighted depreciation of the destination currencies. The first is large whenever a country with a big outward-FDI sector sees its currency fall against everything, which is the combination that could proxy for global episodes rather than for the multinational channel. Controlling for it holds both the size of the outward-FDI sector and the country's own depreciation fixed, so the exposure is identified from the destination-currency term alone, that is, from which destinations the country invested in. Both responses survive and are larger than in the baseline. Finally, controlling for the position-weighted real GDP growth of the destination countries leaves both responses essentially unchanged.

Table B.1: Cross-Country Robustness: Credit and Real GDP Responses

	$h=0$	$h=1$	$h=2$	$N_{h=0}$
Panel A: Direct investment income, gross credit				
Baseline	0.046	0.082***	0.080**	612
	(0.029)	(0.019)	(0.034)	
Lagged shocks controlled	0.048**	0.086***	0.084***	587
	(0.024)	(0.017)	(0.032)	
Local-currency outcome	0.053*	0.087***	0.103**	612
	(0.031)	(0.027)	(0.048)	
Shock years 2017–2024	0.069**	0.081***	0.091***	469
	(0.027)	(0.018)	(0.035)	
One-off: future shocks controlled	0.046	0.089***	0.080**	612
	(0.029)	(0.015)	(0.035)	
One-off, quasi-stayer sample	0.109	0.083	0.514***	237
	(0.084)	(0.087)	(0.138)	
Year fixed effects only	0.015	0.050***	0.052**	612
	(0.018)	(0.017)	(0.025)	
Panel B: Real GDP growth				
Baseline	-0.101**	-0.238***	-0.422***	632
	(0.045)	(0.073)	(0.080)	
Lagged shocks controlled	-0.099**	-0.237***	-0.391***	605
	(0.043)	(0.056)	(0.071)	
Shock years 2017–2024	-0.088**	-0.204***	-0.366***	478
	(0.040)	(0.073)	(0.070)	
One-off: future shocks controlled	-0.101**	-0.203***	-0.391***	632
	(0.045)	(0.067)	(0.064)	
One-off, quasi-stayer sample	-1.544*	-2.416	-3.556*	241
	(0.922)	(1.692)	(2.068)	
Year fixed effects only	-0.076	-0.175*	-0.276***	632
	(0.048)	(0.098)	(0.106)	

Notes: Each row re-estimates the baseline panel local projection for the gross-credit outcome (Panel A) or the real-GDP outcome (Panel B) of Figure 8 with the stated modification; entries are β_h with country-clustered standard errors in parentheses, and N at $h=0$ (N declines with the horizon as in the baseline). "Lagged shocks controlled" adds one lag of the FDI-network exposure and of the trade-weighted depreciation. "Local-currency outcome" converts income and GDP to local currency with the bilateral U.S. dollar rate before differencing and applies to the income outcome only. "One-off" controls for the shocks between $t+1$ and $t+h$, and the quasi-stayer sample additionally restricts to observations whose two lagged absolute exposures lie below the sample median; together these adapt the local-projection difference-in-differences logic of Dube et al. (2023). * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table B.2: Cross-Country Robustness: Constructing the Exposure

	$h=0$	$h=1$	$h=2$	$N_{h=0}$
Panel A: Direct investment income, gross credit				
Baseline	0.046	0.082***	0.080**	612
	(0.029)	(0.019)	(0.034)	
Leave-U.S. exposure	0.067**	0.115***	0.101**	612
	(0.033)	(0.024)	(0.043)	
Leave-dollar exposure	0.075**	0.130***	0.104**	609
	(0.032)	(0.025)	(0.044)	
Scale × own depreciation controlled	0.121***	0.168***	0.138***	612
	(0.028)	(0.031)	(0.047)	
Destination-weighted growth controlled	0.046	0.082***	0.080**	612
	(0.029)	(0.019)	(0.034)	
Panel B: Real GDP growth				
Baseline	-0.101**	-0.238***	-0.422***	632
	(0.045)	(0.073)	(0.080)	
Leave-U.S. exposure	-0.133***	-0.290***	-0.509***	632
	(0.049)	(0.092)	(0.113)	
Leave-dollar exposure	-0.141***	-0.310***	-0.519***	629
	(0.051)	(0.097)	(0.117)	
Scale × own depreciation controlled	-0.168***	-0.291***	-0.455***	632
	(0.055)	(0.102)	(0.110)	
Destination-weighted growth controlled	-0.097**	-0.237***	-0.426***	632
	(0.044)	(0.071)	(0.077)	

Notes: Each row re-estimates the baseline panel local projection of Table B.1 with the exposure rebuilt or with an additional control, for the gross-credit outcome (Panel A) and the real-GDP outcome (Panel B); entries are β_h with country-clustered standard errors in parentheses. "Leave-U.S. exposure" rebuilds the exposure after dropping the United States as a destination, and "Leave-dollar exposure" additionally drops every destination whose currency is de facto pegged to the dollar, defined as a standard deviation of its annual depreciation against the dollar below 0.01 over the sample. "Scale × own depreciation controlled" adds the outward-FDI-to-GDP ratio interacted with the country's own depreciation against the dollar, which is the additive piece of the exposure that reflects the size of the outward-FDI sector rather than its geography, so that the exposure coefficient comes from the destination-currency piece alone. "Destination-weighted growth controlled" adds the position-weighted real GDP growth of the destination countries. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

C Proofs and Regularity Conditions for Section 4

C.1. Proof of Proposition 4.1

We prove the comparative static in Proposition 4.1. The parent solves

$$\max_{S_d, S_a} V(S_d) + \mathscr{E}\xi V^*(S_a) - C(S_d + S_a),$$

with $N = S_d + S_a$. The first-order conditions are

$$V'(S_d) = C'(N),$$
$$\mathscr{E}\xi V^{*\prime}(S_a) = C'(N).$$

Totally differentiating the first condition gives

$$V''(S_d)dS_d = C''(N)dN.$$

Using $dN = dS_d + dS_a$, this becomes

$$V'' dS_d = C''(dS_d + dS_a).$$

Rearranging,

$$dS_d = -\frac{C''}{-V'' + C''} dS_a. \tag{C.1}$$

Because $V'' < 0$ and $C'' > 0$, Equation (C.1) implies that $dS_d < 0$ whenever $dS_a > 0$. Moreover,

$$dN = dS_d + dS_a = \frac{-V''}{-V'' + C''} dS_a. \tag{C.2}$$

Thus $dN > 0$ whenever $dS_a > 0$.

Now totally differentiate the second first-order condition:

$$\xi V^{*\prime} d\mathscr{E} + \mathscr{E}\xi V^{*\prime\prime} dS_a = C'' dN.$$

Using Equation (C.2), we obtain

$$\xi V^{*\prime} d\mathscr{E} = \left[-\mathscr{E}\xi V^{*\prime\prime} + C'' \frac{-V''}{-V'' + C''} \right] dS_a.$$

The term in brackets is strictly positive because $V^{*\prime\prime} < 0$, $V'' < 0$, $C'' > 0$, $\mathscr{E} > 0$, and $\xi > 0$. Since $\xi V^{*\prime} > 0$, an increase in $\mathscr{E}$ implies

$$dS_a > 0.$$

Equation (C.1) and Equation (C.2) then imply

$$dS_d < 0, \qquad dN > 0.$$

Finally, because $W_N'(N) > 0$, the increase in support employment raises the support wage:

$$dW_N = W_N'(N)dN > 0.$$

■

C.2. Production-Labor Regularity Conditions

The main text assumes that optimized production-labor demand is increasing in the support services allocated to that operation:

$$\ell'(S_d) > 0, \qquad \ell^{*\prime}(S_a) > 0.$$

The following proposition gives primitive conditions under which this property holds.

Proposition C.1 (Regularity conditions for production-labor responses)**.** *Suppose domestic production labor and affiliate production labor are chosen optimally from the value functions*

$$V(S_d) = \max_{L} \{R(F(S_d, L)) - W_L L\},$$
$$V^*(S_a) = \max_{L^*} \left\{R^*\left(F(S_a, L^*)\right) - W_L^* L^*\right\}.$$

Assume the domestic and affiliate production-labor choices are interior. Suppose the domestic labor choice satisfies the second-order condition

$$R''(F(S_d, L))\left(\frac{\partial F(S_d, L)}{\partial L}\right)^2 + R'(F(S_d, L))\frac{\partial^2 F(S_d, L)}{\partial L^2} < 0, \tag{C.3}$$

and suppose support services raise the marginal profitability of domestic production labor:

$$R''(F(S_d, L))\frac{\partial F(S_d, L)}{\partial S_d}\frac{\partial F(S_d, L)}{\partial L} + R'(F(S_d, L))\frac{\partial^2 F(S_d, L)}{\partial S_d \partial L} > 0. \tag{C.4}$$

Similarly, suppose the affiliate labor choice satisfies the second-order condition

$$R^{*\prime\prime}\left(F(S_a, L^*)\right)\left(\frac{\partial F(S_a, L^*)}{\partial L^*}\right)^2 + R^{*\prime}\left(F(S_a, L^*)\right)\frac{\partial^2 F(S_a, L^*)}{\partial (L^*)^2} < 0, \tag{C.5}$$

and suppose support services raise the marginal profitability of affiliate production labor:

$$R^{*\prime\prime}\left(F(S_a, L^*)\right)\frac{\partial F(S_a, L^*)}{\partial S_a}\frac{\partial F(S_a, L^*)}{\partial L^*} + R^{*\prime}\left(F(S_a, L^*)\right)\frac{\partial^2 F(S_a, L^*)}{\partial S_a \partial L^*} > 0. \tag{C.6}$$

Then optimized production-labor demand is increasing in allocated support services:

$$\ell'(S_d) > 0, \qquad \ell^{*\prime}(S_a) > 0.$$

Consequently, together with Proposition 4.1, a small depreciation raises affiliate production labor and reduces domestic production labor:

$$dL^* > 0, \qquad dL < 0.$$

Since production is increasing in both inputs, affiliate output rises and domestic output falls:

$$dY^* > 0, \qquad dY < 0.$$

Proof. For domestic operations, production labor is chosen from

$$V(S_d) = \max_L \{R(F(S_d, L)) - W_L L\}.$$

The domestic labor first-order condition is

$$R'(F(S_d, L))\,\frac{\partial F(S_d, L)}{\partial L} = W_L.$$

Define

$$\Phi(S_d, L) \equiv R'(F(S_d, L))\,\frac{\partial F(S_d, L)}{\partial L} - W_L.$$

At an interior optimum, $\Phi(S_d, L) = 0$. The derivative of the first-order condition with respect to L is

$$\frac{\partial \Phi(S_d, L)}{\partial L} = R''(F(S_d, L))\left(\frac{\partial F(S_d, L)}{\partial L}\right)^2 + R'(F(S_d, L))\,\frac{\partial^2 F(S_d, L)}{\partial L^2}.$$

By Equation (C.3), $\partial\Phi/\partial L < 0$. Hence the implicit function theorem implies that the first-order condition locally defines $L = \ell(S_d)$, with

$$\frac{dL}{dS_d} = -\frac{\partial\Phi/\partial S_d}{\partial\Phi/\partial L}.$$

Moreover,

$$\frac{\partial \Phi(S_d, L)}{\partial S_d} = R''(F(S_d, L))\,\frac{\partial F(S_d, L)}{\partial S_d}\,\frac{\partial F(S_d, L)}{\partial L} + R'(F(S_d, L))\,\frac{\partial^2 F(S_d, L)}{\partial S_d \partial L}.$$

By Equation (C.4), $\partial\Phi/\partial S_d > 0$. Therefore

$$\ell'(S_d) = \frac{dL}{dS_d} > 0.$$

The affiliate problem is analogous. Given affiliate support S_a, affiliate production labor is chosen from

$$V^*(S_a) = \max_{L^*} \left\{R^*\big(F(S_a, L^*)\big) - W_L^* L^*\right\}.$$

The affiliate labor first-order condition is

$$R^{*\prime}\big(F(S_a, L^*)\big)\,\frac{\partial F(S_a, L^*)}{\partial L^*} = W_L^*.$$

Define

$$\Phi^*(S_a, L^*) \equiv R^{*\prime}\big(F(S_a, L^*)\big)\,\frac{\partial F(S_a, L^*)}{\partial L^*} - W_L^*.$$

By Equation (C.5), $\partial\Phi^*/\partial L^* < 0$. Hence the implicit function theorem implies that the first-order condition locally defines $L^* = \ell^*(S_a)$, with

$$\frac{dL^*}{dS_a} = -\frac{\partial\Phi^*/\partial S_a}{\partial\Phi^*/\partial L^*}.$$

By Equation (C.6), $\partial\Phi^*/\partial S_a > 0$. Therefore

$$\ell^{*\prime}(S_a) = \frac{dL^*}{dS_a} > 0.$$

Proposition 4.1 implies that a depreciation raises affiliate support and reduces domestic support:

$$dS_a > 0, \qquad dS_d < 0.$$

Since $\ell'(S_d) > 0$ and $\ell^{*\prime}(S_a) > 0$, it follows that

$$dL = \ell'(S_d)dS_d < 0, \qquad dL^* = \ell^{*\prime}(S_a)dS_a > 0.$$

Since F is increasing in both inputs, $dS_a > 0$ and $dL^* > 0$ imply $dY^* > 0$, while $dS_d < 0$ and $dL < 0$ imply $dY < 0$. ■

C.3. Proof of Corollary 4.1

Suppose

$$R(Y) = AY^\rho, \qquad R^*(Y^*) = A^*(Y^*)^\rho, \qquad 0 < \rho < 1,$$

and

$$F(S, L) = S^{\alpha_S} L^{\alpha_L}, \qquad \alpha_S > 0, \quad \alpha_L > 0,$$

with

$$\rho(\alpha_S + \alpha_L) < 1.$$

For domestic operations, the problem is

$$V(S_d) = \max_L \left\{ A S_d^{\rho\alpha_S} L^{\rho\alpha_L} - W_L L \right\}.$$

The first-order condition is

$$\rho\alpha_L A S_d^{\rho\alpha_S} L^{\rho\alpha_L - 1} = W_L.$$

Solving for domestic production labor gives

$$L = \left(\frac{\rho\alpha_L A}{W_L} \right)^{\frac{1}{1-\rho\alpha_L}} S_d^{\frac{\rho\alpha_S}{1-\rho\alpha_L}}.$$

Since $\alpha_S > 0$, $\rho > 0$, and $\rho(\alpha_S + \alpha_L) < 1$, we have $1 - \rho\alpha_L > 0$. Hence

$$\frac{dL}{dS_d} > 0.$$

Substituting optimized labor demand back into the domestic value function implies

$$V(S_d) = \kappa S_d^{\frac{\rho\alpha_S}{1-\rho\alpha_L}},$$

for some constant $\kappa > 0$. The exponent satisfies

$$0 < \frac{\rho\alpha_S}{1-\rho\alpha_L} < 1,$$

because $\rho(\alpha_S + \alpha_L) < 1$. Hence

$$V'(S_d) > 0, \qquad V''(S_d) < 0.$$

For the affiliate, the problem is

$$V^*(S_a) = \max_{L^*} \left\{ A^* S_a^{\rho\alpha_S} (L^*)^{\rho\alpha_L} - W_L^* L^* \right\}.$$

The first-order condition gives

$$L^* = \left(\frac{\rho\alpha_L A^*}{W_L^*} \right)^{\frac{1}{1-\rho\alpha_L}} S_a^{\frac{\rho\alpha_S}{1-\rho\alpha_L}}.$$

Thus

$$\frac{dL^*}{dS_a} > 0.$$

Substituting optimized affiliate labor demand back into the affiliate value function gives

$$V^*(S_a) = \kappa^* S_a^{\frac{\rho\alpha_S}{1-\rho\alpha_L}},$$

for some constant $\kappa^* > 0$. Therefore

$$V^{*\prime}(S_a) > 0, \qquad V^{*\prime\prime}(S_a) < 0.$$

The preceding expressions show that optimized production-labor demand is increasing in allocated support services for both domestic and affiliate operations, and that the induced value functions are increasing and strictly concave. Hence the conditions in Proposition 4.1 and Proposition C.1 are satisfied. ■

D Quantitative Model Details

This appendix presents the detailed two-country general-equilibrium model in Section 5. The model features households, tradable producers, multinational parents and their foreign affiliates, nominal rigidities, international asset trade, and monetary policy. We describe the Home economy below; the Foreign economy is symmetric unless otherwise noted. We then define market clearing, external accounts, and the equilibrium conditions that close the model.

D.1. Households

The Home representative household supplies two types of labor: production labor L_t, supplied to both tradable firms and multinational firms, and non-production support labor N_t, supplied to multinational parents. Support labor is differentiated across parent firms: N_t denotes an aggregate of firm-specific support-labor services $N_t(j)$, whose aggregation and allocation across firms are specified below.

The household maximizes expected lifetime utility under Greenwood–Hercowitz–Huffman (GHH) preferences:

$$E_0 \sum_{t=0}^{\infty} \beta^t \frac{X_t^{1-\sigma}}{1-\sigma}, \qquad X_t = C_t - \chi_L \frac{L_t^{1+\varphi}}{1+\varphi} - \chi_N \frac{N_t^{1+\varphi}}{1+\varphi},$$

where C_t denotes final consumption, σ governs the curvature of utility, χ_L and χ_N govern the disutility of production and support labor, and φ is the inverse Frisch elasticity. GHH preferences eliminate wealth effects on labor supply, so aggregate labor responds to wages rather than directly to exchange-rate-induced income gains.

Consumption structure. Final consumption has a nested CES structure. At the top level, households combine a multinational-goods composite $C_{M,t}$ and a tradable-goods composite $C_{T,t}$:

$$C_t = \left[\omega^{\frac{1}{\eta}} C_{M,t}^{\frac{\eta-1}{\eta}} + (1-\omega)^{\frac{1}{\eta}} C_{T,t}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where $\eta > 0$ governs substitution between multinational and tradable goods, and ω determines the expenditure weight on the multinational sector.

Cost minimization implies

$$C_{M,t} = \omega \left(\frac{P_{M,t}}{P_t} \right)^{-\eta} C_t, \qquad C_{T,t} = (1-\omega) \left(\frac{P_{T,t}}{P_t} \right)^{-\eta} C_t,$$

where P_t, $P_{M,t}$, and $P_{T,t}$ denote the price indexes for final consumption, the multinational composite, and the tradable composite, respectively.

Each sectoral composite $C_{\ell,t}$, for $\ell \in \{M, T\}$, aggregates Home and Foreign components:

$$C_{\ell,t} = \left[\omega_\ell^{\frac{1}{\eta_\ell}} C_{\ell H,t}^{\frac{\eta_\ell-1}{\eta_\ell}} + (1-\omega_\ell)^{\frac{1}{\eta_\ell}} C_{\ell F,t}^{\frac{\eta_\ell-1}{\eta_\ell}} \right]^{\frac{\eta_\ell}{\eta_\ell-1}}.$$

For the multinational composite, $C_{MH,t}$ denotes goods produced in Home by Home-controlled multinational parents, while $C_{MF,t}$ denotes goods produced in Home by Foreign-controlled affiliates. For the tradable composite, $C_{TH,t}$ denotes Home-produced tradable goods and $C_{TF,t}$ denotes imported Foreign tradable goods. Thus, multinational goods are produced locally by either domestically controlled parents or foreign-controlled affiliates, whereas tradable goods may be produced domestically or imported.

The associated sectoral price index is

$$P_{\ell,t} = \left[\omega_\ell P_{\ell H,t}^{1-\eta_\ell} + (1-\omega_\ell) P_{\ell F,t}^{1-\eta_\ell} \right]^{\frac{1}{1-\eta_\ell}}, \qquad \ell \in \{M, T\}.$$

Cost minimization implies the demand functions

$$C_{\ell H,t} = \omega_\ell \left(\frac{P_{\ell H,t}}{P_{\ell,t}}\right)^{-\eta_\ell} C_{\ell,t}, \qquad C_{\ell F,t} = (1-\omega_\ell) \left(\frac{P_{\ell F,t}}{P_{\ell,t}}\right)^{-\eta_\ell} C_{\ell,t}, \qquad \ell \in \{M, T\}.$$

Here $P_{MH,t}$ and $P_{MF,t}$ denote the price indexes for Home-controlled multinational goods and Foreign-controlled affiliate goods sold in Home, while $P_{TH,t}$ and $P_{TF,t}$ denote the corresponding price indexes for Home and Foreign tradable goods sold in Home.

At the variety level, monopolistically competitive firms produce differentiated goods. For Home tradables, the CES composite is

$$C_{TH,t} = \left[\int_0^1 C_{TH,t}(j)^{\frac{\varepsilon_T - 1}{\varepsilon_T}} \, dj\right]^{\frac{\varepsilon_T}{\varepsilon_T - 1}}.$$

The corresponding variety-level aggregators for Foreign tradables and multinational goods are defined analogously. We introduce the implied variety-level demand schedules below when characterizing firms' pricing problems.

Labor supply. Production labor is supplied competitively at the common wage $W_{L,t}$. We focus on the structure of the support labor market, since support services are the input through which multinational activity affects domestic production.

The support labor market features monopsonistic competition, which provides a parsimonious way to generate the upward-sloping firm-level support labor supply curve used in Section 4. Each multinational parent firm j posts a firm-specific nominal support wage $W_{N,t}(j)$, and households allocate aggregate support labor across parent firms $j \in [0,1]$ according to

$$N_t = \left[\int_0^1 N_t(j)^{\frac{1+\varepsilon_N}{\varepsilon_N}} \, dj\right]^{\frac{\varepsilon_N}{1+\varepsilon_N}},$$

where $\varepsilon_N > 0$ governs the elasticity of support labor supply to firm-specific wages. Conditional on aggregate support labor N_t, the household allocates support labor across parent firms in response to firm-specific wages:

$$N_t(j) = \left(\frac{W_{N,t}(j)}{W_{N,t}}\right)^{\varepsilon_N} N_t, \tag{D.1}$$

where the aggregate support wage index is

$$W_{N,t} = \left[\int_0^1 W_{N,t}(j)^{1+\varepsilon_N} dj\right]^{\frac{1}{1+\varepsilon_N}}.$$

Thus, each parent firm faces an upward-sloping support labor supply schedule with elasticity ε_N. The support wage index also implies that aggregate support labor income satisfies

$$\int_0^1 W_{N,t}(j) N_t(j) \, dj = W_{N,t} N_t.$$

Budget constraint. The Home household trades one-period nominal bonds denominated in Home and Foreign currency. Home-currency bonds are traded domestically and pay the gross nominal return R_t. Foreign-currency bonds are traded internationally and pay the gross nominal return R_t^*. The household also pays a quadratic portfolio adjustment cost on its real Foreign-currency bond position, which ensures stationarity of net foreign assets (Schmitt-Grohé and Uribe, 2003). An exogenous risk-premium shock perturbs the return on Foreign-currency bonds and generates deviations from uncovered interest parity.

The nominal budget constraint of the representative household is

$$\begin{aligned} &P_t C_t + B_{H,t} + \mathscr{E}_t B_{F,t} + P_t \frac{\psi_b}{2}\left(\frac{\mathscr{E}_t B_{F,t}}{P_t} - \bar{b}_F\right)^2 \\ &\quad = R_{t-1} B_{H,t-1} + R_{t-1}^* \mathscr{E}_t B_{F,t-1} \exp(\psi_{rp,t-1}) + W_{L,t} L_t + W_{N,t} N_t + \Pi_t^H + \mathscr{I}_t^H. \end{aligned} \tag{D.2}$$

Here $B_{H,t}$ and $B_{F,t}$ denote holdings of Home- and Foreign-currency nominal bonds, respectively, and $\mathscr{E}_t$ is the nominal exchange rate expressed as Home currency per unit of Foreign currency. The parameter $\psi_b > 0$ governs the adjustment cost around the steady-state real foreign bond position $\bar{b}_F$, and $\psi_{rp,t}$ is an exogenous risk-premium shock on Foreign-currency bond returns. Labor income consists of production wage income, $W_{L,t}L_t$, and aggregate support wage income, $W_{N,t}N_t$. The term Π_t^H denotes Home-currency profits from Home-located production accruing to Home households. It includes profits of Home tradable producers, domestic operating profits of Home MNE parents net of their support wage bill, and the share $1-\xi^*$ of operating surplus generated by Foreign-controlled affiliates in Home that accrues to Home residents. The term $\mathscr{I}_t^H$ separately denotes Home-currency direct investment income accruing to Home households from Home-controlled affiliates operating abroad.

D.2. Tradable Producers

A continuum of Home tradable producers $j \in [0,1]$ produces differentiated varieties using Home production labor and sells them in both Home and Foreign markets. Technology is linear:

$$Y_{T,t}(j) = L_{T,t}(j). \tag{D.3}$$

Total output satisfies

$$Y_{T,t}(j) = C_{TH,t}(j) + C_{TH,t}^*(j). \tag{D.4}$$

Demand in each destination is

$$C_{TH,t}(j) = \left(\frac{P_{TH,t}(j)}{P_{TH,t}}\right)^{-\varepsilon_T} C_{TH,t}, \tag{D.5}$$

$$C_{TH,t}^*(j) = \left(\frac{P_{TH,t}^*(j)}{P_{TH,t}^*}\right)^{-\varepsilon_T} C_{TH,t}^*, \tag{D.6}$$

where $P_{TH,t}(j)$ is the Home-currency price charged in the Home market, and $P^*_{TH,t}(j)$ is the Foreign-currency price charged in the Foreign market.

Firms are monopolistically competitive and set destination-specific local-currency prices subject to Rotemberg adjustment costs. The nominal profit of Home tradable producer j is

$$\Pi_{T,t}(j) = P_{TH,t}(j)C_{TH,t}(j) + \mathscr{E}_t P^*_{TH,t}(j)C^*_{TH,t}(j) - W_{L,t}L_{T,t}(j) \\ - \frac{\psi^T_p}{2} P_{TH,t}C_{TH,t}\left(\frac{P_{TH,t}(j)}{P_{TH,t-1}(j)} - 1\right)^2 - \frac{\psi^{T,*}_p}{2}\mathscr{E}_t P^*_{TH,t}C^*_{TH,t}\left(\frac{P^*_{TH,t}(j)}{P^*_{TH,t-1}(j)} - 1\right)^2.$$

Each firm chooses production labor and destination-specific local-currency prices to maximize the expected discounted value of profits, subject to the technology in Equation (D.3), the market-clearing condition in Equation (D.4), and the demand schedules in Equation (D.5)–Equation (D.6). Price adjustment is subject to Rotemberg costs, which are treated as real expenditures on the firm's own sectoral composite. These costs therefore enter goods-market clearing as adjustment absorption and reduce distributable profits. Because they are quadratic in price changes, the associated variety-level adjustment demand is second order around the zero-inflation steady state and does not affect the model's first-order dynamics.

Foreign tradable producers are defined symmetrically.

D.3. Multinational Firms

The multinational block implements the support-allocation mechanism developed in Section 4. A Home multinational parent $j \in [0,1]$ operates a domestic production unit and a Foreign affiliate. The parent hires non-production support labor $N_t(j)$ and transforms it one-for-one into support services,

$$S_t(j) = N_t(j),$$

which are allocated between domestic operations and the Foreign affiliate:

$$S_t(j) = S_{d,t}(j) + S_{a,t}(j).$$

Domestic parent production combines domestic support services and production labor,

$$Y_{MH,t}(j) = F\big(S_{d,t}(j), L_{MH,t}(j)\big), \tag{D.7}$$

while the Home-controlled affiliate abroad produces

$$Y^*_{MH,t}(j) = F\big(S_{a,t}(j), L^*_{MH,t}(j)\big),$$

where $L^*_{MH,t}(j)$ denotes production labor hired locally by the affiliate abroad. We use a constant-returns Cobb–Douglas technology,

$$F(S,L) = S^{\alpha_S}L^{1-\alpha_S}, \qquad \alpha_S \in (0,1).$$

Foreign multinationals are defined symmetrically.

The Home parent's domestic unit sells a differentiated multinational variety in the Home market. Demand is

$$C_{MH,t}(j) = \left(\frac{P_{MH,t}(j)}{P_{MH,t}}\right)^{-\varepsilon_M} C_{MH,t}, \tag{D.8}$$

where $P_{MH,t}(j)$ is the Home-currency price of variety j. The Home-controlled affiliate abroad sells a differentiated variety in the Foreign market. Demand is

$$C^*_{MH,t}(j) = \left(\frac{P^*_{MH,t}(j)}{P^*_{MH,t}}\right)^{-\varepsilon_M} C^*_{MH,t},$$

where $P^*_{MH,t}(j)$ is the Foreign-currency price charged by the affiliate.

Both the domestic unit and the Foreign affiliate set local-currency prices subject to Rotemberg adjustment costs. The nominal operating profit of the domestic unit, before support labor costs, is

$$\begin{aligned}\widetilde{\Pi}_{MH,t}(j) = P_{MH,t}(j)C_{MH,t}(j) - W_{L,t}L_{MH,t}(j)\\ - \frac{\psi_p^M}{2} P_{MH,t}C_{MH,t}\left(\frac{P_{MH,t}(j)}{P_{MH,t-1}(j)} - 1\right)^2.\end{aligned}$$

The nominal operating profit of the Foreign affiliate, before payments to the parent, is

$$\begin{aligned}\widetilde{\Pi}^*_{MH,t}(j) = P^*_{MH,t}(j)C^*_{MH,t}(j) - W^*_{L,t}L^*_{MH,t}(j)\\ - \frac{\psi_p^{M,*}}{2} P^*_{MH,t}C^*_{MH,t}\left(\frac{P^*_{MH,t}(j)}{P^*_{MH,t-1}(j)} - 1\right)^2.\end{aligned}$$

A fraction $\xi \in [0,1]$ of affiliate operating surplus accrues to the Home parent and is recorded as Home direct investment income. The corresponding Home-currency income flow is

$$\mathcal{E}_t \xi \widetilde{\Pi}^*_{MH,t}(j).$$

The remaining share $1-\xi$ accrues to Foreign residents. We treat ξ as a reduced-form income-accrual parameter rather than modeling payout, retained earnings, local ownership, and affiliate capital accumulation separately.

The Home-controlled affiliate abroad chooses local production labor and its Foreign-currency price, taking the support services allocated by the parent as given. Because prices are subject to Rotemberg adjustment costs, the affiliate's pricing problem is forward looking. For any support allocation, the affiliate solves

$$\max_{\{L^*_{MH,\tau}(j),P^*_{MH,\tau}(j)\}_{\tau\geq t}} E_t \sum_{\tau=t}^{\infty} \Lambda^*_{MH,t,\tau}\widetilde{\Pi}^*_{MH,\tau}(j), \tag{D.9}$$

subject to the affiliate production technology and demand schedule, where $\Lambda^*_{MH,t,\tau}$ is the parent's nominal stochastic discount factor converted into Foreign currency; see Equation (D.14). The affiliate prices to maximize its value to the controlling parent; the constant ξ scales this objective without affecting the optimal policy. Let $\widetilde{\Pi}^{*,opt}_{MH,\tau}(j)$ denote affiliate operating profit at date τ, evaluated at the solution to this dynamic problem. The optimized profit depends on parent support, aggregate states, and lagged prices; we suppress these arguments below to keep notation compact.

The Home parent chooses support hiring, support allocation, domestic production labor, and the domestic MNE price to maximize

$$\max_{\{S_{d,\tau}(j),S_{a,\tau}(j),N_\tau(j),L_{MH,\tau}(j),P_{MH,\tau}(j)\}_{\tau\geq t}} E_t\sum_{\tau=t}^{\infty}\Lambda_{t,\tau}\left[\widetilde{\Pi}_{MH,\tau}(j)+\xi\mathscr{E}_\tau\widetilde{\Pi}^{*,opt}_{MH,\tau}(j)-W_{N,\tau}(j)N_\tau(j)\right],$$

subject to the support constraint

$$N_\tau(j)=S_{d,\tau}(j)+S_{a,\tau}(j), \qquad \tau\geq t,$$

the domestic production technology Equation (D.7), the domestic demand schedule Equation (D.8), and the firm-level support labor supply schedule in Equation (D.1). The variable $\Lambda_{t,\tau}$ denotes the nominal stochastic discount factor of Home households between dates t and τ for Home-currency payoffs. The parent takes production wages as given and internalizes that expanding support employment requires posting a higher firm-specific support wage.

Let $Q_{MH,t}(j)$ denote the Home-currency shadow value of relaxing the domestic unit's date-t production constraint in the parent's problem, and let $Q^*_{MH,t}(j)$ denote the Foreign-currency shadow value of relaxing the affiliate's date-t production constraint in the affiliate's problem Equation (D.9). Both shadow values incorporate the forward-looking pricing conditions generated by Rotemberg adjustment costs. An additional unit of domestic support raises domestic output by $F_S(S_{d,t}(j),L_{MH,t}(j))$ and is therefore worth

$$Q_{MH,t}(j)F_S\big(S_{d,t}(j),L_{MH,t}(j)\big)$$

to the parent. An additional unit of affiliate support raises affiliate output by $F_S(S_{a,t}(j),L^*_{MH,t}(j))$, of which the parent receives the share ξ in Home currency, so its valuation is

$$\mathscr{E}_t\xi Q^*_{MH,t}(j)F_S\big(S_{a,t}(j),L^*_{MH,t}(j)\big).$$

The parent therefore allocates support according to

$$Q_{MH,t}(j)F_S\big(S_{d,t}(j),L_{MH,t}(j)\big)=\mathscr{E}_t\xi Q^*_{MH,t}(j)F_S\big(S_{a,t}(j),L^*_{MH,t}(j)\big)=MC^N_t(j), \tag{D.10}$$

where $MC^N_t(j)$ is the marginal cost of firm-specific support labor implied by Equation (D.1). Under the isoelastic support-labor supply schedule,

$$MC^N_t(j)=\left(1+\frac{1}{\varepsilon_N}\right)W_{N,t}(j).$$

Equation (D.10) is the general-equilibrium version of the support-allocation condition in

Section 4. A Home depreciation raises $\mathscr{E}_t$, increasing the Home-currency marginal value of affiliate support and inducing the parent to shift support toward foreign affiliate operations. Because support services are rival within the parent firm, the increase in affiliate support limits the support available for domestic MNE production.

D.4. Monetary Policy and Market Clearing

Monetary policy follows a Taylor rule with interest-rate smoothing:

$$\widehat{R}_t = \rho_R \widehat{R}_{t-1} + (1-\rho_R)\left(\phi_\pi \widehat{\pi}_t + \phi_y \widehat{Y}_t\right) + \varepsilon_{R,t},$$

where

$$\widehat{R}_t \equiv \log\left(\frac{R_t}{\bar{R}}\right), \qquad \widehat{\pi}_t \equiv \log\left(\frac{\pi_t}{\bar{\pi}}\right), \qquad \widehat{Y}_t \equiv \log\left(\frac{Y_t}{\bar{Y}}\right)$$

denote log deviations of the gross nominal interest rate, inflation, and the Home real-GDP quantity index from their steady-state values. Foreign monetary policy is defined symmetrically.

Goods market clearing. Goods markets clear by sector and location. Home tradable output is absorbed by Home households, Foreign households, and the real resource costs of price adjustment. Aggregate Home tradable output satisfies

$$Y_{T,t} = C_{TH,t} + C^*_{TH,t} + \Phi_{TH,t} + \Phi^*_{TH,t},$$

where $\Phi_{TH,t}$ and $\Phi^*_{TH,t}$ denote the real resource costs of adjusting Home tradable prices in the Home and Foreign markets:

$$\Phi_{TH,t} = \frac{\psi_p^T}{2} C_{TH,t}\left(\frac{P_{TH,t}}{P_{TH,t-1}} - 1\right)^2,$$
$$\Phi^*_{TH,t} = \frac{\psi_p^{T,*}}{2} C^*_{TH,t}\left(\frac{P^*_{TH,t}}{P^*_{TH,t-1}} - 1\right)^2.$$

Foreign tradable market clearing is defined symmetrically.

Multinational goods are consumed locally in the market where they are produced. Home-controlled domestic MNE production satisfies

$$Y_{MH,t} = C_{MH,t} + \Phi_{MH,t},$$

where

$$\Phi_{MH,t} = \frac{\psi_p^M}{2} C_{MH,t}\left(\frac{P_{MH,t}}{P_{MH,t-1}} - 1\right)^2.$$

For Home-controlled affiliates abroad,

$$Y^*_{MH,t} = C^*_{MH,t} + \Phi^*_{MH,t},$$

where

$$\Phi^*_{MH,t} = \frac{\psi_p^{M,*}}{2} C^*_{MH,t} \left(\frac{P^*_{MH,t}}{P^*_{MH,t-1}} - 1 \right)^2 .$$

The corresponding market-clearing conditions for the other two multinational production units are defined symmetrically.

Aggregate real GDP. Aggregate Home real GDP is measured as a fixed-weight quantity index of production located in Home. To first order,

$$\widehat{Y}_t = \bar{s}_T \widehat{Y}_{T,t} + \bar{s}_{MH} \widehat{Y}_{MH,t} + \bar{s}_{MF} \widehat{Y}_{MF,t},$$

where

$$\widehat{Y}_{i,t} \equiv \log\left(\frac{Y_{i,t}}{\bar{Y}_i} \right), \qquad i \in \{T, MH, MF\},$$

and $\bar{s}_T$, $\bar{s}_{MH}$, and $\bar{s}_{MF}$ are the corresponding steady-state value-added shares, which sum to one. Here $Y_{MF,t}$ denotes production by Foreign-controlled affiliates located in Home.

Labor market clearing. Production labor market clearing is

$$L_t = \int_0^1 \left[L_{T,t}(j) + L_{MH,t}(j) + L_{MF,t}(j) \right] dj.$$

Production labor is used by Home-located tradable producers, Home-controlled domestic MNE units, and Foreign-controlled affiliates located in Home. Production labor employed by Home affiliates abroad is supplied in the Foreign economy and enters the Foreign labor-market clearing condition.

Support labor market clearing requires that the aggregate support labor supplied by the household equal the CES aggregate of support labor hired by Home multinational parents.

Bond market clearing. Home-currency bonds are traded domestically and are in zero net supply, so

$$B_{H,t} = 0.$$

The Foreign-currency bond is traded internationally. Let $B^*_{F,t}$ denote the Foreign household's net position in the same bond, expressed in Foreign currency. International bond-market clearing requires

$$B_{F,t} + B^*_{F,t} = 0.$$

D.5. External Adjustment and Measurement

Trade balance. The nominal trade balance records cross-border sales of tradable goods:

$$TB_t = \mathcal{E}_t P^*_{TH,t} C^*_{TH,t} - P_{TF,t} C_{TF,t}. \tag{D.11}$$

The first term is Home export revenue from tradable goods sold to Foreign households, converted into Home currency. The second term is Home expenditure on imported Foreign tradable goods. Multinational goods are produced and consumed locally, so cross-border multinational activity affects external earnings through direct investment income rather than through the trade balance.

Combining the Home household budget constraint in Equation (D.2) with goods-market clearing gives the nominal external budget constraint below. In this aggregation, the price-adjustment expenditures subtracted from the profits of Home-located units cancel against the corresponding adjustment absorption of Home sectoral output in the clearing conditions, so no adjustment-cost terms remain in the external constraint apart from those inside the affiliate surpluses that define DII_t:

$$\mathscr{E}_t B_{F,t} + P_t \frac{\psi_b}{2}\left(\frac{\mathscr{E}_t B_{F,t}}{P_t} - \bar{b}_F\right)^2 = R^*_{t-1}\exp(\psi_{rp,t-1})\mathscr{E}_t B_{F,t-1} + TB_t + DII_t. \tag{D.12}$$

Here TB_t is the nominal trade balance defined in Equation (D.11), and DII_t is the nominal direct investment income balance

$$DII_t = \mathscr{E}_t \xi \widetilde{\Pi}^*_{MH,t} - \xi^* \widetilde{\Pi}_{MF,t},$$

where the first term is income received by Home residents from Home-controlled affiliates abroad, and the second term is income paid by Home-located affiliates to Foreign parents.Consistent with the convention in Section 5.4, the model assigns no separate transfer price to parent support. Hence $S_{a,t}$ generates neither a service export for Home nor an intermediate-service import for Foreign, and $S^*_{a,t}$ generates neither a service import for Home nor a service export for Foreign. The Home parent's entire support wage bill is already subtracted in its objective, while affiliate operating surplus enters the external accounts only through the DII_t terms above.

Define the Home-currency value of the model's Foreign-currency bond position as

$$NFA_t \equiv \mathscr{E}_t B_{F,t}.$$

Subtracting $\mathscr{E}_{t-1} B_{F,t-1}$ from both sides of Equation (D.12) gives

$$\underbrace{\mathscr{E}_t B_{F,t} - \mathscr{E}_{t-1} B_{F,t-1}}_{\Delta NFA_t} = \underbrace{CA_t}_{\text{current account}} + \underbrace{\left(\mathscr{E}_t - \mathscr{E}_{t-1}\right) B_{F,t-1}}_{\text{exchange-rate valuation effect}} \underbrace{- P_t \frac{\psi_b}{2}\left(\frac{\mathscr{E}_t B_{F,t}}{P_t} - \bar{b}_F\right)^2}_{\text{portfolio adjustment cost}}, \tag{D.13}$$

where the model current account is

$$CA_t = TB_t + DII_t + \mathscr{E}_t\left[R^*_{t-1}\exp(\psi_{rp,t-1}) - 1\right] B_{F,t-1}.$$

The first term of CA_t is the trade balance, the second term is the direct investment income balance, and the third term is portfolio income on the existing Foreign-currency bond position. Because the model contains only a one-period Foreign-currency bond, the valuation term in

Equation (D.13) reflects exchange-rate revaluation of the principal. Interest payments on the bond are portfolio income and therefore enter the current account.

The mechanism emphasized in the paper concerns the production-related components of the current account: the trade balance and the direct investment income balance. The trade-balance component is linked to domestic tradable production, whereas the direct-investment-income component is linked to foreign affiliate activity. In standard models without multinational income, $DII_t = 0$, so this production-related current-account response operates entirely through the trade balance. In the quantitative exercises, the steady-state Foreign-currency bond position is set to zero, which ensures that portfolio-income and bond-valuation terms do not drive the main transmission results.

Relation to balance-of-payments measurement. The model abstracts from several measurement issues that arise in the data. First, exchange-rate movements affect the Home-currency value of foreign affiliate profits immediately in the model. In firm accounts, however, reported profits and affiliate-to-parent payments may respond more gradually because of transfer-pricing practices, invoicing conventions, hedging, and accounting rules. We therefore interpret the model's exchange-rate translation of affiliate profits as a shadow value of foreign operations rather than as a prediction of one-for-one contemporaneous pass-through into measured payments.

Second, balance-of-payments statistics include reinvested earnings as part of direct investment income even when no cash payment crosses borders. The offsetting accounting entry is recorded in the financial account as reinvestment in the foreign affiliate. Our model abstracts from this reinvestment margin by fixing affiliate capital and treating ξ as a reduced-form Home claim on affiliate operating surplus. In the quantitative exercise, ξ is disciplined using affiliate-to-parent payments relative to affiliate income. The residual share of operating surplus accrues to Foreign residents rather than financing retained earnings or affiliate capital accumulation. This convention focuses the analysis on the short-run exchange-rate transmission channel through affiliate profitability and parent support allocation. Endogenizing payout, retained earnings, and affiliate capital accumulation is beyond the scope of this paper but is a natural extension for studying lower-frequency multinational investment responses.

D.6. Full Equilibrium System

This subsection shows the equilibrium system of the quantitative baseline. Variables without asterisks denote Home and starred variables denote Foreign. M denotes multinational goods and T denotes tradable goods. Within the multinational sector, MH is output of a Home-controlled producer and MF is output of a Foreign-controlled producer; for example, $Y^*_{MH,t}$ is output of the Home affiliate located in Foreign. All quantities and relative prices are expressed in units of the country-specific CPI:

$$q_t \equiv \frac{\mathscr{E}_t P_t^*}{P_t}, \qquad b_t \equiv \frac{\mathscr{E}_t B_{F,t}}{P_t}, \qquad p_{i,t} \equiv \frac{P_{i,t}}{P_t}, \qquad \pi_{i,t} \equiv \frac{P_{i,t}}{P_{i,t-1}}.$$

Similarly, $w_{k,t} = W_{k,t}/P_t$ for $k \in \{L, N\}$. The Foreign demand weights and parent-income claim are ω_M^*, ω_T^*, and ξ^*; in the symmetric benchmark they equal their Home counterparts.

Households and international financial markets. Households have the GHH composite and marginal utility

$$X_t = C_t - \chi_L \frac{L_t^{1+\varphi}}{1+\varphi} - \chi_N \frac{N_t^{1+\varphi}}{1+\varphi}, \qquad \lambda_t = X_t^{-\sigma}.$$

The Euler equations for Home- and Foreign-currency bonds are

$$\begin{aligned} 1 &= \beta \mathbb{E}_t \left[\frac{\lambda_{t+1}}{\lambda_t} \frac{R_t}{\pi_{t+1}} \right], \\ 1 &= \beta \mathbb{E}_t \left[\frac{\lambda_{t+1}}{\lambda_t} \frac{R_t^*}{\pi_{t+1}^*} \frac{q_{t+1}}{q_t} \frac{\exp(\psi_{rp,t})}{1 + \psi_b (b_t - \bar{b})} \right]. \end{aligned}$$

The Foreign household obeys the starred counterpart of the first equation. GHH preferences imply labor-supply conditions without a wealth-effect term:

$$\chi_L L_t^{\varphi} = w_{L,t}, \qquad \chi_N N_t^{\varphi} = w_{N,t},$$

with analogous Foreign conditions.

Demand, prices, and price dynamics. Home consumption is a three-level CES aggregate. In CPI-relative prices, its price indices and demands are

$$\begin{aligned} p_{M,t} &= \left[\omega_M p_{MH,t}^{1-\eta_M} + (1-\omega_M) p_{MF,t}^{1-\eta_M} \right]^{1/(1-\eta_M)}, \\ p_{T,t} &= \left[\omega_T p_{TH,t}^{1-\eta_T} + (1-\omega_T) p_{TF,t}^{1-\eta_T} \right]^{1/(1-\eta_T)}, \\ 1 &= \omega p_{M,t}^{1-\eta} + (1-\omega) p_{T,t}^{1-\eta}, \\ C_{M,t} &= \omega p_{M,t}^{-\eta} C_t, \qquad C_{T,t} = (1-\omega) p_{T,t}^{-\eta} C_t, \\ C_{MH,t} &= \omega_M \left(\frac{p_{MH,t}}{p_{M,t}} \right)^{-\eta_M} C_{M,t}, \qquad C_{MF,t} = (1-\omega_M) \left(\frac{p_{MF,t}}{p_{M,t}} \right)^{-\eta_M} C_{M,t}, \\ C_{TH,t} &= \omega_T \left(\frac{p_{TH,t}}{p_{T,t}} \right)^{-\eta_T} C_{T,t}, \qquad C_{TF,t} = (1-\omega_T) \left(\frac{p_{TF,t}}{p_{T,t}} \right)^{-\eta_T} C_{T,t}. \end{aligned}$$

Foreign demands are the mirror image, with the Foreign-controlled multinational good MF receiving the home-bias weight ω_M^* and the Foreign tradable receiving the weight ω_T^*. Relative prices evolve according to $p_{i,t}/p_{i,t-1} = \pi_{i,t}/\pi_t$ in Home and $p_{i,t}^*/p_{i,t-1}^* = \pi_{i,t}^*/\pi_t^*$ in Foreign.

Production, rival support, and marginal costs. The four multinational production technologies are

$$Y_{MH,t} = A_{M,t} S_{d,t}^{\alpha_S} L_{MH,t}^{1-\alpha_S}, \qquad Y_{MF,t} = A_{M,t} (S_{a,t}^*)^{\alpha_S} L_{MF,t}^{1-\alpha_S},$$
$$Y_{MF,t}^* = A_{M,t}^* (S_{d,t}^*)^{\alpha_S} (L_{MF,t}^*)^{1-\alpha_S}, \qquad Y_{MH,t}^* = A_{M,t}^* S_{a,t}^{\alpha_S} (L_{MH,t}^*)^{1-\alpha_S}.$$

Tradable technologies are linear:

$$Y_{T,t} = A_{T,t} L_{T,t}, \qquad Y_{T,t}^* = A_{T,t}^* L_{T,t}^*.$$

There are no intermediate inputs, support-efficiency shocks, or separate head-office employment category in the baseline.

Support services are rival within each parent and one-for-one in support labor:

$$N_t = S_{d,t} + S_{a,t}, \qquad N_t^* = S_{d,t}^* + S_{a,t}^*.$$

The allocation and support-wage conditions are

$$\frac{mc_{MH,t} p_{MH,t} Y_{MH,t}}{S_{d,t}} = \xi q_t \frac{mc_{MH,t}^* p_{MH,t}^* Y_{MH,t}^*}{S_{a,t}},$$
$$\frac{mc_{MF,t}^* p_{MF,t}^* Y_{MF,t}^*}{S_{d,t}^*} = \frac{\xi^*}{q_t} \frac{mc_{MF,t} p_{MF,t} Y_{MF,t}}{S_{a,t}^*},$$
$$w_{N,t}\left(1 + \frac{1}{\varepsilon_N}\right) = \alpha_S \frac{mc_{MH,t} p_{MH,t} Y_{MH,t}}{S_{d,t}},$$
$$w_{N,t}^*\left(1 + \frac{1}{\varepsilon_N}\right) = \alpha_S \frac{mc_{MF,t}^* p_{MF,t}^* Y_{MF,t}^*}{S_{d,t}^*}.$$

The first two equations are the rival-allocation conditions; the last two embed the parent support-labor monopsony wedge. Production-labor marginal costs are

$$mc_{MH,t} = \frac{1}{1-\alpha_S} \frac{w_{L,t}}{p_{MH,t}} \frac{L_{MH,t}}{Y_{MH,t}}, \qquad mc_{MF,t} = \frac{1}{1-\alpha_S} \frac{w_{L,t}}{p_{MF,t}} \frac{L_{MF,t}}{Y_{MF,t}},$$
$$mc_{MF,t}^* = \frac{1}{1-\alpha_S} \frac{w_{L,t}^*}{p_{MF,t}^*} \frac{L_{MF,t}^*}{Y_{MF,t}^*}, \qquad mc_{MH,t}^* = \frac{1}{1-\alpha_S} \frac{w_{L,t}^*}{p_{MH,t}^*} \frac{L_{MH,t}^*}{Y_{MH,t}^*},$$
$$mc_{TH,t} = \frac{w_{L,t}}{A_{T,t} p_{TH,t}}, \qquad mc_{TH,t}^x = mc_{TH,t} \frac{p_{TH,t}}{q_t p_{TH,t}^*},$$
$$mc_{TF,t} = mc_{TF,t}^* \frac{p_{TF,t}^* q_t}{p_{TF,t}}, \qquad mc_{TF,t}^* = \frac{w_{L,t}^*}{A_{T,t}^* p_{TF,t}^*}.$$

Local-currency price setting. Price adjustment is Rotemberg at the variety level in every market. For a variety i, the pricing equation can be written as

$$(1-\varepsilon_i) + \varepsilon_i mc_{i,t} = \psi_i \left[\pi_{i,t}(\pi_{i,t} - 1) - \beta \mathbb{E}_t \left(\mathcal{Q}_{i,t,t+1} \frac{C_{i,t+1}}{C_{i,t}} \pi_{i,t+1}^2 (\pi_{i,t+1} - 1) \right) \right].$$

Here ε_i is ε_M for multinational varieties and ε_T for tradable varieties. The adjustment-cost coefficients are ψ_p^M for parent multinational goods, ψ_p^A for affiliate multinational goods, ψ_p^T for domestic tradables, and ψ_p^X for export tradables. The relevant discount factors are

$$\begin{aligned}
\mathcal{Q}_{MH,t,t+1} &= \frac{\lambda_{t+1}}{\lambda_t \pi_{t+1}}, & \mathcal{Q}_{MF,t,t+1} &= \Omega_{MF,t,t+1}, \\
\mathcal{Q}_{MF^*,t,t+1} &= \frac{\lambda^*_{t+1}}{\lambda^*_t \pi^*_{t+1}}, & \mathcal{Q}_{MH^*,t,t+1} &= \Omega_{MH^*,t,t+1}, \\
\mathcal{Q}_{TH,t,t+1} &= \frac{\lambda_{t+1}}{\lambda_t \pi_{t+1}}, & \mathcal{Q}_{TH^*,t,t+1} &= \frac{\lambda_{t+1}}{\lambda_t} \frac{q_{t+1}}{q_t \pi^*_{t+1}}, \\
\mathcal{Q}_{TF,t,t+1} &= \frac{\lambda^*_{t+1}}{\lambda^*_t} \frac{q_t}{q_{t+1} \pi_{t+1}}, & \mathcal{Q}_{TF^*,t,t+1} &= \frac{\lambda^*_{t+1}}{\lambda^*_t \pi^*_{t+1}},
\end{aligned}$$

where the affiliate discount factors are the controlling parents' discount factors converted into the currency of the affiliate's revenues,

$$\begin{aligned}
\Omega_{MF,t,t+1} &= \frac{\lambda^*_{t+1}}{\lambda^*_t} \frac{q_t}{q_{t+1}} \frac{1}{\pi_{t+1}}, \\
\Omega_{MH^*,t,t+1} &= \frac{\lambda_{t+1}}{\lambda_t} \frac{q_{t+1}}{q_t} \frac{1}{\pi^*_{t+1}}.
\end{aligned} \tag{D.14}$$

Each affiliate thus discounts exactly as an exporter from its parent's country does, $\Omega_{MF} = \mathcal{Q}_{TF}$ and $\Omega_{MH^*} = \mathcal{Q}_{TH^*}$; the $1-\xi$ local claim on affiliate surplus affects income flows, not the pricing objective.

Goods, labor, and external-account clearing. Let $\mathscr{A}(\pi;\psi) \equiv 1 + \frac{\psi}{2}(\pi-1)^2$ denote the gross-output factor implied by price adjustment: under the convention stated in Appendix D.2, adjustment expenditures purchase the sectoral composite, so sectoral gross output equals household demand plus adjustment absorption. Goods-market and production-labor clearing are

$$\begin{aligned}
Y_{MH,t} &= C_{MH,t}\mathscr{A}(\pi_{MH,t};\psi_p^M), & Y_{MF,t} &= C_{MF,t}\mathscr{A}(\pi_{MF,t};\psi_p^A), \\
Y^*_{MF,t} &= C^*_{MF,t}\mathscr{A}(\pi^*_{MF,t};\psi_p^M), & Y^*_{MH,t} &= C^*_{MH,t}\mathscr{A}(\pi^*_{MH,t};\psi_p^A), \\
Y_{T,t} &= C_{TH,t}\mathscr{A}(\pi_{TH,t};\psi_p^T) + C^*_{TH,t}\mathscr{A}(\pi^*_{TH,t};\psi_p^X), \\
Y^*_{T,t} &= C^*_{TF,t}\mathscr{A}(\pi^*_{TF,t};\psi_p^T) + C_{TF,t}\mathscr{A}(\pi_{TF,t};\psi_p^X), \\
L_t &= L_{T,t} + L_{MH,t} + L_{MF,t}, & L^*_t &= L^*_{T,t} + L^*_{MF,t} + L^*_{MH,t}.
\end{aligned}$$

External balance.

$$
\begin{aligned}
tb_t &= q_t p^*_{TH,t} C^*_{TH,t} - p_{TF,t} C_{TF,t}, \\
dii^{in}_t &= \xi q_t \left[p^*_{MH,t} C^*_{MH,t} - w^*_{L,t} L^*_{MH,t} - \frac{\psi^A_p}{2} p^*_{MH,t} C^*_{MH,t} (\pi^*_{MH,t} - 1)^2 \right], \\
dii^{out}_t &= \xi^* \left[p_{MF,t} C_{MF,t} - w_{L,t} L_{MF,t} - \frac{\psi^A_p}{2} p_{MF,t} C_{MF,t} (\pi_{MF,t} - 1)^2 \right], \\
dii_t &= dii^{in}_t - dii^{out}_t, \\
ca_t &= tb_t + dii_t + \left(R^*_{t-1} \exp(\psi_{rp,t-1}) - 1 \right) \frac{q_t}{q_{t-1}} \frac{b_{t-1}}{\pi^*_t}, \\
b_t + \frac{\psi_b}{2}(b_t - \bar{b})^2 &= R^*_{t-1} \exp(\psi_{rp,t-1}) \frac{q_t}{q_{t-1}} \frac{b_{t-1}}{\pi^*_t} + tb_t + dii_t, \\
tot_t &= \frac{p_{TF,t}}{q_t p^*_{TH,t}}.
\end{aligned}
$$

Thus, DII measures the parent's claim on affiliate operating surplus after local production-labor costs and the affiliate's price-adjustment expenditures; it is not a trade flow. There is no double-counting of adjustment costs in this accounting: production labor produces the affiliate's gross output, the affiliate's adjustment expenditure purchases sectoral output that appears as adjustment absorption in the clearing conditions above, and the same expenditure is subtracted once from the surplus available for distribution.

Aggregate output, policy, shocks, and steady state. The physical-output aggregate used in the Taylor rule, and Home GDP, are

$$
\begin{aligned}
Y_t &= p_{MH,ss} C_{MH,t} + p_{MF,ss} C_{MF,t} + p_{TH,ss} C_{TH,t} + q_{ss} p^*_{TH,ss} C^*_{TH,t}, \\
gdp_t &= p_{MH,t} C_{MH,t} + p_{MF,t} C_{MF,t} + p_{TH,t} C_{TH,t} + q_t p^*_{TH,t} C^*_{TH,t}.
\end{aligned}
$$

GDP is measured as final sales of Home-located production, treating all sectors symmetrically. Under the price-adjustment convention above, adjustment resources are produced and used up within the firm; they are intermediate own-consumption and are therefore excluded from value added. This convention is innocuous for the quantitative results: the adjustment factor $\mathscr{A}(\pi;\psi)$ equals one at the zero-inflation steady state and its first-order deviation vanishes, so gross output and final sales coincide to first order and the GDP measure is identical under either convention in the linearized solution. The Home monetary rule is

$$
\log \frac{R_t}{\bar{R}} = \rho_R \log \frac{R_{t-1}}{\bar{R}} + (1 - \rho_R) \left[\phi_\pi \log \pi_t + \phi_y \log \frac{Y_t}{\bar{Y}} \right] + e^m_t,
$$

with a starred counterpart for Foreign. The exogenous processes are

$$\log A_{M,t} = \rho_{A_M} \log A_{M,t-1} + e_{A_M,t}, \qquad \log A_{T,t} = \rho_{A_T} \log A_{T,t-1} + e_{A_T,t},$$
$$\log A^*_{M,t} = \rho_{A_M} \log A^*_{M,t-1} + e_{A^*_M,t}, \qquad \log A^*_{T,t} = \rho_{A_T} \log A^*_{T,t-1} + e_{A^*_T,t},$$
$$\psi_{rp,t} = \rho_\psi \psi_{rp,t-1} + e_{\psi,t}.$$

Together with the two monetary shocks, this gives seven structural shocks.

The non-stochastic symmetric steady state has $A_M = A^*_M = A_T = A^*_T = 1$, $\psi_{rp} = 0$, $b = \bar{b} = 0$, unit inflation, $q = 1$, and $\bar{R} = \bar{R}^* = 1/\beta$. Multinational marginal costs equal $(\varepsilon_M - 1)/\varepsilon_M$ and tradable marginal costs equal $(\varepsilon_T - 1)/\varepsilon_T$. Symmetry implies equal gross DII receipts and payments, and hence $dii = tb = ca = 0$.

E Mapping Parent-Level Elasticities to the Support-Allocation Block

This appendix describes how the parent-level long-difference evidence disciplines the support block of the quantitative model. The calibration is a regression-to-regression mapping. For a candidate set of support-block parameters, we simulate model parents with the same revenue-side exchange-rate shocks used in the empirical parent regressions, solve the parent block firm by firm in partial equilibrium, and run the corresponding cross-sectional regression on the simulated model outcomes. The support-block parameters are chosen so that the model reproduces the relative responses estimated in the parent data.

This procedure separates two objects. The parent regressions identify how outcomes differ across parents with different revenue-side exchange-rate exposure during 2020–22. The aggregate exchange-rate experiment in Section 6.2 then determines how the calibrated support block responds in general equilibrium to a common depreciation.

E.1. Empirical Regression and Simulated Counterpart

The empirical parent regression relates each parent's 2020–22 log outcome change to its revenue-side foreign-exchange exposure, controlling for purchase-side exposure, industry fixed effects, and pre-period controls. Let E_j denote the revenue-side exchange-rate shock for parent j. For outcome y_j, the empirical specification can be written as

$$\Delta \ln y_j^{data} = a_y + \beta_y^{data} E_j + \Gamma'_y X_j + \delta_{i(j)} + u_{j,y},$$

where X_j contains purchase-side exposure and pre-period controls, and $\delta_{i(j)}$ denotes industry fixed effects. The coefficient β_y^{data} is identified from cross-sectional variation in revenue-side exposure across parents, conditional on these controls.

The model counterpart uses the same cross-sectional object. Parents are ex ante identical in the quantitative model, but in the micro-calibration exercise they are exposed to different realized revenue-side exchange-rate shocks E_j. For a candidate parameter vector $\theta = (\alpha_S, \varepsilon_N)$, we solve the parent block for each observed E_j, holding aggregate prices, wages, and demand shifters fixed at their symmetric steady-state values. Let $\Delta \ln y_j^{model}(\theta)$ denote the resulting simulated log change in the model counterpart of outcome y_j. We then run the simulated regression

$$\Delta \ln y_j^{model}(\theta) = a_y^{model}(\theta) + \beta_y^{model}(\theta) E_j + u_{j,y}^{model}(\theta). \tag{E.1}$$

The calibration chooses α_S and ε_N so that the simulated slopes match the empirical parent-regression slopes for support employment and non-support parent employment:

$$\beta_{H_N}^{model}(\alpha_S, \varepsilon_N) = \beta_{H_N}^{data}, \qquad \beta_{H_O}^{model}(\alpha_S, \varepsilon_N) = \beta_{H_O}^{data}. \tag{E.2}$$

The intercept in Equation (E.1) absorbs the component of the outcome change that is common across parents in the simulated cross section.[6] To see what the intercept does absorb, suppose the revenue-side shock can be written as

$$E_j = E_0 + \widetilde{E}_j,$$

where E_0 is a common component and $\widetilde{E}_j$ is relative exposure. Then

$$a_y^{model} + \beta_y^{model} E_j = \left(a_y^{model} + \beta_y^{model} E_0\right) + \beta_y^{model} \widetilde{E}_j.$$

The common component moves the intercept, and the slope is identified from relative exposure. Because the parent block is nonlinear in exposure, however, the slope is itself a function of the point around which the cross-section is evaluated: β_y^{model} depends on E_0. Since the empirical slopes were estimated on outcomes generated under the 2020–22 depreciation, the model counterpart evaluates the parent block at the same realized exposures, and the aggregate UIP experiment continues to supply the general-equilibrium response to the common exchange-rate component.

The calibration uses the parent-level exchange-rate shocks from the 2020–22 episode.[7] We use the raw shock E_j rather than demeaning it before the nonlinear firm solve. This procedure adds shock-realization heterogeneity across parents without introducing ex-ante firm heterogeneity.

[6] The simulated regression is otherwise unconditional—the model has no counterpart of purchase-side exposure or industry composition—so the calibration matches conditional empirical slopes with unconditional model slopes generated by the same mechanism. This mismatch is quantitatively innocuous here because the two exposures are nearly orthogonal in the shock data (correlation 0.04); partialling the purchase-side exposure out of the simulated regression moves the calibrated parameters by less than 0.1%.

[7] Of the 4,050 parents with a revenue-side shock, we retain the 3,638 with both a revenue- and a purchase-side shock, mirroring the parent regressions' conditioning on purchase-side exposure. The sample-construction table reports a smaller count, 3,416, because it additionally requires the matched parent-survey data used in the regressions; the calibration needs only the exposure distribution, so it retains all parents with constructible shocks. The distinction is immaterial for the calibrated parameters.

E.2. Parent-Block Model Counterpart

The simulated regression is generated from the parent block of the quantitative model in partial equilibrium. The parent allocates support services between domestic MNE production and foreign-affiliate activity. Let $S_{d,j}$ denote support allocated to domestic production and $S_{a,j}$ support allocated to the foreign affiliate. The parent-block support-allocation condition is

$$\frac{mc_{MH,j}\, p_{MH,j}\, Y_{MH,j}}{S_{d,j}} = \xi \exp(E_j) \frac{mc^*_{MH,j}\, p^*_{MH,j}\, Y^*_{MH,j}}{S_{a,j}}.$$

The term $\exp(E_j)$ captures the parent-currency value of affiliate revenue. A higher revenue-side exchange-rate exposure raises the return to support used abroad relative to support used in domestic production. This is the model force that generates the cross-sectional responses in support employment, non-support parent employment, and wages.

For the two-year long-difference calibration, we use the flexible-price parent block. With goods demand elasticity ε_M, define

$$\varrho = \frac{\varepsilon_M - 1}{\varepsilon_M}, \qquad \phi = \frac{\varrho \alpha_S}{1 - \varrho(1-\alpha_S)}, \qquad \psi_Y = \alpha_S + (1-\alpha_S)\phi,$$

where α_S is the support-input share in MNE production and the elasticity ε_N governs the firm-level support-labor supply curve. Combining the allocation condition with the support-labor supply condition gives the closed-form parent responses

$$\frac{S_{d,j}}{S_{a,j}} = \frac{\bar{S}_d}{\bar{S}_a} \exp\left(\frac{E_j}{\phi - 1}\right), \tag{E.3}$$

$$\left(\frac{S_{d,j}}{\bar{S}_d}\right)^{\phi-1} = \left(\frac{N_j}{\bar{N}}\right)^{1/\varepsilon_N}. \tag{E.4}$$

Equation (E.3) is the within-parent allocation margin: higher exposure shifts support between domestic and affiliate use. Equation (E.4) is the scale margin: changes in the return to support affect total support services N_j. Domestic MNE production labor satisfies

$$L_{MH,j} \propto S^{\phi}_{d,j}. \tag{E.5}$$

Thus, for each observed exposure E_j, the parent block produces model outcomes N_j, $L_{MH,j}$, and wages. These simulated outcomes are then mapped into the parent data and used in the simulated regression Equation (E.1).

E.3. Mapping Model Variables to Parent Data

The parent data report employment categories, while the model is written in terms of support services and production labor. We map the model variables into the data as follows. Model support services N_j map into observed non-production support employment $H_{N,j}$. Model domestic MNE production labor $L_{MH,j}$ maps into observed other parent employment $H_{O,j}$, the residual category that excludes both head-office and support employment:

$$\Delta \ln H_{N,j} \longleftrightarrow \Delta \ln N_j, \qquad \Delta \ln H_{O,j} \longleftrightarrow \Delta \ln L_{MH,j}.$$

The two slope targets in Equation (E.2) are estimated on the support and other-employment categories, so their empirical counterparts and the model objects are defined on the same head-office-exclusive basis.

The variable N_j is measured in units of support services, not in worker headcount, and the units of support services are arbitrary up to a scale normalization. We therefore introduce two constant measurement conversions: Z_N converts support-service units into observed support workers, and μ converts the model price of one support-service unit into an observed wage per support worker,

$$H_{N,j} = \frac{N_j}{Z_N}, \qquad w^H_{N,j} = \mu\, w_{N,j}.$$

Here $w^H_{N,j}$ is the observed support wage per support worker, while $w_{N,j}$ is the model price of one unit of support services; because μ is constant, the observed support wage inherits the firm-level monopsony variation of the model wage. Both conversions are measurement constants: they enter no first-order condition and leave firm costs, the support allocation, and the aggregate equilibrium unchanged.

The value of Z_N is chosen to match the observed support-worker share on the model's own denominator. Because the model has two labor categories, the relevant share is support employment relative to support plus other employment, with the head office excluded on both sides. Let s_H denote this share. In steady state,

$$s_H = \frac{\bar{H}_N}{\bar{H}_N + \bar{H}_O} = \frac{\bar{N}/Z_N}{\bar{N}/Z_N + \bar{L}_{MH}}.$$

In the parent data shown in Table A.1, the mean shares of total employment are 43.2% head office, 22.5% support, and 34.3% other, so the model-denominator share is $s_H = 22.5/(22.5+34.3) = 0.396$, and the target identifies $Z_N = 7.26$. The two share concepts are mutually consistent: $0.396 \times (1 - 0.432) = 0.225$ recovers the observed support share of *total* parent employment. Alternative aggregations of the same data give 0.358 (employment-weighted aggregate) and 0.517 (mean of within-parent ratios); we use the ratio of mean shares and avoid the mean of ratios, which is dominated by parents with near-zero denominators, for the same skewness reason discussed for the affiliate payout ratio. Since Z_N is constant across firms, it does not affect support-employment elasticities:

$$\Delta \ln H_{N,j} = \Delta \ln N_j.$$

It only affects levels and the headcount weights used to compute average parent wages.

The observed average parent wage is constructed as the headcount-weighted average of support and non-support parent wages:

$$\bar{w}_j = \frac{w^H_{N,j} H_{N,j} + w_{L,j} L_{MH,j}}{H_{N,j} + L_{MH,j}}.$$

The wage conversion μ is chosen so that the simulated average-wage slope matches the estimated 0.169. The slope identifies μ because the average wage responds to exposure mainly through composition—high-exposure parents shift employment toward the support category—and this composition effect is stronger the higher the support wage relative to the production wage. We report μ through the implied steady-state support-wage premium, $\rho \equiv \bar{w}^H_N / w_L = 1.57$, a plausible magnitude for the commercial and sales staff in the support category.

E.4. Identification of α_S, ε_N, and the Measurement Conversions

The micro calibration uses four targeted moments. The support employment and other employment slopes jointly identify α_S and ε_N, the support-worker share identifies Z_N, and the average-wage slope identifies μ. Because the calibration is exactly identified, the targeted rows of Table E.1 fit by construction, and the informative rows are the untargeted ones.

Identification of the economic parameters is transparent in the closed forms Equation (E.3)–Equation (E.5). Both simulated employment outcomes are exact scalar multiples of a single function of exposure, so the *ratio* of the other-employment slope to the support-employment slope depends on $(\alpha_S, \varepsilon_N)$ only through ϕ and $1/\varepsilon_N$ and pins ε_N given α_S; the *level* of the support-employment slope then pins α_S. Economically, ε_N governs the scale margin—how much total support the parent can add when its return rises—while α_S governs how strongly the within-parent reallocation of support drags production labor.[8]

Table E.1 reports the fit. Among the untargeted checks, the total-employment slope is worth noting. Total observed parent employment weights the support and other categories at their headcount shares, and at the corrected support share the model delivers a slope of −0.009 against −0.014 in the data: the model reproduces the flat total-employment response in Table 2 as a near-cancellation of the support and other-employment margins.

E.5. Robustness to the Other-Employment Slope Target

The other-employment slope is the less precisely estimated of the two employment targets, and its point estimate varies across specifications: −0.270 without controls and fixed effects, −0.297 in the baseline, and −0.349 in the manufacturing-only sample, with the exposure-construction robustness estimates interior to this range. Table E.2 recalibrates $(\alpha_S, \varepsilon_N)$ at each of these values, and at a stress point one exposure-robust standard error above the baseline estimate (−0.124),

[8] The measurement conversions are block-recursive: (Z_N, μ) are computed after $(\alpha_S, \varepsilon_N)$, and no aggregate result in the paper depends on them.

Table E.1: Parent-level micro calibration: data and model moments

Moment	Data	Model	Role
Support-employment slope	0.384	0.384	Targeted (α_S, ε_N)
Other-employment slope	−0.297	−0.297	Targeted (α_S, ε_N)
Support-worker share, support+other denominator	0.396	0.396	Targeted (Z_N)
Average-wage slope	0.169	0.169	Targeted (μ)
Total-employment slope	−0.014	−0.009	Untargeted
Support p90–p10 effect	0.093	0.102	Untargeted
Other-employment p90–p10 effect	−0.070	−0.079	Untargeted
Average-wage p90–p10 effect	0.042	0.045	Untargeted

Notes: "Slope" is the coefficient from a long-difference regression on the parent revenue-side exchange-rate shock. Data coefficients are the conditional parent-regression estimates of Table 2, estimated on outcome-specific samples (support N=1,125; other N=1,446; average wage and total employment N=1,811); model coefficients come from intercept-only regressions on the 3,638 parents with both exchange-rate shocks. Calibrated values: $\alpha_S = 0.361$, $\varepsilon_N = 2.336$, $Z_N = 7.26$, and μ reported as the steady-state support-wage premium $\rho = 1.57$. The model holds aggregate objects fixed in this calibration exercise.

Table E.2: Robustness to the other-employment slope target

Other-employment target	α_S	ε_N	DII	DII endog.	Fixed RGDP	Fixed RGNI	CPI
No controls (−0.270)	0.343	2.443	0.518	0.248	0.111	0.629	3.725
Baseline (−0.297)	0.361	2.336	0.515	0.245	0.108	0.623	3.729
Manufacturing only (−0.349)	0.395	2.171	0.509	0.239	0.103	0.612	3.736
Stress: +1 robust SE (−0.124)	0.234	3.626	0.530	0.260	0.120	0.650	3.683

Notes: Each row recalibrates $(\alpha_S, \varepsilon_N)$ to the indicated other-employment slope, holding the support-employment target at 0.384, re-pins (ω, ω_T) to the gross-flow targets, and re-runs the UIP experiment normalized to a 10% impact real depreciation. DII is the impact net direct-investment-income response in percentage points of steady-state GDP; the endogenous component subtracts the pure currency-valuation term, which equals 0.27 percentage points in every row because the re-pinned calibration holds steady-state DII credits at 2.7% of GDP. Fixed-price RGDP and RGNI are impact percent changes; CPI is the annualized impact inflation response in percent. The stress row targets the baseline estimate plus one exposure-robust standard error.

holding the support-employment target at 0.384, re-pinning (ω, ω_T) to the gross-flow targets, and re-running the baseline depreciation experiment.

The aggregate results are robust across these targets: over all rows of Table E.2, the impact direct-investment-income response stays within 0.509–0.530 percentage points of GDP, real GDP, real GNI, and inflation vary comparably little, and the organizational reallocation is preserved at every grid point (affiliate support rises by 6.3–7.3% and domestic support falls by 0.58–0.74% on impact). The support-employment slope, held at its estimated value throughout, is the primary discipline on the aggregate response; the other-employment slope only divides it between α_S and ε_N, and the division largely cancels in the aggregate.

Figure F.1: External adjustment after a Foreign monetary policy shock

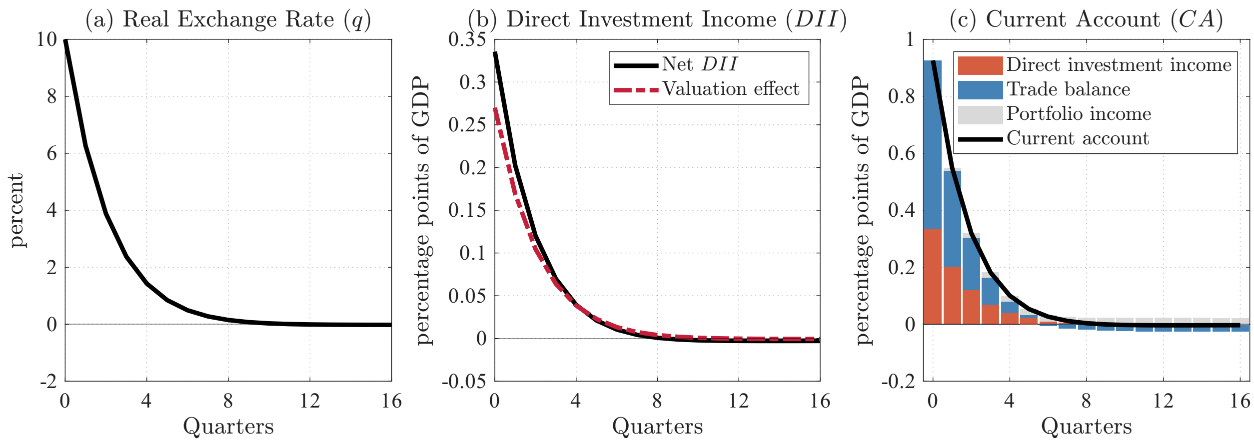

Notes: Impulse responses to a Foreign monetary policy shock, normalized to a 10% impact real depreciation. Panel (a) is in percent deviation from steady state; panels (b)–(c) are in percentage points of steady-state GDP. Panel (b) plots net direct-investment income against its pure currency-valuation component, $\xi\, q\, \bar{\tilde{\Pi}}^*_{MH} - \xi^* \bar{\tilde{\Pi}}_{MF}$, computed with the exchange rate on its dynamic path but the affiliate operating surpluses fixed at steady state. Panel (c) decomposes the current account into the trade balance, direct investment income, and portfolio income.

F Alternative Depreciation Shocks: Valuation and Affiliate Production

The baseline UIP experiment isolates the exchange-rate transmission mechanism by generating a Home depreciation without directly shocking Foreign demand. In this appendix, we consider two alternative shocks that generate the same 10% impact depreciation but move the profitability of foreign affiliate expansion in opposite directions. A contractionary Foreign monetary policy shock weakens Foreign demand and attenuates the real affiliate-income response. A Foreign discount-factor shock strengthens current Foreign demand and amplifies the real affiliate-income response. The comparison clarifies which part of the direct-investment-income response is tied mechanically to the exchange rate and which part reflects endogenous changes in affiliate production.

F.1. Foreign Monetary Policy Shock

We first consider a contractionary Foreign monetary policy shock. The shock raises the Foreign policy rate and, through uncovered interest parity, depreciates the Home currency while weakening Foreign demand. We normalize the shock so that the impact real depreciation is 10%, matching the baseline UIP experiment.

Figure F.1 shows that the mechanical valuation component of direct-investment income is nearly unchanged from the baseline. Because the impact depreciation is the same, the Home-currency value of existing affiliate income again rises by about 0.27 percentage points of GDP. Net direct-investment income, however, increases by only 0.34 percentage points of GDP, implying that valuation accounts for roughly 80% of the impact response. In contrast to the baseline UIP

Figure F.2: Impulse responses: UIP shock vs. Foreign monetary policy shock

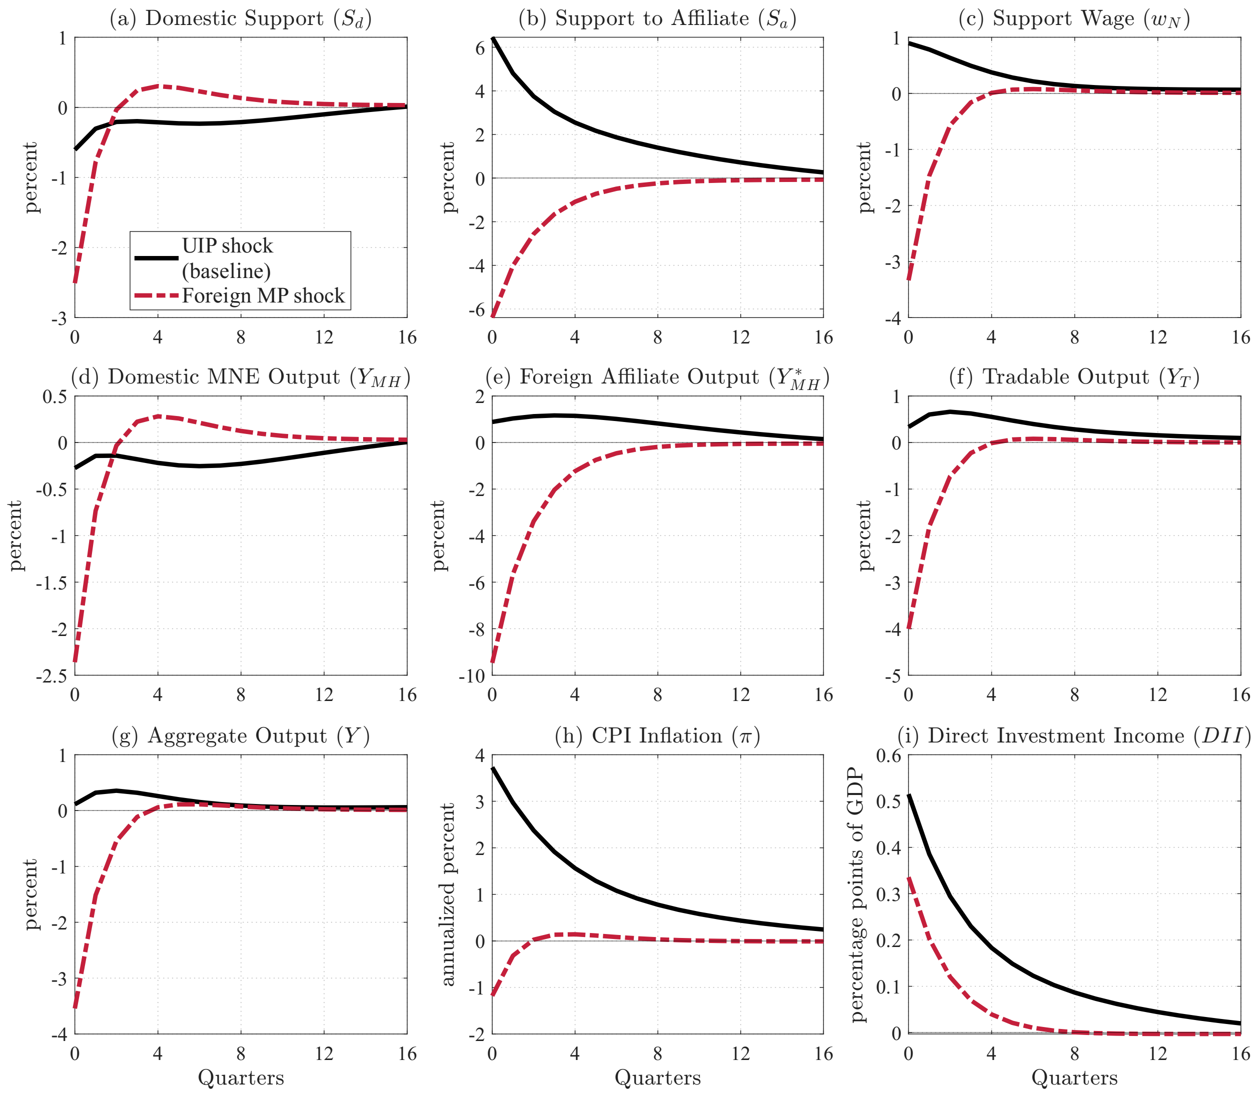

Notes: Each economy is normalized to a 10% impact real depreciation. Black solid: UIP risk-premium shock (baseline). Red dash-dotted: Foreign monetary policy shock. Panels follow the layout of Figure 7. Panel (h) reports annualized CPI inflation; panel (i) is in percentage points of steady-state GDP; the remaining panels are percent deviations from steady state.

experiment, the endogenous affiliate-income component is therefore very small.

The gross decomposition clarifies where the small residual comes from. On impact, DII receipts rise by 0.33 percentage points of GDP, of which 0.27 is pure translation; the endogenous outward-surplus component is therefore only +0.06. Affiliate revenue falls with host demand, but the affiliate's local wage bill falls even faster in the Foreign recession, so the affiliate's Foreign-currency operating surplus is approximately unchanged despite the sharp decline in output. The remaining +0.01 comes from lower payments on Foreign-owned affiliates located at Home, whose surplus also weakens. The proximity of net DII to the valuation benchmark thus reflects two small and partly offsetting gross margins, not a mechanical absence of endogenous responses.

The comparison in Figure F.2 explains why the endogenous component is so small. Unlike the UIP shock, tighter Foreign monetary policy contracts Foreign demand, causing Home-owned affiliates to sell less in their local market. Affiliate output falls sharply, and with it the incentive to reallocate scarce support toward foreign affiliates; support instead shifts back toward domestic operations. The exchange rate continues to generate a mechanical valuation gain on existing affiliate income, but the production response that amplified direct-investment income in the baseline is largely absent.

F.2. Foreign Discount-Factor Shock

We next consider a Foreign discount-factor shock that shifts Foreign expenditure toward the present and, through the Foreign Taylor rule and uncovered interest parity, depreciates the Home currency. We introduce the shock as a preference shifter in Foreign household utility. Let ζ_t^* denote the shock, with

$$\zeta_t^* = \rho_\zeta \zeta_{t-1}^* + \varepsilon_t^{\zeta^*}, \qquad \rho_\zeta = 0.9.$$

The Foreign household's marginal utility becomes

$$\lambda_t^* = \exp\left(\zeta_t^*\right)\left(X_t^*\right)^{-\sigma},$$

so the Foreign Euler equation is

$$1 = \beta \mathbb{E}_t \left[\frac{\lambda_{t+1}^*}{\lambda_t^*} \frac{R_t^*}{\pi_{t+1}^*} \right].$$

A positive realization of ζ_t^* raises the marginal utility of current consumption. Because the shock is mean-reverting, it increases current Foreign expenditure relative to future expenditure. Through the Foreign Taylor rule and uncovered interest parity, this expansion depreciates the Home currency while strengthening demand for Home-owned affiliate output. As before, the innovation is normalized to a 10% impact real depreciation.

Figure F.3 shows that the pure valuation component is again essentially unchanged, reflecting the identical impact depreciation. Net direct-investment income, however, rises to 1.31 percentage points of GDP on impact, of which only about 20% is accounted for by mechanical valuation. The remaining increase comes from higher operating surpluses of Home-owned affiliates abroad and, in somewhat larger part, from lower payments on Foreign-owned affiliates located at Home.

The comparison with the baseline UIP experiment in Figure F.4 shows the mechanism. Stronger Foreign demand raises affiliate sales and operating profits. The higher return to foreign operations induces the parent to reallocate considerably more support toward its affiliate, further expanding affiliate production and income. Relative to the UIP experiment, the exchange-rate

Figure F.3: External adjustment after a Foreign discount-factor shock

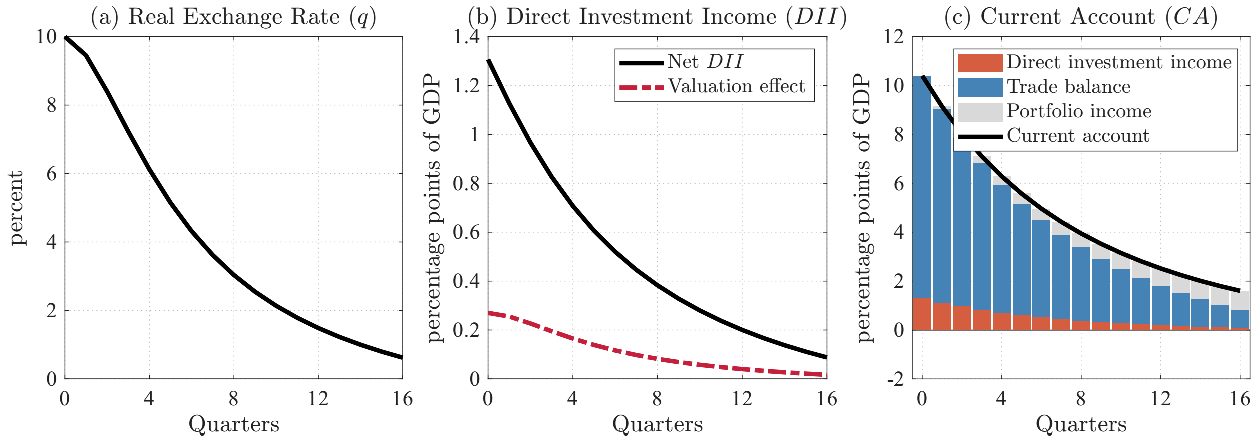

Notes: Impulse responses to a Foreign discount-factor shock, normalized to a 10% impact real depreciation. Panel (a) is in percent deviation from steady state; panels (b)–(c) are in percentage points of steady-state GDP. Panel (b) plots net direct-investment income against its pure currency-valuation component, $\xi q \bar{\tilde{\Pi}}^*_{MH} - \xi^* \bar{\tilde{\Pi}}_{MF}$, computed with the exchange rate on its dynamic path but the affiliate operating surpluses fixed at steady state. Panel (c) decomposes the current account into the trade balance, direct investment income, and portfolio income.

incentive to serve Foreign markets is reinforced rather than offset by the surrounding economic environment, so the real affiliate-income channel becomes the dominant source of direct-investment-income adjustment.

Generating a 10% impact depreciation from a discount-factor shock requires a large innovation, so the responses are correspondingly large; the informative object is the decomposition of the direct-investment-income response rather than its scale. The exercise shows how the direct-investment-income decomposition changes when the same depreciation is accompanied by stronger conditions for foreign affiliate expansion.

Taken together, the three experiments identify two distinct margins governing multinational external adjustment. The exchange rate determines the mechanical valuation gain on existing affiliate income. The real affiliate-income response, by contrast, depends on the economic forces that make foreign affiliate expansion more or less profitable. The Foreign monetary policy shock illustrates how weaker conditions for affiliate production suppress the real-income channel, whereas the Foreign discount-factor shock shows how stronger conditions amplify it.

Figure F.4: Impulse responses: UIP shock vs. Foreign discount-factor shock

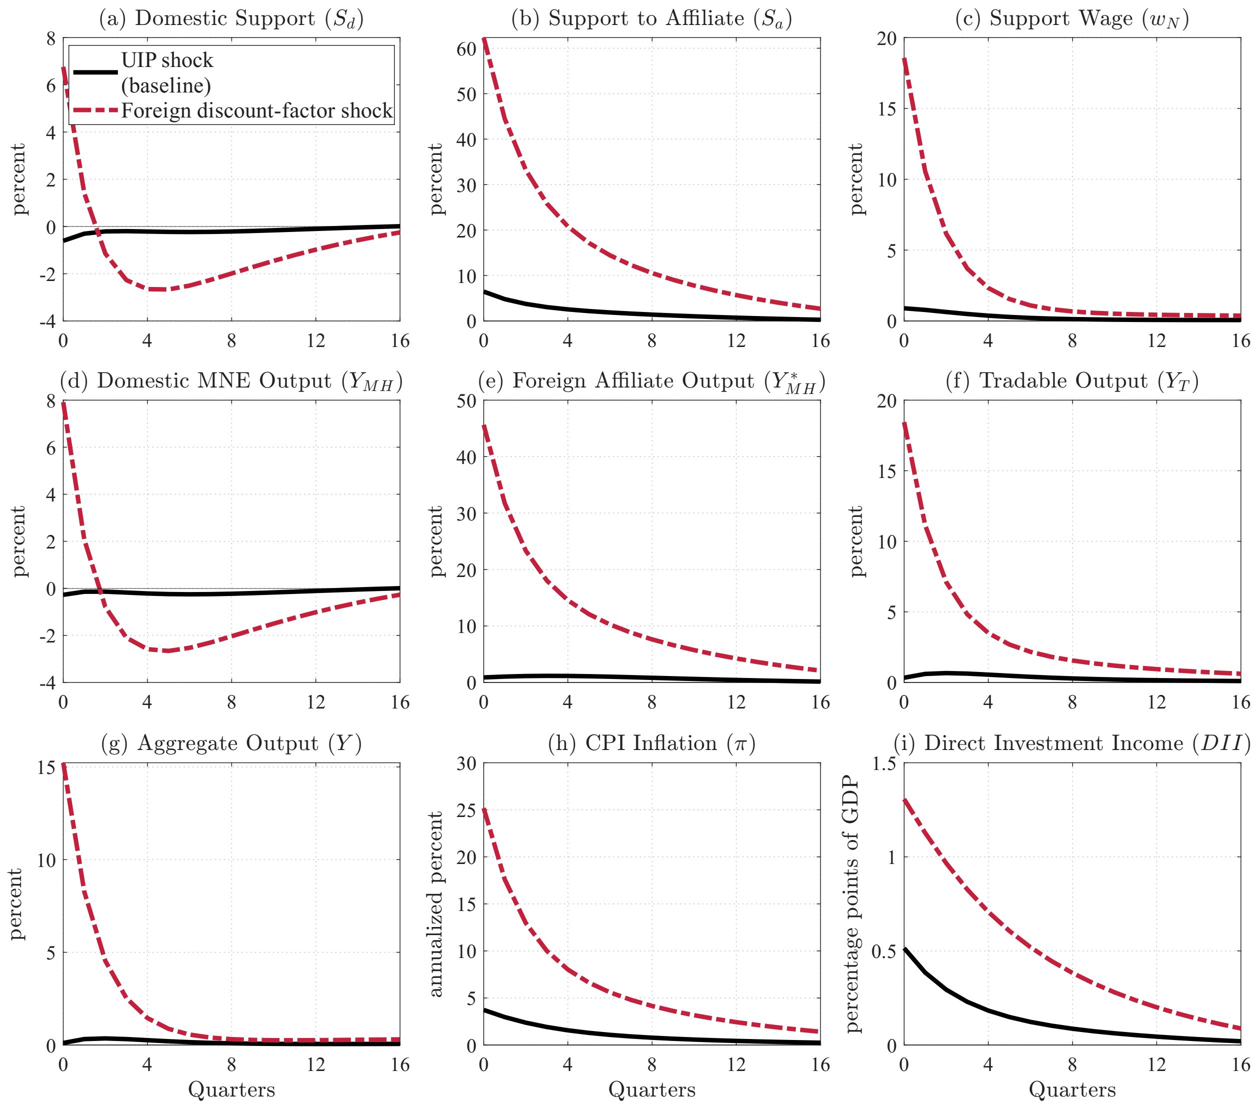

Notes: Each economy is normalized to a 10% impact real depreciation. Black solid: UIP risk-premium shock (baseline). Red dash-dotted: Foreign discount-factor shock. Panels follow the layout of Figure 7.

G Quantitative Model Robustness

G.1. Sensitivity to the Affiliate-Income Pass-Through

The baseline calibration sets the effective affiliate-income pass-through parameter to the mean ratio of affiliate-to-parent payments to affiliate current income, $\xi = 0.743$; the corresponding median ratio is 0.315. These observed payments include several types of intrafirm transfers and do not provide an exact accrual-based measure of direct investment income, so we use the ratio only as a proxy for the strength of the affiliate-to-parent income linkage. The parameter ξ determines both the steady-state direct-investment-income credits generated per unit of affiliate surplus and the weight placed on affiliate surplus in the parent's support-allocation condition.

Table G.1: Sensitivity of aggregate responses to the effective affiliate-income pass-through

ξ	ω	$\frac{A}{X+A}$	Net DII impact (ppt of GDP) Total	Valuation	Endogenous	Organizational	TB	CA	ERPT
0.743 (mean)	0.292	0.305	0.515	0.270	0.245	0.189	2.16	2.68	0.29
0.500	0.457	0.395	0.523	0.270	0.253	0.208	2.15	2.67	0.30
0.315 (median)	0.772	0.509	0.532	0.270	0.262	0.226	2.12	2.65	0.31

Notes: Impact responses to the baseline UIP shock, normalized to a 10% real depreciation, in percentage points of steady-state GDP. At each ξ, (ω, ω_T) are re-solved so that the steady state reproduces exports of 17.7% of GDP and DII credits of 2.7% of GDP; all other parameters are at baseline. $A/(X+A)$ is the affiliate share of steady-state foreign-facing sales. "Valuation" is the mechanical translation component computed with affiliate surpluses fixed at steady state; "Endogenous" is the total net DII response minus the valuation component; "Organizational" is the difference between the baseline endogenous-support model and the proportional-support counterfactual with the same steady state. TB and CA are impact responses of the trade balance and current account; ERPT is cumulative CPI exchange-rate pass-through at four quarters.

We therefore ask whether the aggregate results are sensitive to its calibration.

To answer this, we re-solve the model at $\xi \in \{0.315, 0.500, 0.743\}$. At each value, ω and ω_T are re-calibrated so that the steady state re-hits the two gross-flow targets, exports of 17.7% of GDP and direct-investment-income credits of 2.7% of GDP; all other parameters remain at baseline. Moving from the mean to the median payment share raises steady-state affiliate sales from 7.8% to 18.3% of GDP and the affiliate share of foreign-facing sales from 0.30 to 0.51.

Table G.1 reports the results. Because the credits target is re-imposed at every ξ, the mechanical valuation component of the impact response is 0.27 percentage points of GDP throughout. The total net direct-investment-income response is essentially invariant, rising slightly from 0.51 to 0.53 percentage points of GDP as ξ falls from the mean to the median. The endogenous component rises from 0.24 to 0.26, and the organizational component—the part attributable to the endogenous reallocation of rival support, measured against the proportional-support counterfactual—rises from 0.19 to 0.23. The trade balance, current account, and exchange-rate pass-through are essentially unchanged, as is the impact response of real GDP (about 0.11% at every ξ).

The invariance reflects two offsetting forces. A lower ξ reduces the Home parent's effective income claim per unit of affiliate surplus, but holding measured direct-investment-income credits fixed then requires a larger affiliate base, which strengthens the production side of the mechanism. The observed gross-flow targets discipline the product of the effective affiliate-income share and affiliate surplus, and the aggregate responses depend primarily on that product. Calibrating ξ to the mean payment share is therefore not driving the aggregate results; if anything, it is mildly conservative for the organizational channel.

G.2. National-Accounting Treatment of Parent Support Deployed Abroad

Our baseline model assumes the transfer price of parent support is normalized to zero. Support deployed to the foreign affiliate, $S_{a,t}$, is an internal, nontraded intra-firm input: the affiliate is charged nothing for it, so its productive contribution is recorded inside affiliate operating surplus and therefore in *Foreign* GDP, and it reaches Home only as the ξ share of that surplus, through direct investment income and hence through gross national income. At Home the support wage bill $w_{N,t}S_{a,t}$ is paid to Home households but is exactly offset by lower parent operating surplus, so it nets to zero in Home GDP. This section asks how much of the paper's ownership–location result depends on that normalization.

An alternative convention. Index the convention by an imputed transfer price $\tau_t = \chi\, w_{N,t}$ per unit of support. Measured Home service exports and imports are

$$X_t^S = \chi\, w_{N,t} S_{a,t}, \qquad M_t^S = \chi\, q_t\, w_{N,t}^* S_{a,t}^*, \tag{G.1}$$

and the accounts become $gdp_t^\chi = gdp_t + X_t^S - M_t^S$, $tb_t^\chi = tb_t + X_t^S - M_t^S$, and $dii_t^\chi = dii_t - \xi X_t^S + \xi^* M_t^S$. We compare the baseline, $\chi = 0$, with cost-basis imputation, $\chi = 1$, which treats support as a service exported to the affiliate and values it at the support wage the parent actually pays.

The fee in Equation (G.1) is evaluated at *aggregate* $w_{N,t}$ and $S_{a,t}$, which the atomistic parent takes as given. It is therefore lump sum from the individual parent's viewpoint and enters no first-order condition, so the real allocation is invariant to χ. We verify this numerically rather than assert it: re-solving the model at each χ leaves every impulse response of q, C, Y_{MH}, Y_{MH}^*, S_a, S_d, N, w_N, w_L, inflation, and the policy rate unchanged to 2.7×10^{-15}. A literal *marginal* transfer price would not be innocuous: the parent would capture the full fee but only ξ of the surplus it displaces, which for $\xi < 1$ drives a wedge into the allocation condition and induces over-deployment of support abroad. That case is a change in the mechanism, not in the accounting, and we do not impose it.

An implicit intra-firm transfer. Under the re-recording, Home GDP and the trade balance each rise by $X_t^S - M_t^S$, while net direct investment income falls by only $\xi X_t^S - \xi^* M_t^S$, because just the ξ share of the displaced surplus was accruing to Home. The residual $(1-\xi)X_t^S - (1-\xi^*)M_t^S$ is a genuine unrequited transfer: the zero-transfer-price normalization implicitly hands the non-repatriated share of the support Home supplies to foreign residents, worth 0.31% of Home GDP each period in the symmetric steady state. A balance-of-payments-consistent recording of the imputation therefore has three entries—a service export, a reduction in the direct-investment-income credit, and a secondary-income debit—and with the transfer entry included the current account and gross national disposable income are invariant to χ exactly in the model.

Quantitative results. Figure G.1 shows the responses. On impact, the measured fixed-price real GDP response rises from 0.11% under the baseline to 0.27% under cost-basis imputation, while

Figure G.1: External adjustment under alternative accounting conventions

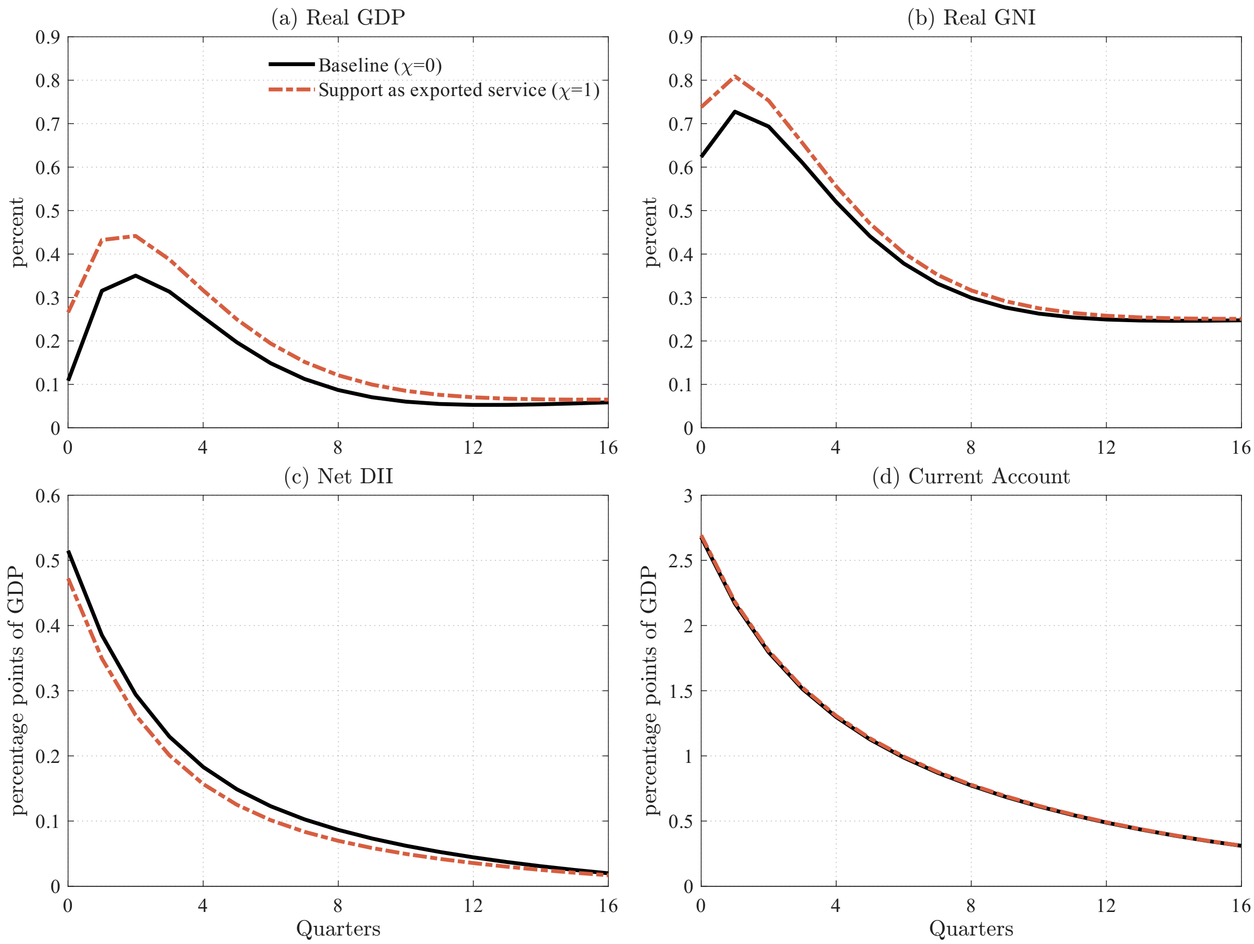

Notes: Impulse responses to a UIP risk-premium shock normalized to a 10% impact real depreciation, under two transfer-price conventions for parent support deployed abroad: $\chi = 0$ (baseline, support expensed at Home and not recorded as a service export) and $\chi = 1$ (support recorded as a service export at the support wage). Real GDP and real GNI are fixed-price quantity indices; direct investment income and the current account are in percentage points of steady-state GDP, and $\Delta RGNI^F = \Delta RGDP^F + \Delta DII + \Delta PI$ as in Section 6.2. Panel (d) plots the current account without the secondary-income entry; with that entry recorded the two paths coincide exactly. The underlying real allocation is identical across conventions.

net direct investment income falls from 0.51 to 0.47 percentage points of GDP. The two move in opposite directions by construction: the convention reclassifies the same activity between domestic production and foreign income. The current account is essentially unaffected on impact (2.68 against 2.69 percentage points), and exactly unaffected once the secondary-income entry is recorded. The real content of the mechanism is convention-free—support reallocates abroad, affiliate output expands, and domestic parent output contracts by identical amounts under either convention—while the measured production–income split moves: gross national income responds 5.8 times more than domestic production under the baseline and 2.8 times under cost-basis imputation.

Figure G.2: Measured real GDP across foreign-market organizations, by accounting convention

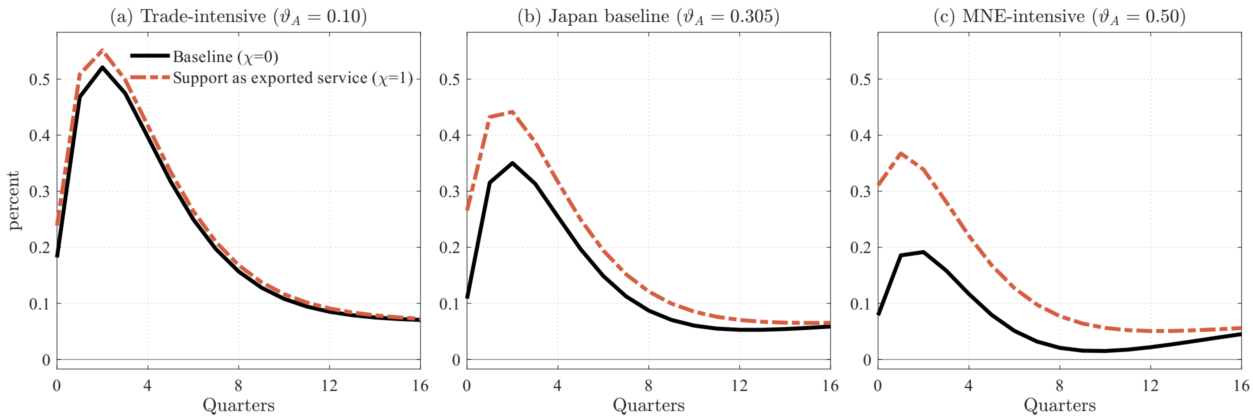

Notes: Fixed-price real GDP responses to a UIP risk-premium shock normalized to a 10% impact real depreciation, for the three foreign-market-organization economies of Section 6.4, under the baseline convention ($\chi = 0$, solid) and with support recorded as an exported service ($\chi = 1$, dashed). Each economy is re-solved at its own counterfactual calibration, and the allocation-invariance check is repeated economy by economy. The vertical scale is common across panels.

Interaction with foreign-market organization. The size of the reclassification scales with the support actually deployed abroad, so the convention matters most exactly where affiliate production matters most. Figure G.2 repeats the exercise for the three foreign-market-organization economies of Section 6.4, re-solving the model at each counterfactual calibration. Under the baseline convention the measured real-GDP response declines with affiliate intensity: impacts of 0.18, 0.11, and 0.08 percent in the trade-intensive economy, the Japan baseline, and the MNE-intensive economy, with peaks of 0.52, 0.35, and 0.19. Recording support as an exported service raises all three responses and reverses the impact ordering—0.24, 0.27, and 0.31—because the imputed service export scales with the support reallocation, which is largest where the affiliate sector is largest. The reversal is confined to short horizons: the peak responses under $\chi = 1$ are 0.55, 0.44, and 0.37, so the trade-intensive economy again shows the strongest measured response, because its expenditure-switching gains build over several quarters while the imputed service export mean-reverts with the exchange rate.

The counterfactual conclusion should therefore be stated in convention-robust form: shifting foreign-market activity from exports toward affiliate production weakens the domestic goods-production response and shifts external adjustment toward income generated abroad. Recording parent support as an exported service raises measured Home GDP, particularly at short horizons, but does not eliminate the distinction between production for export at Home and production undertaken by foreign affiliates.

G.3. Asymmetric Inward and Outward Multinational Positions

The benchmark calibration is symmetric across countries, so the inward multinational block—Foreign-owned affiliates producing in Home, with output Y_{MF}—is the same size as the outward block Y^*_{MH}, and steady-state net direct investment income is zero. Japan is the most asymmetric of the G7 economies on this dimension. In 2019 its direct-investment-income credits were 2.7% of GDP against debits of 0.66%, a debit-to-credit ratio of 0.245, reflecting an inward direct-investment position far smaller than the outward one. Because the headline objects in Section 6.2 are the *net* direct-investment-income and current-account responses, and because both are differences between a credit and a debit leg, the symmetric benchmark abstracts from the asymmetry that shapes them. This appendix therefore re-solves the model with a Japan-consistent inward block.

An asymmetric steady state. Dropping symmetry costs the closed form: with $\omega_M \neq \omega^*_M$, the real exchange rate, the trade balance, and net direct investment income are no longer pinned at 1, 0, and 0 by construction, and the steady state must be solved numerically. We hold preferences fixed—the scale parameters χ_L and χ_N keep their baseline values, so Home and Foreign share preferences and technology, and the four labor-supply conditions become genuine equilibrium restrictions rather than definitions—and solve for w_L, w^*_L, q, C, and C^*. Four demand weights are then chosen so that the asymmetric steady state reproduces four moments: ω_M delivers the Japanese debit-to-credit ratio of 0.245, ω^*_M holds direct-investment-income credits at 2.7% of GDP, ω^*_T holds exports at 17.7% of GDP, and ω_T normalizes $\bar{q} = 1$. Holding the credits and export targets at their baseline values, rather than re-reading them from the data, is what makes the exercise informative: the mechanical valuation benchmark $\xi \bar{q} \bar{\bar{\Pi}}^*_{MH}$ is then numerically identical across the two calibrations, so any difference in the net response is attributable to the debit side and to the endogenous credit response rather than to a rescaled steady-state level. All remaining parameters, including ξ and the micro-calibrated $(\alpha_S, \varepsilon_N)$, are unchanged.

The asymmetry has to be absorbed somewhere in the steady-state external accounts. With $\bar{b}_F = 0$, net-foreign-asset stationarity requires $\overline{TB} = -\overline{DII}$: the steady-state direct-investment-income surplus of 2.0% of GDP finances a trade deficit of the same size. Japan's actual configuration—a current-account surplus of about 3.5% of GDP alongside a positive net foreign asset position—cannot be a stationary steady state of a model without growth, so some closure must be imposed. We choose $\bar{b}_F = 0$ because it keeps the portfolio position, and hence any portfolio valuation gain, out of the exchange-rate experiment, leaving the direct-investment-income decomposition clean. The alternative closure, matching Japan's near-zero trade balance by giving Home a large net portfolio debt, would introduce a second valuation channel and confound the decomposition.

Two features of the resulting steady state are worth noting because they bound how much else the exercise changes. First, the outward block is unchanged: affiliate sales abroad remain 7.8% of

Table G.2: Symmetric benchmark versus a Japan-consistent asymmetric calibration

	Symmetric benchmark	Japan-consistent asymmetric
Steady state		
DII credits (% of GDP)	2.70	2.70
DII debits (% of GDP)	2.70	0.66
Inward/outward affiliate sales	1.00	0.25
Net DII (% of GDP)	0.00	2.04
Trade balance (% of GDP)	0.00	−2.04
Impact response to a 10% *depreciation (ppt of GDP)*		
Gross DII credits	0.392	0.397
Gross DII debits	−0.122	−0.032
Net DII	0.515	0.429
Valuation component	0.270	0.270
Endogenous component	0.245	0.159
of which organizational	0.189	0.126
Trade balance	2.163	2.227
Current account	2.678	2.656
Real GDP (percent)	0.108	0.031
First year, cumulative (ppt of annual GDP)		
Net DII	0.356	0.292
Valuation component	0.178	0.178
Current account	2.039	2.020

Notes: Responses to the baseline UIP shock, normalized to a 10% impact real depreciation. In the asymmetric calibration $(\omega_M, \omega_M^*, \omega_T, \omega_T^*) = (0.935, 0.718, 0.726, 0.751)$ against $(0.723, 0.723, 0.750, 0.750)$ in the benchmark; credits, exports, and $\bar{q}$ are held at their benchmark values, so only the inward/outward asymmetry changes. "Organizational" is the difference between the endogenous-support model and the proportional-support counterfactual with the same steady state. Impact rows are quarterly flows relative to quarterly steady-state GDP; the final block is the first-year cumulative flow relative to annual GDP.

Home GDP, while inward affiliate sales fall from 7.8% to 1.9%. Second, the two countries remain almost exactly the same size—Home GDP falls by 1.0% and Foreign GDP rises by 1.2% relative to the symmetric benchmark—so the comparison does not confound asymmetric multinational positions with a change in relative country size.

The credit side is unchanged; the net response falls. Table G.2 and Figure G.3 report the results. The gross credit response is essentially identical across the two calibrations, 0.392 against 0.397 percentage points of GDP on impact, and the two paths are visually indistinguishable in panel (b) of Figure G.3. This is the central diagnostic: the mechanism of the paper, and the micro evidence that disciplines it, both live on the outward block, whose steady-state size and whose response to the depreciation are unaffected by how large the inward block is.

The net response nonetheless falls, from 0.515 to 0.429 percentage points of GDP on impact,

Figure G.3: Direct investment income under symmetric and asymmetric multinational positions

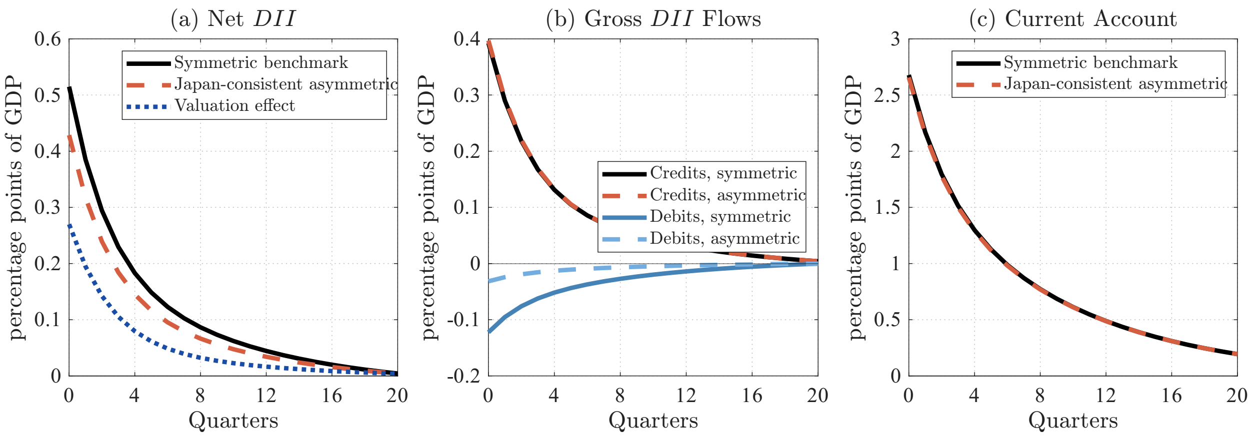

Notes: Impulse responses to the baseline UIP shock, normalized to a 10% impact real depreciation, in percentage points of steady-state GDP. Black solid: symmetric benchmark. Red dashed: Japan-consistent asymmetric calibration. Panel (a) plots net direct investment income against the common pure currency-valuation component. Panel (b) plots the gross credit and debit legs separately; the two credit paths coincide. Panel (c) plots the current account.

a reduction of about one sixth. The entire difference comes from the debit leg. In the symmetric benchmark, a Home depreciation lowers payments to Foreign parents by 0.122 percentage points of GDP—the mirror image of the paper's own mechanism, with the Foreign parent withdrawing support from its now less profitable Home affiliate—and because these payments are a debit, their decline *adds* to net direct investment income. Shrinking the inward block to Japan's actual size shrinks that contribution to 0.032. The symmetric benchmark is therefore not a conservative lower bound for the net response: it modestly *overstates* it, by crediting Japan with an inward block roughly four times larger than the data support.

The current account and the mechanism decomposition are robust. Two further results matter for the paper's claims. First, the current-account response is almost unchanged, 2.678 against 2.656 percentage points of GDP on impact, because the smaller direct-investment-income response is offset by a slightly larger trade-balance response: with a smaller inward block, more of Home's multinational demand is met by Home parents, and the expenditure-switching response of the trade balance is correspondingly stronger. The headline external-adjustment result is thus not an artifact of symmetry.

Second, the organizational decomposition survives. The endogenous component of the net direct-investment-income response falls from 0.245 to 0.159 percentage points of GDP, and the part attributable to the endogenous reallocation of rival support—measured against the proportional-support counterfactual—falls from 0.189 to 0.126. As a share of the endogenous component, however, the organizational channel is essentially unchanged, 77% against 80%. The claim that the endogenous response is driven by the reallocation of scarce parent support, rather than by a scale effect, is therefore invariant to the asymmetry.

Figure G.4: Impulse responses under symmetric and asymmetric multinational positions

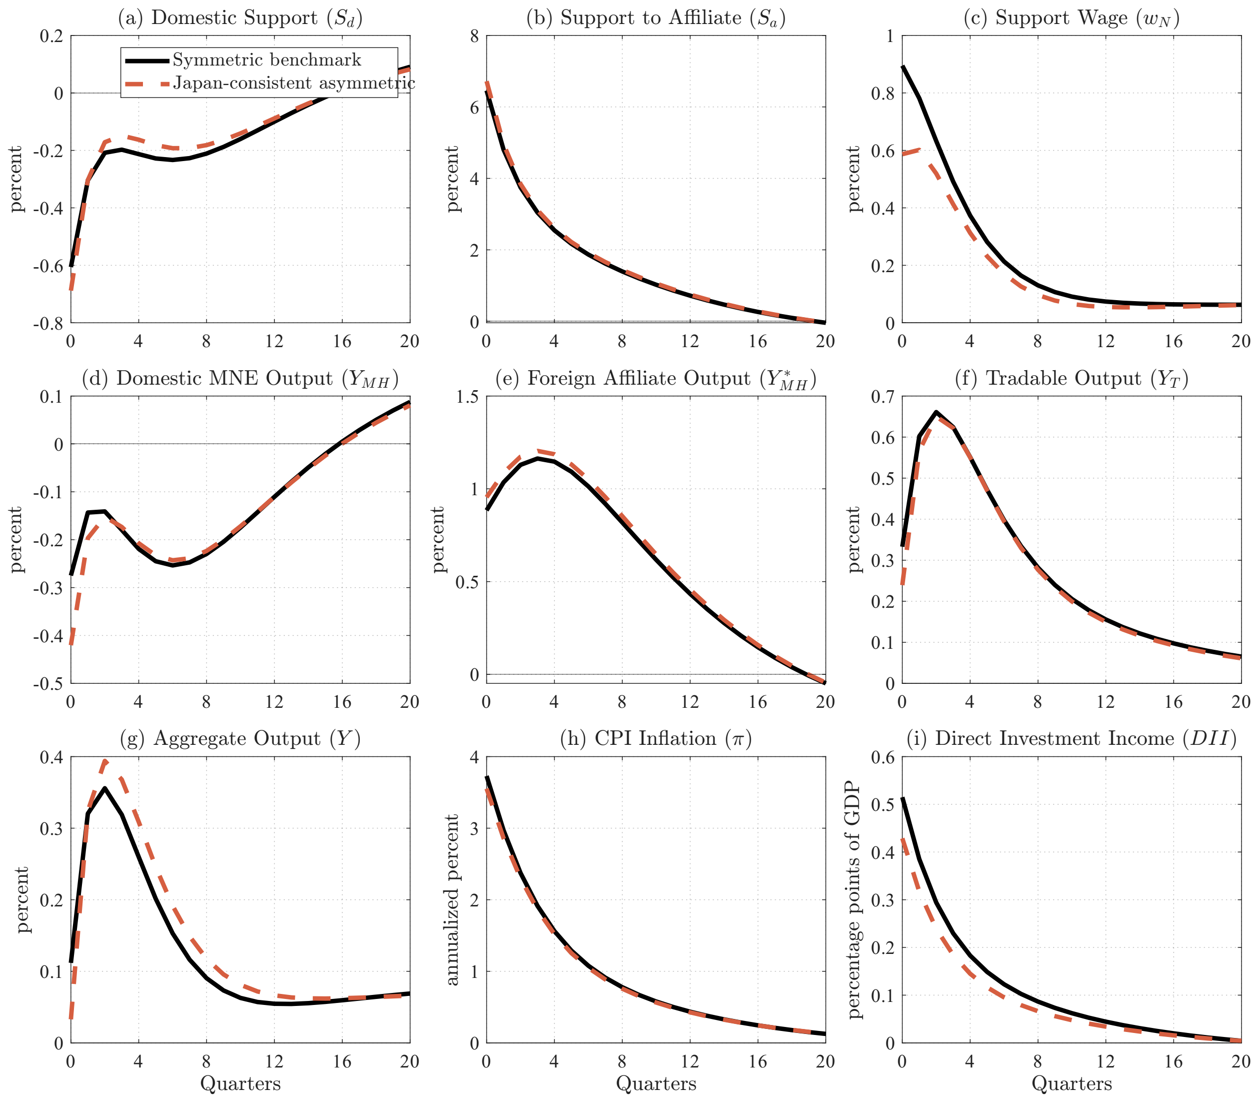

Notes: Each economy is normalized to a 10% impact real depreciation. Black solid: symmetric benchmark. Red dashed: Japan-consistent asymmetric calibration. Panels follow the layout of Figure 7. Panel (h) reports annualized CPI inflation; panel (i) is in percentage points of steady-state GDP; the remaining panels are percent deviations from steady state.

Not everything is invariant. The impact response of real GDP falls from 0.108% to 0.031%, because in the symmetric benchmark a substantial part of the Home real-activity response comes from Foreign-owned affiliates located in Home expanding after the depreciation, and that block is now small. The full set of impulse responses is compared in Figure G.4: the support-allocation, affiliate-output, tradable-output, and inflation responses are nearly identical across the two calibrations, and the differences are concentrated in the Home-located multinational aggregates, whose steady-state composition is what the asymmetry changes. Since the paper's quantitative claims concern the external accounts rather than the level of the domestic output response, we read this as a caveat on the real-activity numbers rather than on the mechanism.

Implication for the Japanese benchmark. The comparison with the 2022 episode in Section 6.2 is affected in the expected direction. Under the asymmetric calibration the model's first-year net direct-investment-income response is 0.292 percentage points of annual GDP per 10% depreciation, against 0.356 in the benchmark, while the valuation component is unchanged at 0.178. Scaling proportionally to a 25–30% depreciation gives 0.7–0.9 percentage points of annual GDP rather than 0.9–1.1, so the model accounts for roughly 40% of the observed increase rather than about half. The qualitative reading is unchanged—the response is about 1.6 times the valuation-only path, and the endogenous margin remains essential to explaining the size of the observed move—but the asymmetric calibration is the more appropriate basis for a Japan-specific magnitude, and it is the more conservative one.

Appendix References

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M. Taylor, "A Local Projections Approach to Difference-in-Differences Event Studies," Working Paper 31184, National Bureau of Economic Research 2023.

Giroud, Xavier and Holger M. Mueller, "Firm Leverage, Consumer Demand, and Employment Losses During the Great Recession," *The Quarterly Journal of Economics*, 2017, *132* (1), 271–316.

Schmitt-Grohé, Stephanie and Martín Uribe, "Closing Small Open Economy Models," *Journal of International Economics*, 2003, *61* (1), 163–185.