Sovereign Default and the Scarring of Export Relationships*

Diogo Sampaio Lima
Maastricht University

Last updated: March 26, 2026

Abstract

[TODO: Abstract — write after introduction and results are stable.]

JEL classification: F14, F34, F41.

Keywords: sovereign default, export relationships, extensive margin, trade scarring, search frictions

*Email address: d.sampaiolima@maastrichtuniversity.nl.

1 Introduction

Sovereign defaults are costly events. Output falls, consumption contracts, and access to international capital markets is lost for years. Yet the quantitative sovereign debt literature, following Arellano (2008) and Aguiar and Gopinath (2006), has largely treated these costs as a reduced-form penalty — an exogenous drop in output or exclusion from credit markets that deters default in equilibrium. This approach sidesteps a harder question: *what, specifically, does default break in the real economy?* If the answer is "export relationships," then the costs of default are not simply a contemporaneous punishment but a persistent scar on the economy's productive capacity — one that feeds back into sovereign borrowing conditions long after the debt crisis has been resolved.

This paper provides evidence that sovereign default persistently damages the network of export relationships, and develops a quantitative model in which this damage is the primary channel through which default is costly. On the empirical side, I use product-level bilateral trade data from BACI-CEPII, covering over 5,000 HS6 product categories across all trading partners, to study the evolution of exports around five post-default sovereign restructuring episodes identified in Asonuma and Trebesch (2016). On the theoretical side, I build a sovereign debt model in the tradition of Chatterjee and Eyigungor (2012) in which the real economy features a microfounded trade block with heterogeneous firms, search-and-matching frictions for export relationships, relationship-specific capital, and dormant links that can be reactivated. Default damages the trade network through four economically interpretable wedges — profitability, working capital, new-match efficiency, and reactivation efficiency — and the slow rebuilding of destroyed relationships generates an endogenous, persistent output cost.

The empirical analysis documents four facts. First, export product variety — the count of distinct products a country exports — falls sharply at the time of default and remains depressed a decade later, even as total export value largely recovers. Local projection estimates show a decline of approximately 12 log points on impact, deepening to 17 log points at the ten-year horizon. Second, this extensive-margin scar operates primarily through products and product-destination pairs rather than through loss of destination markets: countries continue exporting to most of their pre-crisis partners, but the density of products shipped to each destination thins persistently. Third, pre-crisis export relationships — product-destination pairs that were active before default — survive at significantly lower rates than comparable relationships in non-default control episodes. Fourth, relationships that are lost re-enter more slowly after default: cumulative re-entry rates are persistently depressed relative to controls. Together, these facts point to a mechanism rooted in the destruction and slow rebuilding of relationship-specific assets, rather than a broad loss of market access.

The model generates these patterns endogenously. Default triggers an increase in trade costs and a tightening of trade finance, which raises the productivity threshold for profitable exporting and pushes marginal relationships into dormancy. Simultaneously, matching frictions for both new and reactivated relationships worsen. The combined effect is a sharp fall in the stock of active export relationships and a sluggish recovery: dormant links lose relationship capital over time, new matches are structurally harder to form than reactivations, and reactivations themselves are impeded. The resulting decline in export revenue depresses the sovereign's future output, tightening borrowing constraints and raising spreads — creating a two-way feedback between trade-network damage and sovereign debt sustainability.

The central insight is that persistence arises not from the default wedges themselves — which normalise once the country regains market access — but from the slow-moving dynamics of the relationship stocks. Export relationships embody sunk investments: product adaptation, distribution networks, buyer-specific certifications, and accumulated trust. These assets are costly to build, depreciate without use, and cannot be rebuilt overnight. A sovereign that has defaulted therefore faces a damaged export network that depresses output for years, worsening the very borrowing conditions that led to default. The state variables tracking active and dormant relationships are the channels through which history enters the sovereign's problem, generating a form of hysteresis that is absent from models with reduced-form output costs.

This paper connects three literatures. First, it contributes to the quantitative sovereign debt literature (Arellano, 2008; Aguiar and Gopinath, 2006; Chatterjee and Eyigungor, 2012) by replacing the standard exogenous output cost with a microfounded trade channel. Mendoza and Yue (2012) is the closest predecessor: they model default costs through disruption of imported intermediate inputs, generating an endogenous output decline. My paper differs in focusing on the *export* side and, specifically, on the extensive margin of trade — the destruction and slow rebuilding of individual product-destination relationships rather than aggregate trade volumes. Second, the paper builds on the empirical literature on trade and sovereign debt. Rose (2005) documented that bilateral trade falls after sovereign default; Martinez and Sandleris (2011) found more muted effects. I revisit this question with product-level data and a local projection design that separates the extensive margin (variety) from the intensive margin (volume per product), revealing that the persistent scar is concentrated in the extensive margin. Third, the trade block draws on the heterogeneous-firm trade literature originating with Melitz (2003) and extended to relationship dynamics by Eaton et al. (2011) and Fitzgerald et al. (2024). The innovation is to embed a relationship-based trade block — with search frictions, dormancy, reactivation, and learning-by-doing in relationship capital — inside a sovereign default model, allowing the two-way feedback between default decisions and trade-network dynamics.

The remainder of this paper is organized as follows. Section 2 discusses related literature in detail. Section 3 documents the empirical evidence on export scarring after sovereign default. Section 4 presents the model. Section 5 describes the calibration. Section 6 presents the main results. Section 7 concludes.

2 Related Literature

[TODO: Literature review - pull from notes/literature/summaries.tex]

3 Empirical Evidence: Export Scars from Sovereign Default

This section documents four stylized facts about the trade consequences of sovereign default. The evidence draws on product-level bilateral trade data from BACI-CEPII and five post-default restructuring episodes identified in Asonuma and Trebesch (2016). The analysis proceeds by zooming in from the aggregate to the micro level: we first establish that export variety falls persistently after default while volume recovers (Section 3.2), then ask which margin is scarred (Section 3.3), examine what happens to individual trade relationships (Section 3.4), characterize who drives the recovery (Section 3.6), and assess interpretation and robustness (Section 3.7). Each fact maps to a feature of the model developed in Section 4.

3.1 Data and Episode Construction

Trade data. We use BACI-CEPII HS92 (version V202601), a harmonized version of UN Comtrade that reconciles exporter and importer reports to improve accuracy of bilateral trade flows.[1] The dataset covers 1995–2024; our estimation sample uses data through 2022. The data are at the exporter $\times$ importer $\times$ HS6 product $\times$ year level. We use the HS92 classification throughout, which provides a single consistent product nomenclature across all years and avoids concordance problems across HS revisions.

Outcome variables. Our primary outcome is *export variety*: the count of distinct HS6 product codes exported by country i in year t with annual value exceeding USD 10,000, excluding HS chapters 26–27 (fuels and minerals).[2] Our secondary outcome is *export volume*: the log of total export value in USD. The comparison between variety and volume reveals the share of the trade decline attributable to the extensive margin.

Default episodes. We identify post-default sovereign restructuring episodes from the Asonuma and Trebesch (2016) chronology (2020 update) that satisfy three criteria: (i) classified as post-default (the sovereign missed payments before restructuring), (ii) at least four pre-default years available in BACI, and (iii) at least five post-restructuring years available. We further exclude episodes with commodity-dominated export baskets (Russia 1998, Venezuela 2017) or confounding armed conflict (Ukraine 2015). This yields five episodes in four countries:

Ecuador enters the sample twice (1999 and 2008) as separate episodes sharing the same country fixed effect. The post-restructuring window of the first episode overlaps the pre-default window of the second; results are robust to dropping either episode.

[1] See Gaulier and Zignago (2010). BACI's mirror-flow reconciliation is particularly valuable for measuring the extensive margin, where detecting zero flows requires reliable bilateral coverage.

[2] The USD 10,000 threshold eliminates noise from single shipments and Comtrade rounding artifacts. HS chapters 26 (ores, slag, ash) and 27 (mineral fuels, oils, waxes) are excluded because entry and exit in commodity spot markets reflects price movements and mining operations rather than the sunk-cost mechanism central to the model. Appendix D reports robustness to including fuels and to a higher threshold of USD 50,000.

Table 1: Post-default restructuring episodes. Episodes are from the Asonuma and Trebesch (2016) chronology (2020 update). Pre-default and post-restructuring years refer to data availability in BACI-CEPII HS92.

Country	Default start	Restruct. end	Pre-default yrs	Post-restruct. yrs
Ecuador	Jan 1999	Aug 2000	4	22
Argentina	Dec 2001	Jun 2005	6	17
Dominican Rep.	Feb 2005	May 2005	9	17
Ecuador	Dec 2008	Jun 2009	9	13
Belize	Aug 2012	Mar 2013	16	9

Control group. The panel includes 15 non-defaulting emerging-market countries that provide variation for year fixed effects.[3] The full panel has 19 country clusters.

Table 2 reports descriptive statistics for the treated and control groups. Treated countries are smaller on average, with lower variety counts and fewer product-destination pairs, reflecting the sample composition (Ecuador, Dominican Republic, and Belize are small open economies).

Table 2: **Descriptive statistics**

Variable	Group	Mean	Median	SD	Min	Max
Variety (count)	Treated	1,801	1,791	1,026	296	3,665
	Control	3,714	3,984	708	1,654	4,701
Log variety	Treated	7.272	7.491	0.747	5.690	8.207
	Control	8.199	8.290	0.214	7.411	8.456
Exports (000 USD)	Treated	17,132,583	8,645,865	21,872,524	195,627	84,698,313
	Control	114,719,401	80,361,159	105,120,851	4,716,291	605,146,416
Log export value	Treated	15.598	15.973	1.764	12.184	18.255
	Control	18.086	18.202	1.066	15.367	20.221
Product-dest. pairs	Treated	8,424	6,252	8,253	454	29,565
	Control	53,482	52,357	38,242	5,095	175,499

Treated: 120 country-years (4 countries). Control: 445 country-years (15 countries).
Variety = count of distinct HS6 products with annual exports > USD 10,000, excl. HS 26–27.
Panel covers 1995–2022. Source: BACI-CEPII HS92 (V202601).

We emphasize that BACI records bilateral product-level trade flows, not input-output linkages or production networks. The extensive margin measured here—distinct HS6 export products—captures the breadth of a country's export portfolio, not the structure of its domestic supply chains.

[3] The control group includes Brazil, Colombia, Peru, Mexico, Thailand, Malaysia, Philippines, South Africa, Turkey, Morocco, India, Indonesia, and others. Preemptive restructurings (Greece 2012, Jamaica 2010 and 2013) are excluded from the main sample and used as a placebo check.

3.2 Baseline Fact: Export Breadth Weakens Persistently After Default

We estimate the dynamic response of export variety and volume to sovereign default using local projections (Jordà, 2005). For each horizon $h = -5, -4, \ldots, +10$, we estimate:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i^h + \gamma_t^h + \beta_h D_{it} + \varepsilon_{it}^h \tag{1}$$

where y_{it} is either log variety or log volume for country i in year t, D_{it} is an indicator equal to one in the year country i enters a post-default restructuring, α_i^h are country fixed effects, and γ_t^h are year fixed effects. The coefficients $\{\beta_h\}$ trace out the impulse response of the outcome variable relative to the pre-default year $t-1$.

The specification is deliberately parsimonious. Country fixed effects absorb time-invariant differences in export structure, and year fixed effects absorb global trade shocks including the 2008–09 Great Trade Collapse. We omit macroeconomic controls (GDP growth, GDP per capita) from the baseline to avoid absorbing the very deterioration that default episodes produce; results are similar when these controls are included.

Standard errors are clustered at the country level. Given the small number of treated countries, we report 90% confidence intervals constructed from wild bootstrap p-values (Rademacher weights, 999 repetitions), which is standard practice for inference with few clusters.[4]

Figure 1 presents the central empirical finding. Panel A shows that export variety falls by approximately 12 log points at impact ($\hat{\beta}_0 = -0.119$) and remains depressed by 17 log points a decade later ($\hat{\beta}_{10} = -0.172$). Total export value also falls initially ($\hat{\beta}_0 = -0.097$) but largely recovers by $h = 10$ ($\hat{\beta}_{10} = -0.069$). The asymmetry between the two series is the key fact: the extensive margin is persistently scarred while the intensive margin rebounds.

Confidence intervals are wide, which is expected with five treated episodes and 19 clusters. However, the sign and magnitude of the variety response are consistent across horizons, and the pattern is reinforced by convergent evidence from the Argentina case study and the relationship-level analysis below. The pre-default coefficients ($h < 0$) are statistically indistinguishable from zero for both series, providing no evidence of differential pre-trends.

Panel B illustrates the same pattern for Argentina alone. Following the December 2001 default and June 2005 first-round restructuring, export variety remained depressed well into the 2010s, while total export value recovered to trend substantially earlier. The Argentine devaluation of January 2002—which mechanically boosted export competitiveness—works *against* finding a variety drop, making this estimate conservative.[5]

[4] See Cameron et al. (2008) on the small-cluster problem. With 19 clusters (4 treated countries, ∼15 controls), conventional cluster-robust standard errors may be unreliable. Wild bootstrap provides more accurate size control.

[5] A sharp devaluation should increase the profitability of exporting and encourage entry into export markets, boosting the extensive margin. If variety falls despite this tailwind, the true structural effect of default is larger than what the data show.

Figure 1: **Export variety collapses and recovers slowly after sovereign default.** *Panel A* shows local-projection impulse responses of log export variety (blue, solid) and log export volume (red, dashed) around post-default restructurings, estimated on a panel of five episodes with country and year fixed effects. Shaded bands are 90% confidence intervals from wild-bootstrapped standard errors clustered at the country level. Both series are normalized to zero at $h = -1$. *Panel B* shows the raw normalized series for Argentina (2001 default, 2005 first-round restructuring). Source: BACI-CEPII HS92 (1995–2022); Asonuma and Trebesch (2016) default chronology.

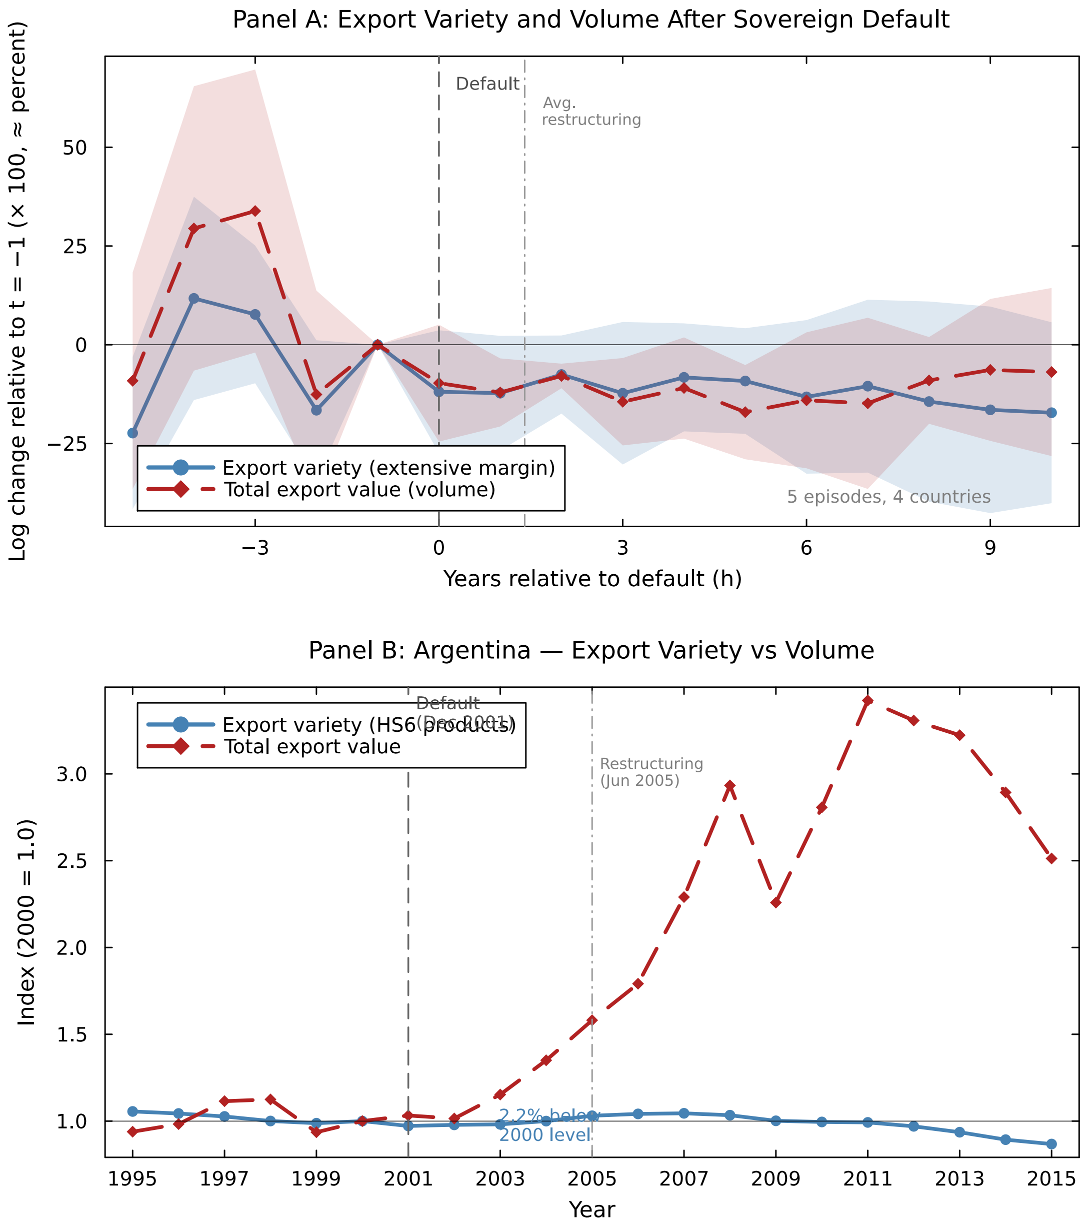

3.3 Which Margin Is Scarred?

The aggregate variety result raises a natural question: which component of the extensive margin is actually moving? We decompose the trade response into four margins—products, destinations, product-destination pairs, and total export value—and find that the scar operates primarily through products and product-destination pairs, not through broad destination loss.

Figure 2: **Extensive-margin decomposition.** Event-study paths for products, destinations, product-destination pairs, and total export value around post-default restructurings. The strongest and most persistent declines are in products and product-destination pairs. Destination reach is relatively stable.

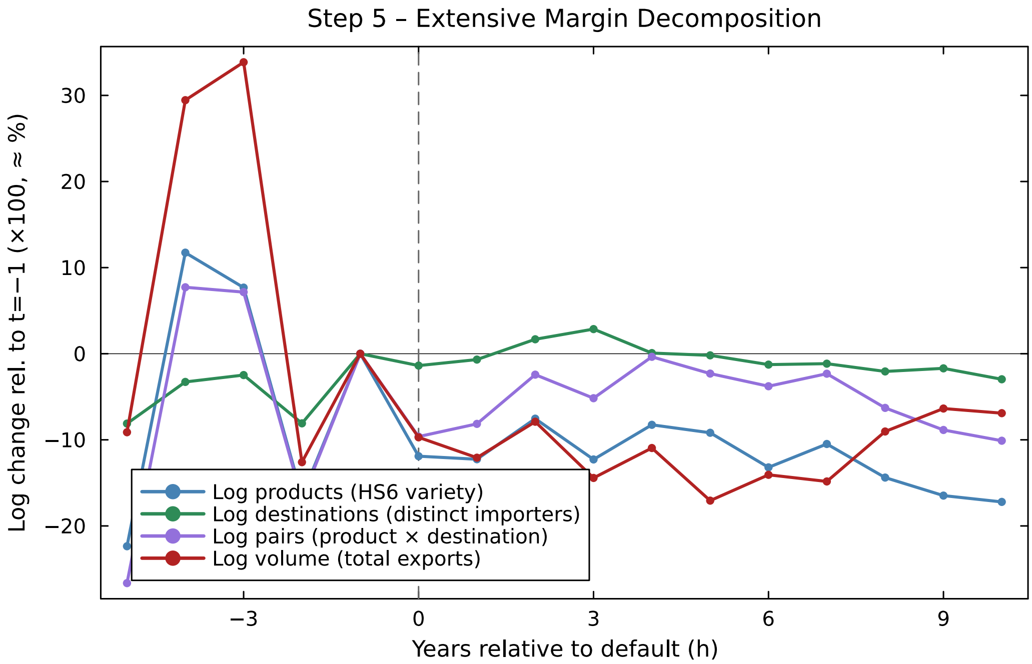

Figure 2 shows that while destinations remain relatively stable after default, the number of exported products and the density of product-destination pairs fall persistently. This distinction matters for interpretation: the post-default export network does not contract primarily by withdrawing from foreign markets. Instead, it becomes thinner—exporters serve similar destinations but with fewer product lines.

The within-destination nature of the scar is confirmed by a direct decomposition. Using the identity $\log(\text{pairs}) = \log(\text{destinations}) + \log(\text{products per destination})$, Figure 3 shows that the decline in product-destination pairs is driven mainly by a fall in within-destination product depth rather than by shrinkage in destination breadth.

This finding points toward a relationship-level mechanism. If the scar were driven by destination-level shocks—for instance, creditor countries imposing trade sanctions—we would expect the destination margin to move. Instead, the damage occurs at the product-destination level, consistent with disruption of specific commercial relationships and the sunk costs of maintaining product lines in individual markets.

This pattern also speaks to the mechanism behind the trade decline. Rose (2005) documented that

Figure 3: **Destination breadth versus within-destination depth.** The main extensive-margin scar reflects thinning *inside* destinations—fewer products per destination—rather than a broad collapse in destination coverage.

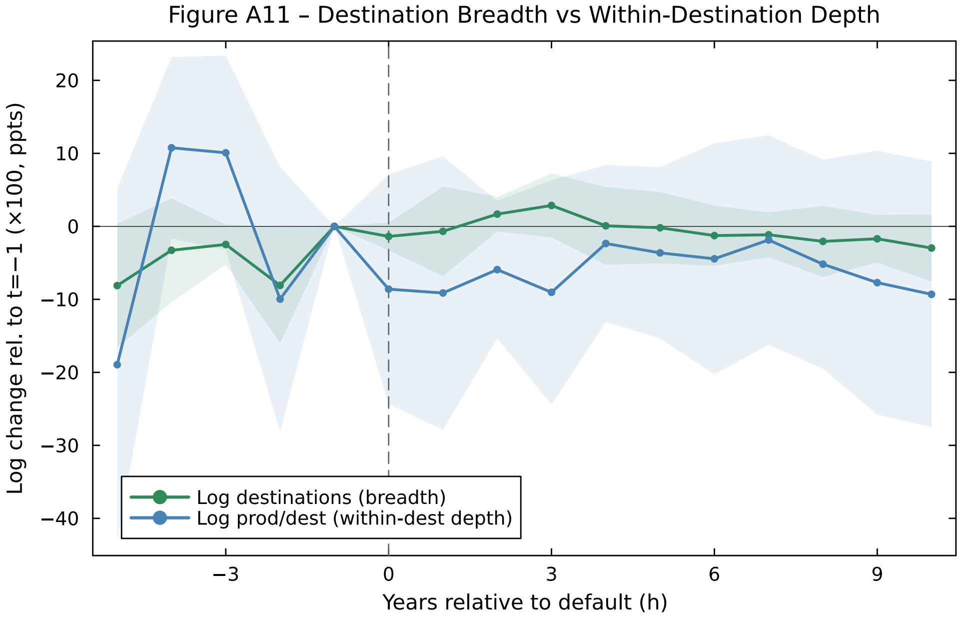

bilateral trade falls between debtor and creditor countries after default, raising the possibility of creditor retaliation. However, Martinez and Sandleris (2011) showed that the decline extends to non-creditor trade partners, ruling out a pure sanctions story. Our within-destination thinning result reinforces this conclusion: the scar is not about losing access to specific markets but about a broad erosion of product-level relationships across all destinations.

3.4 Relationship Dynamics: Survival and Re-entry

We next zoom into the link level and track individual product-destination export relationships before and after default, comparing default episodes with matched non-default controls.[6]

Two facts emerge. First, pre-crisis export relationships survive at markedly lower rates after default than in control episodes (Figure 4a). Default does not merely compress trade contemporaneously—it pushes export relationships into a more fragile regime in which links continue to disappear at elevated rates during the recovery period.

Second, relationships that are lost do not quickly rebuild. Figure 4b shows that cumulative re-entry of exited links is persistently lower in default episodes. This result is among the strongest in the analysis and is central to the interpretation. A pure one-off disruption story is insufficient: the data support a dynamic mechanism in which default both destroys relationships and slows the process by which dormant relationships become active again.

Heterogeneity across links reinforces the interpretation. Larger and older pre-default links survive at much higher rates, while peripheral links are disproportionately fragile (Appendix D,

[6]Controls are non-defaulting EM countries matched on region and income level. Confidence bands are constructed by bootstrapping at the episode level. Details are in Appendix D.

Figure 4: **Relationship destruction and slow rebuilding after default.** *Panel (a)*: the fraction of pre-default export links that remain active falls more steeply in default episodes (treated) than in matched controls. *Panel (b)*: links that exit after default return more slowly—cumulative re-entry is persistently lower for treated episodes. Shaded bands are episode-level bootstrap confidence intervals.

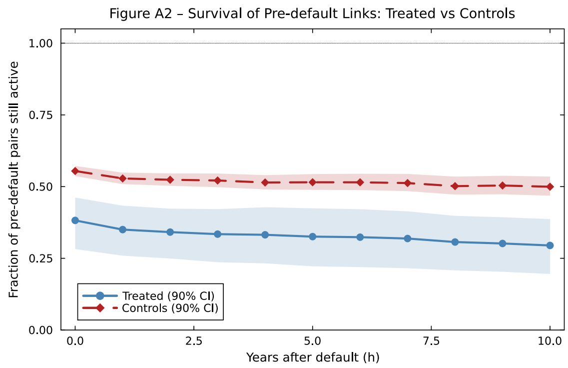

(a) Survival of pre-default links

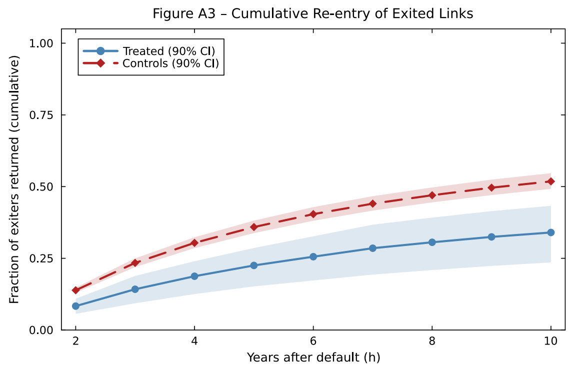

(b) Cumulative re-entry of lost links

Figures 19 and 20). Interestingly, conditional on exit, re-entry does not differ as dramatically across pre-default size quartiles—suggesting that the key asymmetry is between destruction (where size protects) and reactivation (which is broadly difficult for all links).

3.5 Network Structure and the Anatomy of the Export Scar

The facts documented in Sections 3.2–3.4 establish that export relationships are lost in large numbers after default and rebuilt slowly. This section addresses a question they leave open: if so many relationships are destroyed, why does their direct contribution to the aggregate export shortfall appear modest? Answering this question clarifies what kind of cost channel the model needs to explain.

Marginal links are many but small. The resolution lies in the pre-default characteristics of the links that eventually exit. Figure 5 classifies all pre-default active product-destination pairs by their post-default fate and compares their count share to their value share. Links that exit by horizon $h = 5$ account for 64.6 percent of pre-default pair counts but only 16.3 percent of pre-default export value. The reverse holds for survivors: 35 percent of links carry 84 percent of value.

This disproportion reflects a structural feature of the export network rather than anything specific to default. The distribution of export values across product-destination pairs is highly skewed — a small number of large, established relationships account for the bulk of shipments, while the majority of active pairs are near-marginal in value. Default amplifies exit across the network, but the links that disappear are disproportionately small because they were already near the zero-profit threshold before the shock.[7] Figure 6 shows that default shifts the survival

[7]The median pre-default value of eventual exiters is approximately one-sixth that of survivors (USD 36,000 vs. USD 206,000). A link-level regression of exit probability on log pre-default value yields a coefficient of -0.12 (standard error clustered at the episode level), confirming that larger links are substantially less likely to exit.

Figure 5: **Exiters are many but small.** Share of pre-default link counts (left bar) and pre-default export value (right bar) accounted for by eventual exiters, classified at horizon $H = 5$. Exiters represent 64.6 percent of pre-default links by count but only 16.3 percent by value. Source: BACI-CEPII HS92; authors' calculations on the five post-default restructuring episodes.

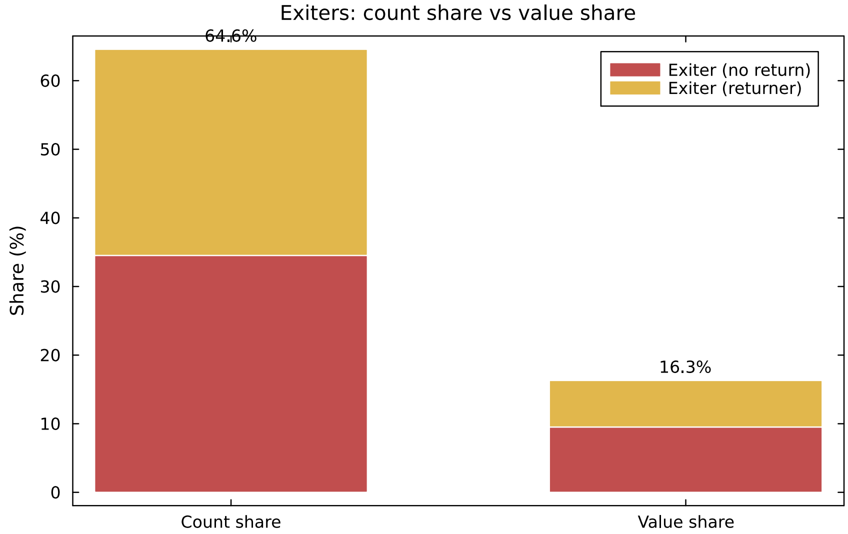

schedule downward broadly across the size distribution — the treated curve lies uniformly below the control curve at every size decile — while the positive size gradient in survival is present in both default and non-default episodes. The links that disappear are small not because default specifically targets them, but because the network's value is concentrated at the top of the size distribution.

Hysteresis lives in the relationship stock, not in per-link revenue. That exiting links were already small explains why their disappearance registers little in aggregate revenue accounting. It does not explain why total exports remain persistently depressed a decade after default. Figure 7 addresses this directly by decomposing total exports into two components: $\log X_t = \log A_t + \log \bar{x}_t$, where A_t is the count of active product-destination pairs and $\bar{x}_t = X_t / A_t$ is average export value per active pair.[8]

The two components behave asymmetrically after default. Average revenue per active product, $\bar{x}_t$, dips at default and returns toward its pre-default level within three to four years; by $h = 10$ it turns positive, indicating that surviving relationships are generating more revenue per link than in the pre-default period. The active product count, A_t, declines sharply on impact and remains depressed throughout the post-default window, reaching -17 log points at $h = 10$. Figure 8 shows this asymmetry in additive form: the extensive margin (blue bars) dominates the negative contribution to $\Delta \log X_t$ at every post-default horizon, while the per-link margin (red bars) partially offsets it and eventually turns positive. Hysteresis in aggregate exports is driven by a persistent failure to rebuild the stock of active relationships, not by lasting weakness within continuing ones.

[8]The figure plots local-projection coefficients estimated using the same specification as equation (1), with country and year fixed effects and standard errors clustered at the country level. Each series is normalized to zero at $h = -1$.

Figure 6: **Default shifts survival downward across the size distribution.** Survival rate at $H = 5$ by pre-default value decile, separately for default episodes (red) and matched control episodes (blue). The positive size gradient is present in both groups; default shifts the entire schedule downward rather than differentially destroying small links.

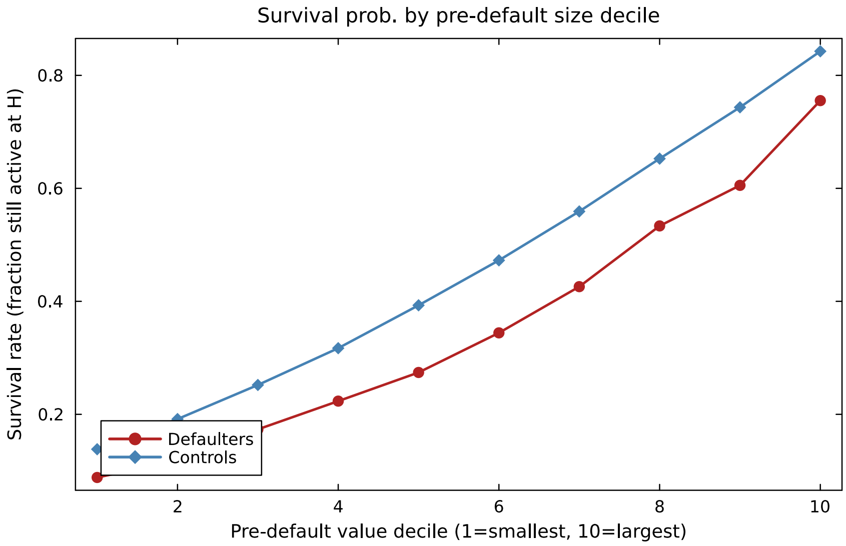

Surviving incumbents grow more slowly after default. A complementary question is whether the revenue path of surviving incumbents — links that were active before default and remain active afterward — also reflects the damage. Figure 9 tracks normalized mean revenue for the cohort of links active at $h = -1$ that remain active through $h = 3$, separately for default and control episodes. Both cohorts grow in absolute terms, as surviving relationships tend to be established and expanding. A gap between the two series opens around $h =$ 3–5, after restructuring is complete and the direct wedge effects have normalized, and persists through $h = 10$: control incumbents reach approximately 2.5 times their pre-default revenue by $h = 10$, while treated incumbents reach 2.3 times theirs — a gap of roughly 8 percentage points in cumulative revenue growth.[9]

The timing of the divergence — emerging after restructuring rather than at default — is consistent with gradual erosion of relationship-specific capital during the default episode rather than a pure contemporaneous cost channel. We interpret this evidence cautiously: the surviving cohort is endogenously selected, limiting causal identification. The robust finding from this section is that hysteresis is primarily a stock phenomenon (the network fails to rebuild breadth) rather than a revenue-per-link phenomenon, and that the links absorbed by the network's scar were already near-marginal before the default occurred.

Bridge to the model. The three findings jointly characterize the empirical environment the model must explain. Default destroys many relationships, but the destroyed links were already near-marginal, so their direct revenue contribution is small. The persistent export scar reflects

[9] The surviving incumbent panel contains only links that cleared the survival selection. The surviving set in default episodes is therefore more positively selected than in control episodes, which attenuates the gap and biases against finding an effect. We interpret the revenue divergence as a lower bound on the true impairment of continuing relationships.

Figure 7: **The relationship stock remains depressed; per-link revenue recovers.** Local-projection impulse responses for log X (total exports, black), log A (active product count, blue), and log $\bar{x}$ (average value per active product, red), normalized to zero at $h = -1$. Shaded bands are 90% confidence intervals from wild-bootstrapped standard errors clustered at the country level (Rademacher weights, 999 repetitions). The active product count remains persistently below its pre-default level while per-link revenue recovers and turns positive. Source: BACI-CEPII HS92; five post-default episodes.

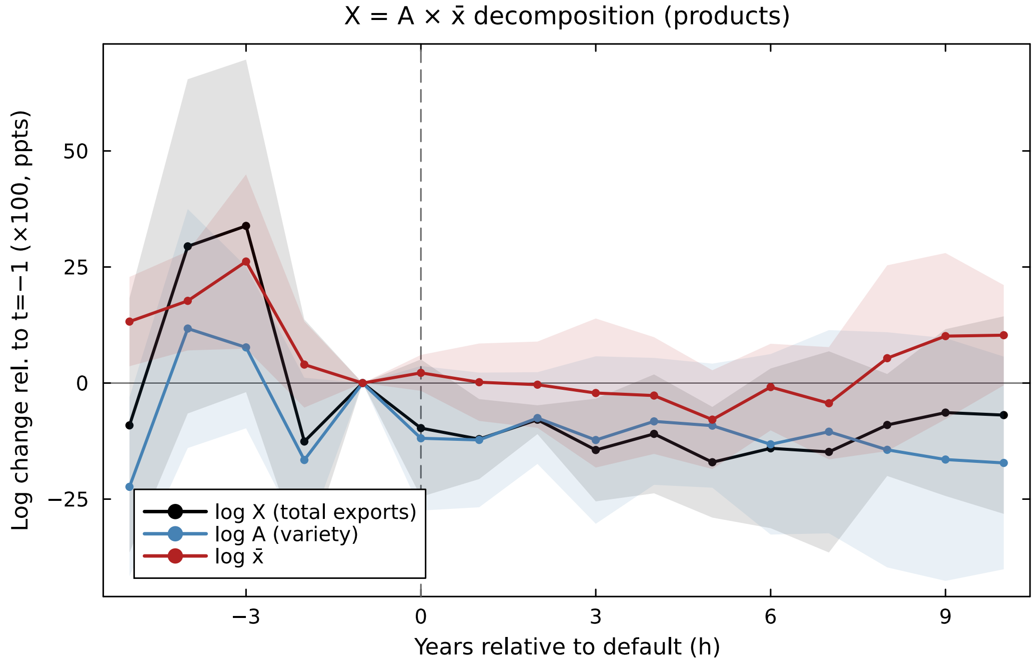

a failure to rebuild the stock of active relationships — A_t remains depressed long after the direct default costs have normalized — with a secondary contribution from slower growth within surviving incumbents. The model in Section 4 interprets this through two mechanisms: a threshold φ_A^* that rises with default wedges and selectively removes marginal links from the active pool, and slow-moving relationship stocks (A_t, D_t) whose rebuilding is impeded by search frictions even after the wedges themselves normalize. The per-link revenue recovery visible in Figure 7 is consistent with the model: once the threshold falls back and wedges normalize, continuing links earn close to their steady-state revenue, but the stock cannot rapidly replenish because new matching and reactivation are structurally slow. The output cost of default therefore operates through the stock channel — depressed A_t lowers aggregate export profit Π_t^X — rather than through contemporaneous revenue impairment within surviving links.

3.6 Composition of Recovery

The final micro-level question is: who drives the recovery that does occur? We decompose post-default export value into contributions from continuing relationships (active before and after default), re-entering relationships (active before default, exited, then returned), and genuinely new relationships (first appearance after default).

Figure 10 delivers a clear message. New links are not absent—they appear in substantial

Figure 8: **The extensive margin dominates the persistent export shortfall.** Additive decomposition of $\Delta \log X_t$ into contributions from the active product count ($\Delta \log A_t$, blue bars) and average value per active product ($\Delta \log \bar{x}_t$, red bars); the black line plots the total $\Delta \log X_t$. At every post-default horizon the variety contribution is negative and larger in magnitude than the per-link contribution, which partially recovers and turns positive at longer horizons.

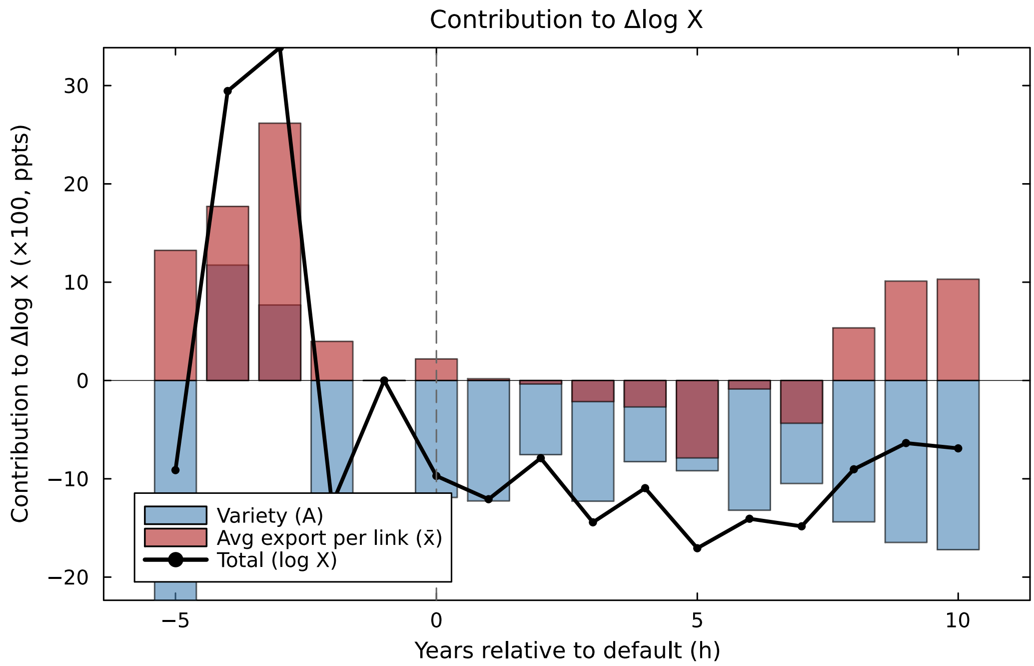

numbers—but they remain quantitatively small in value terms. Export recovery is overwhelmingly dominated by continuing relationships and, to a lesser extent, by re-entering pre-existing relationships. The recovery is *incumbent-heavy*: surviving commercial ties carry most of the economic weight, and genuinely new matches contribute little to the value rebound.

This finding has direct implications for modeling. It suggests that post-default trade dynamics are shaped by a stock of relationship-specific capital that depreciates during default and is costly to rebuild. The model in Section 4 incorporates this through a customer-capital mechanism in which incumbent relationships carry accumulated value that new matches lack.

3.7 Interpretation, Scope, and Robustness

Interpretation. The four facts documented above support a coherent interpretation: sovereign default persistently damages the export trade network, not primarily by eliminating destination reach, but by thinning the network through product-scope losses, lower survival of existing product-destination relationships, and slower re-entry of dormant trade links. Recovery in export value occurs mainly through surviving and returning incumbent relationships rather than through broad creation of new trade links.

Scope. Several limitations should be noted. First, with five treated episodes in four countries, these findings are best understood as stylized facts rather than precise causal estimates. The confidence intervals in the local-projection analysis are wide by construction, and the convergent

Figure 9: **Surviving incumbents grow more slowly after default.** Mean export value per link for the cohort of pre-default active links that remain active through $h = 3$, normalized to 1 at $h = -1$. Treated episodes (red) and matched control episodes (blue) both grow in levels; a gap opens around $h = 3$–5 and persists through $h = 10$. Shaded bands are episode-level bootstrap confidence intervals. The positive selection of surviving incumbents in default episodes attenuates the measured gap; the true impairment is at least as large as shown.

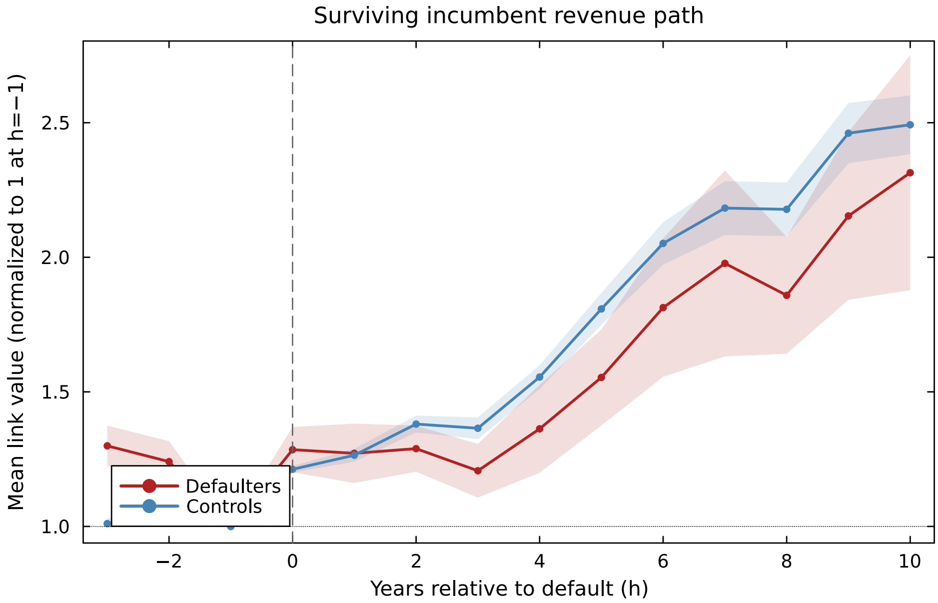

evidence from multiple analytical approaches (LP, case study, relationship dynamics) is more informative than any single point estimate. Second, the analysis documents the trade-network consequences of default but does not directly quantify their output cost. The model in Section 4 provides the structural framework for that quantification. Third, the import-side evidence for a complementary channel through intermediate and capital-good variety is weak in the current data (Appendix E).

Robustness. The baseline results are robust to a range of alternative specifications, reported in Appendix D:

- Alternative timing definitions (default start, restructuring end, midpoint);
- Higher value threshold (USD 50,000);
- Including fuels and minerals (HS chapters 26–27);
- Bilateral variety (product $\times$ destination pairs);
- Dropping Argentina.

A particularly informative check uses preemptive restructurings—episodes where the sovereign restructured without missing payments—as a placebo. Export variety does *not* fall in preemptive episodes, consistent with the interpretation that the mechanism is specific to hard default and its associated disruption of trade finance and commercial relationships.

Figure 10: **Composition of export value after default.** Value shares by link type for $h \geq 1$. Continuing relationships account for the overwhelming majority of export value during recovery. New links appear in meaningful counts but remain economically small.

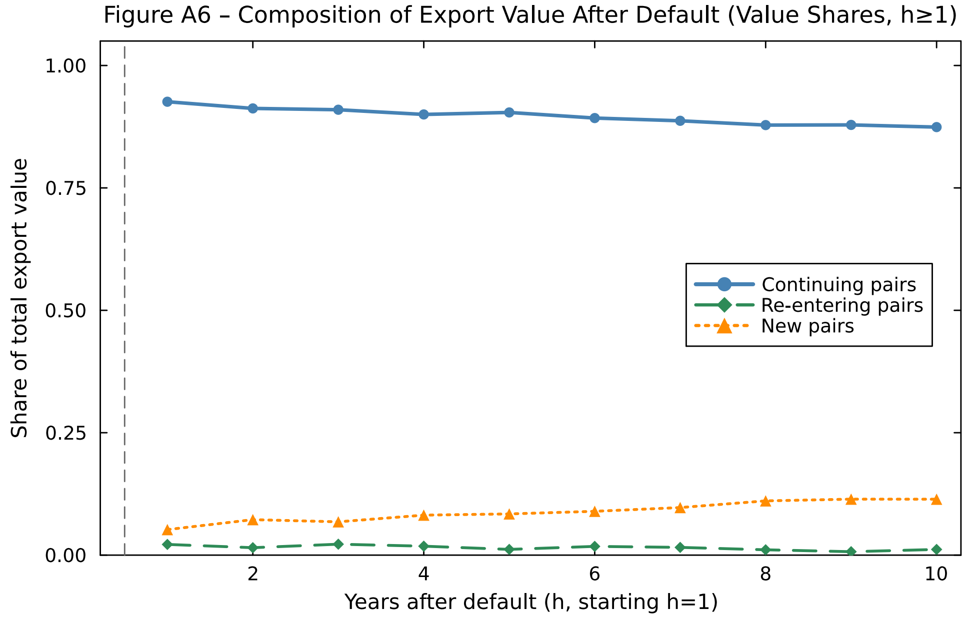

4 Model

The empirical evidence in **??** isolates a specific mechanism: sovereign default causes a persistent contraction in the stock of active export relationships that recovery in export volumes does not reverse. Standard sovereign default models cannot generate this pattern. In **?** and **?**, the output cost of default is a reduced-form endowment penalty that depends only on the current shock y; it reverses when exclusion ends and leaves no trace in future productive capacity. The scar we document is different in kind, not only in magnitude. It lives in a slow-moving network stock, and rebuilding that stock requires time, search, and relationship investment that cannot be commanded by simply returning to markets.

This section presents a model designed to match that mechanism. It has three layers. The first is a *demand system*: quality-adjusted constant-elasticity-of-substitution (CES) preferences over differentiated varieties generate a revenue function that depends jointly on firm productivity and bilateral relationship capital from first principles. The second is a *trade relationship block*: each export opportunity occupies one of three states—active, dormant, or unmatched—with transitions governed by a search-and-matching technology derived from a bilateral meeting function. This block generalises **?** with dormancy, reactivation frictions, and learning-by-doing in relationship capital. The third is a *sovereign debt block*: the **?** long-duration bond framework, extended so that realised default damages the trade block through four economically interpretable wedges.

A key architectural feature is the block-recursive structure. Firms respond to the realised sovereign status $\zeta \in \{0, 1\}$, not to the bond price schedule q. Consequently, individual firm decisions depend on the aggregate state $\Omega \equiv (y, A, D, \zeta)$—where A and D denote the masses of active and dormant relationships—but not on current sovereign debt b. This separation permits a two-stage solution strategy in which micro transition rates are precomputed as lookup tables over Ω, and the sovereign's value function iteration proceeds on the full five-dimensional state (b, y, A, D, ζ) treating those tables as given.[10]

4.1 Environment

Time and uncertainty. Time is discrete, $t = 0, 1, 2, \ldots$, with one period equal to one quarter. Aggregate uncertainty is a domestic productivity endowment y_t following a stationary AR(1) in logs,

$$\ln y_{t+1} = (1 - \rho_y)\mu_y + \rho_y \ln y_t + \sigma_y \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \overset{\text{iid}}{\sim} \mathcal{N}(0, 1). \tag{2}$$

Agents. Four classes of agents interact. The *sovereign government* is a benevolent planner that maximises the present discounted value of household utility and finances consumption by borrowing from international bond markets. *Domestic producers* are a continuum of monopolistically competitive firms, each identified with a single differentiated variety and endowed with productivity φ drawn from a Pareto distribution $G(\varphi)$ on $[\underline{\varphi}, \infty)$; productivity is fixed for the life of the firm.[11] *Foreign buyers* are a representative consumer in each destination with preferences over imported varieties who are passive in the sovereign's problem but one active side of the matching friction in the trade block. *International lenders* are risk-neutral and competitive, pricing bonds at break-even with opportunity cost r^*.

The model abstracts from domestic labour markets, domestic variety demand, physical capital, and monetary policy. The domestic economy is an endowment economy augmented by export profits; general-equilibrium feedback from factor markets is absent.

4.2 Demand, Technology, and the Revenue Function

Quality-adjusted CES preferences. The representative consumer in destination n has preferences

$$U_n = \left(\int_{\omega \in \Omega_n} a_n(\omega)^{1/\sigma} \, q_n(\omega)^{(\sigma-1)/\sigma} \, \mathrm{d}\omega \right)^{\sigma/(\sigma-1)}, \qquad \sigma > 1, \tag{3}$$

where $q_n(\omega)$ is quantity of variety ω, σ is the elasticity of substitution, and $a_n(\omega) \geq 0$ is a buyer-specific quality weight. We specify

$$a_n(\omega) = \exp(\xi \, k_n(\omega)), \qquad \xi > 0, \tag{4}$$

[10]The computational strategy is described in detail in **??**. The block-recursive structure is valid precisely because firms respond to ζ. If firms responded to $q(y, A, D, b')$, the micro objects would depend on b' and the two stages could not be separated.

[11]We abstract from free entry of new exporters. This is consistent with the empirical evidence: Fact (v) shows that genuinely new export links contribute negligibly to post-default recovery, suggesting the relevant extensive margin is the survival and reactivation of existing relationships rather than the entry of new exporters. A free-entry condition would pin the mass of active firms through a zero-profit equation but would not change the qualitative mechanism, since the paper's contribution lies in the transition dynamics of (A, D) rather than in the determination of the steady-state firm mass.

where $k_n(\omega) \geq 0$ is bilateral relationship capital between the producer of variety ω and buyers in destination n. When $k = 0$, preferences reduce to standard CES. When $k > 0$, the buyer places additional weight on the variety—reflecting established distribution channels, adapted product specifications, and current certifications—generating higher demand at any given price. This is the demand-side foundation of the incumbent revenue advantage documented in Fact (vi).

Technology and pricing. Each firm produces with a linear technology at marginal cost τ_t/φ, where $\tau_t \geq 1$ is an iceberg trade cost and the domestic wage is normalised to one. Under monopolistic competition with CES demand, the firm charges a constant markup $\sigma/(\sigma-1)$ over marginal cost:

$$p_n(\varphi) = \frac{\sigma}{\sigma - 1} \cdot \frac{\tau_t}{\varphi}. \tag{5}$$

The markup is independent of relationship capital: k raises demand but not the cost-based optimal price.

Revenue function. Substituting eq. (5) and $a_n(\omega) = \exp(\xi k)$ into the demand function derived from eq. (3), and collecting the destination-level demand shifter $\Xi_n \equiv [(\sigma-1)/\sigma]^{\sigma-1} P_n^{\sigma-1} R_n$ and the macro demand index $Z(y_t) = (y_t/\bar{y})^{\eta}$ ($\eta > 0$), revenue is

$$r(\varphi, k; y_t, \tau_t) = \Xi_n \, Z(y_t) \cdot \exp(\xi k) \cdot \left(\frac{\varphi}{\tau_t}\right)^{\sigma-1}. \tag{6}$$

Revenue has three multiplicative components: a destination-level factor $\Xi_n Z(y_t)$ common to all active relationships in n; a bilateral component $\exp(\xi k)$ specific to each firm-buyer pair; and a firm-level component $(\varphi/\tau_t)^{\sigma-1}$. The bilateral component is the mechanism through which relationship capital creates revenue dispersion across otherwise identical firms. The derivation of eq. (6) from the utility maximisation problem in eq. (3) is provided in **??**.

Operating profit. Including a working-capital wedge $\omega_t \in [0, 1)$, which captures the fraction of revenue absorbed by trade-finance costs (letters of credit, export credit insurance, pre-shipment financing), and a per-period fixed market access cost $f > 0$, operating profit for an active link is

$$\pi^A(\varphi, k; \Omega) = (1 - \omega_t)\, r(\varphi, k; y_t, \tau_t) - f. \tag{7}$$

The working-capital wedge scales with revenue and therefore affects the sovereign's resource constraint as well as the firm's participation decision; we return to this in section 4.6.

4.3 Trade Relationships: States, Matching, and Structural Frictions

Relationship states. Each potential export opportunity has a state $s \in \{A, D, U\}$. An *active* (A) relationship is one in which the firm and buyer are currently trading; the firm earns revenue $r(\varphi, k)$, pays fixed cost f, and accumulates relationship capital through learning-by-doing. A *dormant* (D) relationship is one in which the pair has traded previously but is currently not transacting; some match-specific capital survives—product certifications, adapted specifications, buyer knowledge—and reactivation is possible at lower cost than forming a new relationship from scratch. Capital depreciates faster in dormancy than in activity, so the reactivation advantage fades with time since separation. An *unmatched* (U) relationship represents a potential export opportunity with no prior connection; the firm must search for a buyer, pay a higher setup cost, and begin with minimal initial capital.

The aggregate stocks are A_t, D_t, and $U_t = \bar{M} - A_t - D_t$, where $\bar{M}$ is the total mass of potential export opportunities.[12] The state pair (A_t, D_t) is a sufficient statistic for the export network; recovering U_t from $\bar{M}$ requires no additional tracking.

A relationship in the model is a bilateral link between a domestic firm and a foreign buyer, specific to a product and destination. The aggregate stocks (A_t, D_t) therefore count *product-destination* opportunities, not merely product opportunities. The quantitative model restricts attention to a single composite destination, which is equivalent to assuming symmetric destinations with identical demand conditions Ξ_n, buyer capacity V_n, and setup costs F_n^N, F_n^R. This is the standard aggregation in single-destination Melitz models and is consistent with our empirical strategy, which pools across destinations when constructing variety counts. The restriction sacrifices the ability to separately identify within-destination thinning from across-destination withdrawal, both of which are documented in Fact (iv) and discussed in section ??. A multi-destination extension that allows $N > 1$ destinations with heterogeneous Ξ_n is a natural direction for future work; we leave it aside here because tracking (A_{nt}, D_{nt}) per destination would add N state variables and render the current solution strategy computationally infeasible.

New matching. An unmatched firm with productivity φ invests search effort $x \geq 0$ at convex cost $c(x) = c\, x^{1+1/\alpha_x}/(1+1/\alpha_x)$. The bilateral meeting function is Cobb-Douglas:

$$M^N = \chi(\zeta) \cdot (X \cdot U_t)^{\alpha} \cdot V_n^{1-\alpha}, \qquad \alpha \in (0,1), \tag{8}$$

where X is average search effort in the unmatched pool, $V_n > 0$ is buyer-side matching capacity in destination n, and $\chi(\zeta)$ is matching efficiency that falls during default. Dividing by pool mass and attributing individual effort x to firm φ yields the per-firm match probability[13]

$$\lambda^N(x; \Omega) = \chi(\zeta) \cdot x^{\alpha} \cdot \left(\frac{V_n}{U_t}\right)^{1-\alpha}. \tag{9}$$

Match probability is increasing in own effort and decreasing in congestion. If a match forms, the firm pays setup cost $F^N > 0$ and begins trading at initial capital k_0.

[12] $\bar{M}$ is calibrated to match the pre-default steady-state utilization rate $A^*/\bar{M}$, which equals the observed ratio of active product-destination pairs (annual export value > USD 10,000, excluding HS chapters 26–27) to all potential pairs (HS6 codes $\times$ destination countries in the sample). In the data, this utilization rate is approximately [TODO: XX]%.

[13] In the single-destination model, V_n is a scalar constant not separately identified from $\chi(\zeta)$: the product $\chi(\zeta) \cdot V_n^{1-\alpha}$ is observationally equivalent to a rescaled matching efficiency $\tilde{\chi}(\zeta)$ with $V_n = 1$. We normalize $V_n = 1$ throughout; the parameter $\chi(\zeta)$ absorbs the scale of buyer-side capacity. In a multi-destination extension, V_n would vary across destinations and provide a source of cross-sectional identification.

Reactivation. A dormant link retains a memory of its buyer, making reactivation structurally easier than new matching. The bilateral meeting function for the dormant pool is

$$M^R = \chi^R(\zeta) \cdot \exp\big(\psi \, \bar{k}^D\big) \cdot D_t^{\alpha_R} \cdot V_n^{1-\alpha_R}, \tag{10}$$

where $\bar{k}^D$ is representative dormant capital (defined in section 4.5) and $\exp(\psi \bar{k}^D)$ captures that links with more surviving capital are easier for buyers to re-engage. The per-link reactivation probability is

$$\lambda^R(\bar{k}^D; \Omega) = \chi^R(\zeta) \cdot \exp\big(\psi \, \bar{k}^D\big) \cdot \left(\frac{V_n}{D_t}\right)^{1-\alpha_R}. \tag{11}$$

If reactivation succeeds, the firm pays the lower setup cost $F^R < F^N$ and resumes trading.

For parsimony, we impose $\alpha_R = \alpha$ in the baseline: the congestion elasticity is the same for new matching and reactivation. Separate identification of α_R would require a moment that is specifically sensitive to congestion in the dormant pool independently of congestion in the unmatched pool, which is not available in our data.[14]

Structural ranking. The model embeds the following assumption about relative frictions:

$$\chi^R(\zeta) > \chi(\zeta), \qquad F^R < F^N, \qquad \bar{k}^D > k_0. \tag{12}$$

Reactivation is structurally more likely, cheaper, and more rewarding than new matching at any given ζ. This ranking is the model's counterpart to Fact (v) in **??**: the near-zero contribution of genuinely new links to post-default recovery is explained by the high cost of new matching relative to reactivation, not by the absence of unmatched firms. Default tightens the wedge further by depressing both $\chi(\zeta)$ and $\chi^R(\zeta)$, amplifying the differential.

4.4 Relationship Capital

Relationship capital $k \geq 0$ accumulates through trade and depreciates without it. When a link is active, capital evolves according to

$$k_{t+1} = (1 - \delta_A)\, k_t + \iota_A, \tag{13}$$

with active depreciation $\delta_A \in (0, 1)$ and learning-by-serving flow $\iota_A > 0$. The steady-state active capital is $k^{A,ss} = \iota_A / \delta_A$. When a link is dormant, no trade occurs, so capital depreciates at the faster rate $\delta_D > \delta_A$:

$$k_{t+1} = (1 - \delta_D)\, k_t. \tag{14}$$

The ranking $\delta_D > \delta_A$ has two consequences. First, the reactivation advantage of a dormant link declines with time since separation, as $\bar{k}^D$ falls and therefore λ^R in eq. (11) falls. Second, default episodes accelerate the erosion of dormant capital by keeping links in dormancy longer, compounding the matching friction during recovery.

[14] Economically, one might expect $\alpha_R < \alpha$: reactivating firms target a specific prior buyer rather than competing for any available buyer, making reactivation less sensitive to pool-level congestion. Appendix ?? reports a sensitivity check that allows $\alpha_R \neq \alpha$ and confirms that the main results are not sensitive to this restriction.

New matches begin at capital $k_0 \ll k^{A,ss}$. Since revenue is $\exp(\xi k)$ times a baseline (from eq. (6)), a newly formed relationship earns $\exp(\xi k_0)$ while a mature active relationship earns $\exp(\xi k^{A,ss})$—generating an incumbent revenue advantage that survives as long as the gap between $k^{A,ss}$ and k_0 is large. This is the mechanism behind Fact (vi): recovery in export value is driven by surviving incumbents, not by the cohort of new or reactivated links.

4.5 Firm Decisions, Thresholds, and Aggregation

Micro value functions. Each firm takes aggregate conditions $\Omega = (y, A, D, \zeta)$ as given. Active, dormant, and unmatched relationships solve the following Bellman equations.

An active firm with characteristics (φ, k) chooses between continuing to export and going dormant:

$$V^A(\varphi, k; \Omega) = \max\Big\{\pi^A(\varphi, k; \Omega) + \beta\,\mathbb{E}\Big[V^A(\varphi, k'; \Omega')\Big],\ \ V^{D,m}(\varphi, k; \Omega)\Big\}, \tag{15}$$

where $k' = (1-\delta_A)k + \iota_A$ and $V^{D,m}$ is the value of a dormant relationship defined below.

A dormant firm with characteristics (φ, k) chooses between staying dormant (and attempting reactivation) and abandoning the relationship entirely:

$$\begin{aligned} V^{D,m}(\varphi, k; \Omega) = \max\Big\{ & -m_D + \lambda^R(\bar{k}^D; \Omega)\Big[-F^R + \beta\,\mathbb{E}V^A(\varphi,\, k^R_{t+1}; \Omega')\Big] \\ & + (1-\lambda^R(\bar{k}^D; \Omega))\,\beta\,\mathbb{E}V^{D,m}(\varphi,\, k^D_{t+1}; \Omega'),\ \ V^U(\varphi; \Omega)\Big\}, \end{aligned} \tag{16}$$

where $m_D = 0$ (dormancy maintenance costs are set to zero in the baseline[15] see below); $k^R_{t+1} = (1-\delta_A)k + \iota_A$ is capital upon successful reactivation, $k^D_{t+1} = (1-\delta_D)k$ is capital if reactivation fails and the link remains dormant, and V^U is the value of an unmatched opportunity defined next.

An unmatched firm with productivity φ chooses search effort $x \geq 0$:

$$V^U(\varphi; \Omega) = \max_{x\geq 0}\Big\{-c(x) + \lambda^N(x; \Omega)\Big[-F^N + \beta\,\mathbb{E}V^A(\varphi, k_0; \Omega')\Big] + (1-\lambda^N(x; \Omega))\,\beta\,\mathbb{E}V^U(\varphi; \Omega')\Big\}. \tag{17}$$

The first-order condition for x yields optimal effort $x^*(\varphi; \Omega)$ that is increasing in φ: more productive firms have a larger payoff from forming a new match and therefore search harder. The φ-dependence of x^* complicates the aggregation of the new-creation rate, as discussed below.

Thresholds. The Bellman equations generate two endogenous productivity cutoffs. The *separation threshold* $\varphi^*_A(\Omega)$ is the productivity level at which an active firm is indifferent between

[15] Setting $m_D = 0$ removes an unidentified free parameter. A positive m_D would create a direct resource drain during post-default episodes when D is large, competing with the Π^X compression channel as an explanation for observed output dynamics. Since we do not have a moment that separately identifies "links die because maintenance is costly" from "links die because reactivation is unprofitable," we impose $m_D = 0$ and let the dormancy-death threshold be determined entirely by the option value of reactivation relative to starting fresh. Appendix ?? reports robustness to $m_D > 0$.

continuing and going dormant. Under the model's regularity condition $\varphi_A^* > \varphi_D^*$ (which holds whenever the structural ranking eq. (12) is satisfied at calibrated values), the dormant option at the separation margin resolves to staying dormant rather than abandoning, so the threshold condition simplifies to

$$V^{A,\text{cont}}(\varphi_A^*;\,\Omega) = V^{D,\text{stay}}(\varphi_A^*;\,\Omega), \tag{18}$$

where $V^{A,\text{cont}}$ and $V^{D,\text{stay}}$ are stationary-Ω approximations to the continuation branches of eqs. (15) and (16); their closed-form expressions are given in **??**. The *dormancy-death threshold* $\varphi_D^*(\Omega)$ is the productivity level at which a dormant firm is indifferent between staying dormant and abandoning:

$$V^{D,\text{stay}}(\varphi_D^*;\,\Omega) = V^U(\varphi_D^*;\,\Omega). \tag{19}$$

Both thresholds are found by bisection.[16] Default raises $\varphi_A^*(\Omega)$ and, to a lesser degree, $\varphi_D^*(\Omega)$, so default both increases separation and expands the range of productivity that finds it worthwhile to preserve a dormant link.

Compressed conditional closure. The micro problems eqs. (15) to (17) are defined over individual states (φ, k). Carrying the full within-pool distribution of k as a state variable is computationally infeasible and, given the six empirical targets in **??**, unnecessary: the targets are driven by aggregate stock dynamics (A, D) and the default wedges, not by fine differences in the k cross-section within each pool. We therefore summarise each pool by a single *representative capital*: $\bar{k}^A(\Omega)$ for the active pool and $\bar{k}^D(\Omega)$ for the dormant pool.

The representative capitals are determined by mean-capital accounting: the average capital in each pool must be consistent with the inflows and outflows implied by the aggregate transition rates. For the active pool, links arrive from three sources— surviving active links (capital grew by learning-by-serving), reactivated dormant links (entering at $\bar{k}^D$), and new matches (entering at k_0)—giving[17]

$$\bar{k}^A = \frac{(1-\mathcal{S})\left[(1-\delta_A)\bar{k}^A + \iota_A\right]A + \mathcal{R}\,\bar{k}^D\,D + \mathcal{N}\,k_0\,U}{(1-\mathcal{S})\,A + \mathcal{R}\,D + \mathcal{N}\,U}. \tag{20}$$

For the dormant pool, links arrive from newly separated active relationships (entering at $\bar{k}^A$)

[16] Equations (18) and (19) are derived under the approximation $\Omega' = \Omega$: current aggregate conditions are assumed to persist when evaluating threshold indifference conditions. This stationarity approximation is standard in the literature (see, e.g., **??**) and is adopted here because replacing it with forward-looking thresholds would require iterating the micro problem jointly with the sovereign VFI, destroying the two-stage solution strategy. The direction of bias is conservative for the paper's main result: at the start of a default episode, the approximation overstates the persistence of bad wedges, raising φ_A^* above its rational-expectations value and causing more separation than would occur under correct beliefs about wedge recovery. This overstates the initial shock to A but does not artificially extend its duration, since persistence comes from (A, D) stock dynamics rather than from the threshold level.

[17] The accounting equation for $\bar{k}^A$ is a post-transition average: it includes capital inflows from reactivations and new matches that occur *during* period t. Current-period revenue $\mathcal{X}$, however, is produced by links that were active at the *start* of the period. A strictly correct treatment would evaluate $\mathcal{X}$ at the beginning-of-period capital $(1-\delta_A)\bar{k}^A + \iota_A$ (capital of incumbents after one period of learning-by-serving), rather than the post-transition $\bar{k}^A$. The error is proportional to the reactivation and new-match inflow shares relative to A_t; it is quantitatively small in steady state and bounded in sign (post-transition $\bar{k}^A$ understates current revenue since entrants carry less capital than incumbents). We use the post-transition $\bar{k}^A$ throughout for consistency with the compressed closure.

and surviving dormant links (capital depreciated at rate δ_D):

$$\bar{k}^D = \frac{\mathcal{S}\,\bar{k}^A\,A + (1-\mathcal{R}-\mathcal{L})(1-\delta_D)\,\bar{k}^D\,D}{\mathcal{S}\,A + (1-\mathcal{R}-\mathcal{L})\,D}. \tag{21}$$

Here $\mathcal{S}$, $\mathcal{R}$, $\mathcal{N}$, and $\mathcal{L}$ are aggregate transition rates defined momentarily. Equations eqs. (20) and (21) form a 2×2 fixed point in $(\bar{k}^A, \bar{k}^D)$ given the rates and pool stocks, solved jointly by alternating bisection.

Aggregate transition rates. With Pareto-distributed productivity, the separation rate $\mathcal{S}(\Omega)$ and the dormancy-death rate $\mathcal{L}(\Omega)$ have closed forms:

$$\mathcal{S}(\Omega) = G\big(\varphi_A^*(\Omega)\big) = 1 - \left(\frac{\underline{\varphi}}{\varphi_A^*(\Omega)}\right)^{k_P}, \qquad \mathcal{L}(\Omega) = G\big(\varphi_D^*(\Omega)\big) = 1 - \left(\frac{\underline{\varphi}}{\varphi_D^*(\Omega)}\right)^{k_P}, \tag{22}$$

where $k_P > \sigma - 1$ is the Pareto shape parameter. The reactivation rate is also available in closed form, because the compressed closure makes λ^R independent of individual φ:

$$\mathcal{R}(\Omega) = \lambda^R\big(\bar{k}^D;\,\Omega\big) \cdot \left(\frac{\underline{\varphi}}{\varphi_D^*(\Omega)}\right)^{k_P}. \tag{23}$$

The new-creation rate $\mathcal{N}(\Omega)$ does not have a closed form because $x^*(\varphi;\,\Omega)$ varies with φ, making $\lambda^N(x^*(\varphi);\,\Omega)$ a non-power function of φ. It is evaluated by Gauss–Legendre quadrature over the Pareto tail above the search-participation threshold $\varphi_U^*(\Omega)$ (see **??**).

Laws of motion and aggregate revenue. The aggregate stocks evolve as

$$A_{t+1} = (1-\mathcal{S}_t)\,A_t + \mathcal{R}_t\,D_t + \mathcal{N}_t\,U_t, \tag{24}$$

$$D_{t+1} = \mathcal{S}_t\,A_t + (1-\mathcal{R}_t-\mathcal{L}_t)\,D_t, \tag{25}$$

which we write compactly as $A_{t+1} = \Gamma^A(\Omega_t)$ and $D_{t+1} = \Gamma^D(\Omega_t)$. Aggregate export revenue integrates individual revenues over the Pareto distribution above the separation threshold:

$$\mathcal{X}(\Omega) = A \cdot Z(y)\,\exp(\xi\,\bar{k}^A)\,\tau(\zeta)^{-(\sigma-1)} \cdot \frac{k_P\,\underline{\varphi}^{k_P}}{k_P-(\sigma-1)} \cdot \big(\varphi_A^*\big)^{-[k_P-(\sigma-1)]}. \tag{26}$$

The closed form follows from integrating $r(\varphi, \bar{k}^A; \Omega)$ against the Pareto density above φ_A^*, which delivers a power-law integral because revenue is proportional to $\varphi^{\sigma-1}$ and the Pareto density is $k_P\underline{\varphi}^{k_P}\varphi^{-(k_P+1)}$.

4.6 The Sovereign Block

Long-duration bonds. The sovereign issues non-contingent long-duration bonds following **?**. Each bond pays a coupon κ each period and matures with probability δ (geometric amortisation). The normalisation $\kappa = \delta + r^*$ sets the risk-free price to one. Outstanding debt entering period t

is b_t; new issuance is b_{t+1} at equilibrium price $q(y_t, A_t, D_t, b_{t+1})$.

The bond price depends on the trade network state because low A and high D signal lower future export revenue, lower future output, and therefore higher future default probability. A sovereign that has damaged its export network cannot borrow at the same price as an otherwise identical sovereign with an intact network. This is the second-round feedback from the trade block to bond markets that distinguishes the model from a reduced-form endowment specification.

Domestic output. Total output available to the sovereign is the sum of domestic endowment and net export profits:

$$Y^{\text{tot}}(y, A, D, \zeta) = Y^d(y) + \Pi^X(A, D, y, \zeta), \tag{27}$$

where $Y^d(y) = \exp(y)$ is the non-export endowment and net export profits aggregate individual operating profits over the active pool:[18]

$$\Pi^X(A, D, y, \zeta) = \frac{1 - \omega(\zeta)}{\sigma} \mathcal{X}(\Omega) - f\,A. \tag{28}$$

The factor $1/\sigma$ is the Melitz operating margin (the fraction of revenue left after covering variable trade costs under constant-markup pricing). The working-capital wedge $\omega(\zeta)$ multiplies this entire margin: when ω rises during default, the sovereign's resource constraint tightens both through reduced firm-level participation (via the threshold channel) and through a direct compression of per-unit revenue retained by the household. Entry and reactivation costs are second-order at calibrated values and are omitted from eq. (28).[19]

Resource constraints. Under repayment ($\zeta = 0$), consumption is

$$c_t = Y^{\text{tot}}(y_t, A_t, D_t, 0) - \kappa\, b_t + q(y_t, A_t, D_t, b_{t+1}) \left[b_{t+1} - (1-\delta)\, b_t \right]. \tag{29}$$

Under default ($\zeta = 1$), debt is fully repudiated, so the sovereign retains the entire output but faces a direct default cost:

$$c_t = Y^{\text{tot}}(y_t, A_t, D_t, 1) - C^D(y_t, A_t, D_t). \tag{30}$$

The function $C^D(y, A, D) \geq 0$ is any residual direct cost of default—reputational or legal penalties not already embedded in the wedges.[20] The central mechanism of the paper is not C^D

---

[18] Equation (28) is obtained by integrating $\pi^A(\varphi, k; \Omega) = (1 - \omega) r(\varphi, k; \Omega)/\sigma - f$ over the active pool. Because f is a constant, it factors out of the Pareto integral; the remaining integral of the conditional active-pool density over its support equals one, giving $f \times A$ exactly. Dormancy maintenance costs $m_D D$ enter symmetrically: each of the D dormant links pays m_D per period regardless of productivity. Aggregate revenue $\mathcal{X}(\Omega)$ in the first term is the Pareto integral of $r(\varphi, \bar{k}^A; \Omega)$ over $\varphi \geq \varphi^*_A(\Omega)$, given in closed form as eq. (26).

[19] Entry costs $F^R \mathcal{R}_t D_t$ and $F^N \mathcal{N}_t U_t$ are small relative to the flow of operating profits at our calibrated values. In the model, entry costs account for less than [TODO: XX]% of export-sector value added. This approximation follows the standard in Melitz-style aggregation (**?**).

[20] In the baseline calibration we set $C^D = 0$: the endogenous deterioration of (A, D) through the default wedges is the principal cost channel, and we do not need an additional reduced-form penalty to match the empirical targets. The sensitivity of results to nonzero C^D is reported in **??**.

but the endogenous path of $\mathcal{X}$ during and after default, driven by the slow-moving dynamics of (A_t, D_t).

The sovereign's problem. The sovereign chooses whether to repay or default to maximise the present discounted value of household utility $u(c) = (c^{1-\gamma} - 1)/(1-\gamma)$, $\gamma > 0$. The full state is $s = (b, y, A, D, \zeta)$. The value function satisfies

$$V(s) = \max\Big\{ V^R(s),\ V^D(y, A, D) \Big\}, \tag{31}$$

where $\zeta = 0$ inside the repayment value because any repayment period is, by definition, a non-default period. The repayment value is

$$V^R(s) = \max_{b_{t+1}} \Big\{ u(c_t) + \beta\, \mathbb{E}_{y'|y} \big[V(b_{t+1}, y', A', D', 0) \big] \Big\}, \tag{32}$$

where $A' = \Gamma^A(y, A, D, 0)$ and $D' = \Gamma^D(y, A, D, 0)$ are the trade stocks under no default. The default value is

$$V^D(y, A, D) = u(c^D) + \beta\, \mathbb{E}_{y'|y} \Big[\theta\, V(0, y', A^{D\prime}, D^{D\prime}, 0) + (1-\theta)\, V^D(y', A^{D\prime}, D^{D\prime}) \Big], \tag{33}$$

where θ is the per-period probability of regaining market access, $A^{D\prime} = \Gamma^A(y, A, D, 1)$ and $D^{D\prime} = \Gamma^D(y, A, D, 1)$ are trade stocks under $\zeta = 1$. Note that V^D is defined over (y, A, D) only: since debt is fully repudiated upon default, $b = 0$ and current debt drops out as an argument. This reduces the storage requirement for V^D from five dimensions to three.

Bond pricing. Risk-neutral competitive lenders price bonds at break-even. The bond price satisfies

$$q(y, A, D, b') = \frac{1}{1+r^*}\, \mathbb{E}_{y'|y} \Big[\big(1 - d(b', y', A', D')\big) \Big(\kappa + (1-\delta)\, q(y', A', D', b'') \Big) \Big], \tag{34}$$

where $d(\cdot) \in [0, 1]$ is the equilibrium default probability at the next-period state, $b'' = B(b', y', A', D')$ is the sovereign's optimal debt choice from that state, and A', D' are the trade stocks consistent with b' being chosen today. Current debt b does not appear as an argument of q: lenders care about future default risk, which depends on (y', A', D') and the new debt level b', not on the current stock b. This observation reduces the storage requirement for the bond price from five to four dimensions ($21 \times 30 \times 20 \times 200$ in the baseline grid), yielding a substantial memory saving relative to the naive specification.

4.7 Default Wedges and the Propagation Mechanism

The four wedges. Sovereign default ($\zeta = 1$) shifts four economically distinct wedges that alter firm participation decisions:

$$\tau(\zeta) = \bar{\tau}\exp(\nu_\tau \zeta), \qquad \text{profitability wedge,} \tag{35}$$
$$\omega(\zeta) = \bar{\omega} + \nu_\omega \zeta, \qquad \text{working-capital wedge,} \tag{36}$$
$$\chi(\zeta) = \bar{\chi}\exp(-\nu_{\chi N}\zeta), \qquad \text{new-match efficiency,} \tag{37}$$
$$\chi^R(\zeta) = \bar{\chi}^R \exp(-\nu_{\chi R}\zeta), \qquad \text{reactivation efficiency.} \tag{38}$$

All parameters $\nu_\tau, \nu_\omega, \nu_{\chi N}, \nu_{\chi R} \geq 0$; all wedges revert to their baseline values $\bar{\tau}, \bar{\omega}, \bar{\chi}, \bar{\chi}^R$ when $\zeta = 0$.

The four wedges are economically and empirically distinct. The profitability wedge $\tau(\zeta)$ is a broad export delivery and profitability wedge: it captures the deterioration in the effective return to serving foreign customers during default and exclusion, including market-access penalties, weaker delivery reliability, and the disruption of trade-supporting institutions. It acts on the *separation margin*: higher τ raises the profitability threshold φ_A^*, pushing marginal active links into dormancy. The working-capital wedge $\omega(\zeta)$ is a trade-finance wedge that tightens the cash-flow margin on surviving shipments—letters of credit become harder to obtain, export credit insurance is withdrawn, and pre-shipment financing costs rise. It acts on the *intensive margin*: holding the active set fixed, higher ω compresses the fraction of shipment revenue that is converted into domestic resources. Because τ and ω operate on distinct margins—threshold movement versus revenue compression—they are not redundant, and the calibration section maps each to a separate set of empirical moments (table 5).

On the rebuilding side, $\chi(\zeta)$ reduces the efficiency of forming genuinely *new* export relationships, while $\chi^R(\zeta)$ reduces the efficiency of reactivating *previously served* links. The distinction is central: the empirical evidence points much more strongly to a shortfall in re-entry of lost relationships (Fact iv) than to a broad suppression of new-link creation (Fact v). These two moments therefore target the two wedges separately. Default tightens both, with $\nu_{\chi R} < \nu_{\chi N}$: buyers with prior experience of a supplier are less deterred than entirely new buyers.

The wedges are held constant throughout the exclusion spell. This is conservative for the paper's persistence claim: if the wedges normalised to zero immediately upon restructuring and export damage still persisted, that would constitute cleaner identification of the (A, D) mechanism. Appendix **??** confirms that a specification with time-decaying wedges leaves the persistence results intact.[21]

The mechanism. Default triggers a cascade of adjustments across the export network. When ζ switches from 0 to 1, the wedges τ and ω rise, compressing export profitability for every active link via eq. (7). The separation threshold φ_A^* rises, so marginal active links—those with

[21] Specifically, we consider $\tau(\zeta, s) = \bar{\tau}\exp(\nu_\tau e^{-\lambda_\tau s}\zeta)$ where s is years since default. The baseline constant-wedge specification corresponds to $\lambda_\tau \to 0$.

productivity near the pre-default threshold—find exporting unprofitable and move to dormancy: $\mathcal{S}$ rises, A falls, D rises. Simultaneously, the matching efficiencies χ and χ^R fall: reactivation becomes harder ($\mathcal{R}$ falls) and new-match creation becomes harder ($\mathcal{N}$ falls). The active stock contracts from both ends: it loses members at the bottom of the productivity distribution (separation) and receives fewer replacements from above (reduced reactivation and creation). Aggregate export revenue $\mathcal{X}$ falls in eq. (26) both through the lower stock A and through the higher threshold φ_A^* (which truncates the pool from below). Total output Y^{tot} falls, consumption under repayment falls, and the default set expands. Lenders anticipate higher default probability and the bond price q falls, tightening the sovereign's budget constraint further.

Persistence. The scar persists after default ends. When ζ returns to 0, all four wedges normalise immediately. But the network stocks (A, D) do not: rebuilding A requires that $\mathcal{R}$ and $\mathcal{N}$ replenish the active stock, which is slow for three compounding reasons. First, the dormant pool partially drains to U during the default episode as φ_D^* rises and low-productivity dormant links are abandoned, reducing the pool of reactivatable links. Second, representative dormant capital $\bar{k}^D$ depreciates while links wait in dormancy, lowering λ^R after the episode ends (eq. (11)). Third, new matching through the unmatched pool—the backstop source of active-link replenishment—is structurally slow at any ζ because of the high setup cost F^N and the low initial capital k_0. The consequence is that export revenue and total output remain below their pre-default path for many periods after the exclusion spell ends. This is the hysteresis mechanism documented empirically in **??**: the failure of export variety to expand even as aggregate export volumes recover reflects the slow reconstruction of relationship stocks, not a slow recovery in the profitability of individual surviving links.

4.8 Recursive Equilibrium

Definition 4.1 (Recursive Markov Perfect Equilibrium)**.** *A recursive equilibrium consists of:*

(i) Sovereign value functions $V(s)$, $V^R(s)$, and $V^D(y, A, D)$.

(ii) Sovereign policy functions: default probability $d(s) \in [0, 1]$ and debt choice $b'(s)$.

(iii) A bond price schedule $q(y, A, D, b')$.

(iv) Micro threshold rules $\varphi_A^*(\Omega)$ and $\varphi_D^*(\Omega)$, and search effort $x^*(\varphi; \Omega)$.

(v) Conditional representative capitals $\bar{k}^A(\Omega)$ and $\bar{k}^D(\Omega)$.

such that: (a) given q and the aggregate laws of motion Γ^A, Γ^D, the sovereign's policy and value functions solve eqs. (31) to (33); (b) given Ω, the threshold rules solve eqs. (18) and (19) and x^* solves the firm's problem in eq. (17); (c) given the micro policies, the representative capitals solve eqs. (20) and (21); (d) the aggregate laws of motion eqs. (24) and (25) are consistent with the micro policies; and (e) the bond price schedule satisfies the lenders' break-even condition eq. (34).

The equilibrium is block-recursive by construction: conditions (b)–(d) define a fixed point in micro objects over $\Omega = (y, A, D, \zeta)$ that does not involve b, while condition (e) feeds back into condition (a) through the bond price. Condition (b) does not depend on b because firms respond to ζ, not to q. The two-stage solution algorithm exploits this structure: Stage 1 solves conditions (b)–(d) on a coarse grid over Ω and stores the results as lookup tables; Stage 2 solves conditions (a) and (e) on the full five-dimensional grid, interpolating the Stage 1 tables. Details are in **??**.

5 Calibration

5.1 Functional Forms

[TODO: Specify utility, production functions if not done in model section]

5.2 Externally Set Parameters

Table 3: Externally Calibrated Parameters

Parameter	Description	Value	Source
β	Discount factor	0.96	Standard
[TODO: Add parameters]			

5.3 Internally Calibrated Parameters

Table 4: Internally Calibrated Parameters

Parameter	Description	Value	Target
[TODO: Add parameters]			

Identification of the default wedges. The four default-wedge intensities $(\nu_\tau, \nu_\omega, \nu_{\chi N}, \nu_{\chi R})$ are internally calibrated using moments that correspond as directly as possible to the empirical margin each wedge is designed to capture. Table 5 summarises the mapping. The working-capital wedge ν_ω is disciplined by an external moment—the observed tightening of trade-finance spreads during sovereign default episodes—rather than by an extensive-margin trade moment, ensuring that ν_ω and ν_τ do not both fit the same short-run trade contraction. [22]

5.4 Model Fit

Wedge contributions. To verify that each wedge earns its place and to assess the relative importance of each channel, table 7 reports the model's fit along six key empirical dimensions under the baseline calibration and under four counterfactual specifications in which each default-wedge intensity is shut down in turn ($\nu = 0$).

The pattern that emerges from table 7 is as follows. [TODO: one paragraph describing what each shutdown reveals once the model is solved. Template: "Shutting down ν_{χ^R} alone accounts

[22] Trade-finance spreads (the excess cost of letters of credit and export credit insurance during default episodes relative to normal periods) average approximately [TODO: X] basis points across our episodes, based on BIS and IMF data. This pins ν_ω before the remaining calibration is run.

Table 5: Identification of default-wedge parameters

Parameter	**Targeted moment**	**Eco**
ν_τ (profitability)	Impact decline in active-link stock; separation rate at $t = 0$	τ ra the sep
ν_ω (working capital)	Trade-finance spread tightening during default episodes [external moment]	ω c gin, the dur
ν_{χ^R} (reactivation)	Post-default re-entry shortfall: cumulative reactivation gap at horizons $h = 3, 5$	χ^R entr this piri
ν_{χ^N} (new matching)	Weak contribution of genuinely new links to recovery	χ c new and mar

Notes: Each row maps one default-wedge parameter to one primary empirical moment and states why that moment is informative about that parameter specifically. The external discipline on ν_ω prevents τ and ω from jointly absorbing the same short-run trade contraction.

Table 6: Targeted Moments: Data vs Model

Moment	Data	Model
[TODO: Add moments]		

Table 7: Wedge-shutdown diagnostics

	Baseline	$\nu_\tau = 0$	$\nu_\omega = 0$	$\nu_{\chi^N} = 0$	$\nu_{\chi^R} = 0$
Impact drop in active links (%)	[TODO]	[TODO]	[TODO]	[TODO]	[TODO]
Survival gap at $h = 3$ (pp)	[TODO]	[TODO]	[TODO]	[TODO]	[TODO]
Re-entry shortfall at $h = 5$ (%)	[TODO]	[TODO]	[TODO]	[TODO]	[TODO]
New-link share of recovery (%)	[TODO]	[TODO]	[TODO]	[TODO]	[TODO]
Incumbent revenue share (%)	[TODO]	[TODO]	[TODO]	[TODO]	[TODO]
Output loss at $h = 10$ (%)	[TODO]	[TODO]	[TODO]	[TODO]	[TODO]
Data target					
Impact drop in active links (%)	[TODO]	—	—	—	—
Re-entry shortfall at $h = 5$ (%)	[TODO]	—	—	—	—

Notes: Each column shuts down one default-wedge intensity ($\nu = 0$) while holding all other parameters at their baseline values. The table shows which empirical dimensions break when each wedge is removed, confirming that no single wedge explains all six facts and that the four-wedge specification is not redundant.

for X% of the re-entry shortfall, confirming that the reactivation wedge is the primary driver of the slow-rebuilding fact. Shutting down ν_τ alone roughly halves the impact drop in active links, confirming that the profitability wedge drives the contemporaneous separation margin. The working-capital wedge ν_ω has negligible effect on the extensive-margin moments but materially affects the intensive-margin response in aggregate revenue, consistent with its identification from trade-finance data rather than from survival moments. Shutting down all four wedges simultaneously ($\nu_\tau = \nu_\omega = \nu_{\chi N} = \nu_{\chi R} = 0$) eliminates the model's ability to match any of the six facts, confirming that the default-wedge mechanism as a whole is the driver of the results."]

6 Results

6.1 Main Results

Figure 11: Model simulations

(a) Debt dynamics

(b) Spread response

[TODO: Main quantitative finding]

6.2 Mechanism

[TODO: Mechanism analysis]

6.3 Robustness

[TODO: Robustness checks]

6.4 Policy Implications

[TODO: Policy experiments or additional results]

7 Conclusion

[TODO: Summarize paper: question, method, main finding]

[TODO: Why does this matter? Policy implications?]

[TODO: What didn't you do? What's next?]

References

Aguiar, Mark and Gita Gopinath, "Defaultable debt, interest rates and the current account," *Journal of International Economics*, 2006, *69* (1), 64–83.

Arellano, Cristina, "Default risk and income fluctuations in emerging economies," *American Economic Review*, 2008, *98* (3), 690–712.

Asonuma, Tamon and Christoph Trebesch, "Sovereign debt restructurings: Preemptive or post-default," *Journal of the European Economic Association*, 2016, *14* (1), 175–214.

Cameron, A. Colin, Jonathon B. Gelbach, and Douglas L. Miller, "Bootstrap-based improvements for inference with clustered errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Chatterjee, Satyajit and Burcu Eyigungor, "A seniority arrangement for sovereign debt," *American Economic Review*, 2012, *102* (6), 2604–2649.

Eaton, Jonathan, Samuel Kortum, and Francis Kramarz, "An anatomy of international trade: Evidence from French firms," *Econometrica*, 2011, *79* (5), 1453–1498.

Fitzgerald, Doireann, Stefanie Haller, and Yaniv Yedid-Levi, "How exporters grow," *Review of Economic Studies*, 2024, *91* (4), 2184–2214.

Gaulier, Guillaume and Soledad Zignago, "BACI: International trade database at the product-level. The 1994-2007 version," Working Paper 2010-23, CEPII 2010.

Jordà, Òscar, "Estimation and inference of impulse responses by local projections," *American Economic Review*, 2005, *95* (1), 161–182.

Martinez, Jose Vicente and Guido Sandleris, "Once bitten, twice shy? Investment and sovereign debt restructuring," *Journal of International Economics*, 2011, *84* (2), 233–245.

Melitz, Marc J., "The impact of trade on intra-industry reallocations and aggregate industry productivity," *Econometrica*, 2003, *71* (6), 1695–1725.

Mendoza, Enrique G. and Vivian Z. Yue, "A general equilibrium model of sovereign default and business cycles," *Quarterly Journal of Economics*, 2012, *127* (2), 889–946.

Rose, Andrew K., "One reason countries pay their debts: Renegotiation and international trade," *Journal of Development Economics*, 2005, *77* (1), 189–206.

A Proofs

Proof of Proposition **??**. [TODO: Proof - pull from notes/model/derivations.tex and clean up] □

B Derivations

B.1 Household First-Order Conditions

[TODO: FOC derivations]

C Data Appendix

[TODO: Data details - pull from notes/data/data-notes.tex]

D Empirical Robustness

This appendix reports the full regression table and robustness checks for the baseline empirical results in Section 3. All figures use the same local-projection specification as the main text (equation 1) unless noted otherwise.

D.1 Alternative Timing Definitions

Figure 12 compares results using three timing conventions: default start (baseline), restructuring end, and midpoint.

Table 8: **Local-projection estimates: variety and volume responses to sovereign default**

	$h=0$	$h=1$	$h=3$	$h=5$	$h=7$	$h=10$
Panel A: Log variety						
$\hat{\beta}_h$	−0.119	−0.123	−0.123	−0.092	−0.105	−0.172
	(0.095)	(0.088)	(0.110)	(0.081)	(0.133)	(0.139)
BS p-value	[0.352]	[0.222]	[0.376]	[0.250]	[0.621]	[0.368]
N	546	527	489	451	413	356
Panel B: Log volume						
$\hat{\beta}_h$	−0.097	−0.121	−0.144	−0.171	−0.148	−0.069
	(0.090)	(0.052)	(0.067)	(0.073)	(0.132)	(0.129)
BS p-value	[0.502]	[0.141]	[0.226]	[0.354]	[0.495]	[0.493]
N	546	527	489	451	413	356

Local-projection estimates of equation (1). Country and year fixed effects included.
Standard errors (parentheses) clustered at country level. Wild bootstrap p-values [brackets] with Rademacher weights (999 reps).
Significance: * $p<0.10$, ** $p<0.05$, *** $p<0.01$ (bootstrap).
19 clusters (4 treated countries, 15 controls). 5 post-default episodes.

Figure 12: Variety impulse responses under alternative event-timing definitions.

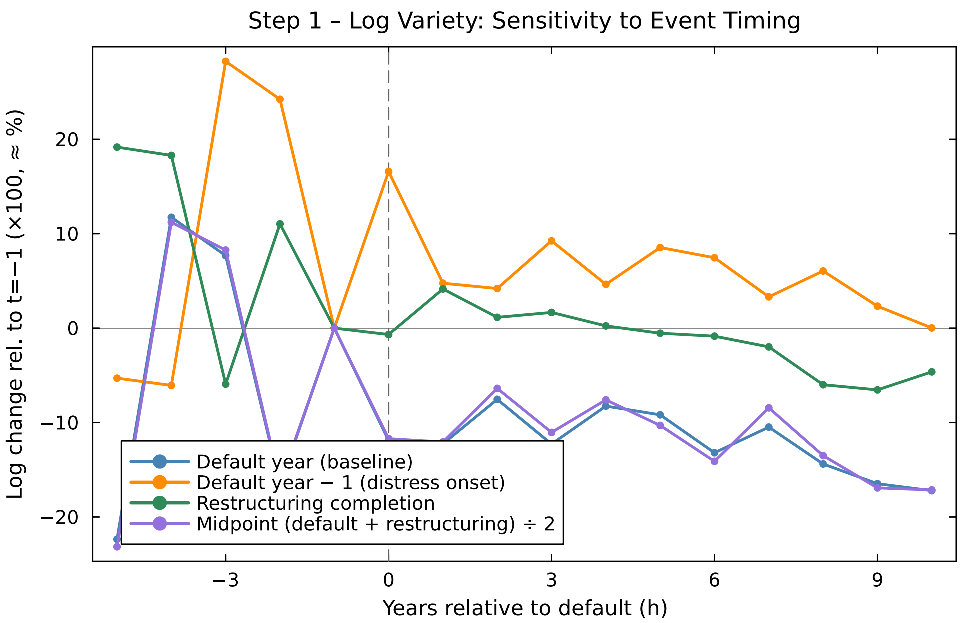

D.2 Specification Comparison

Figure 13 compares the baseline country-plus-year fixed-effects specification with a pre-detrended specification.

Figure 13: Specification robustness for log variety.

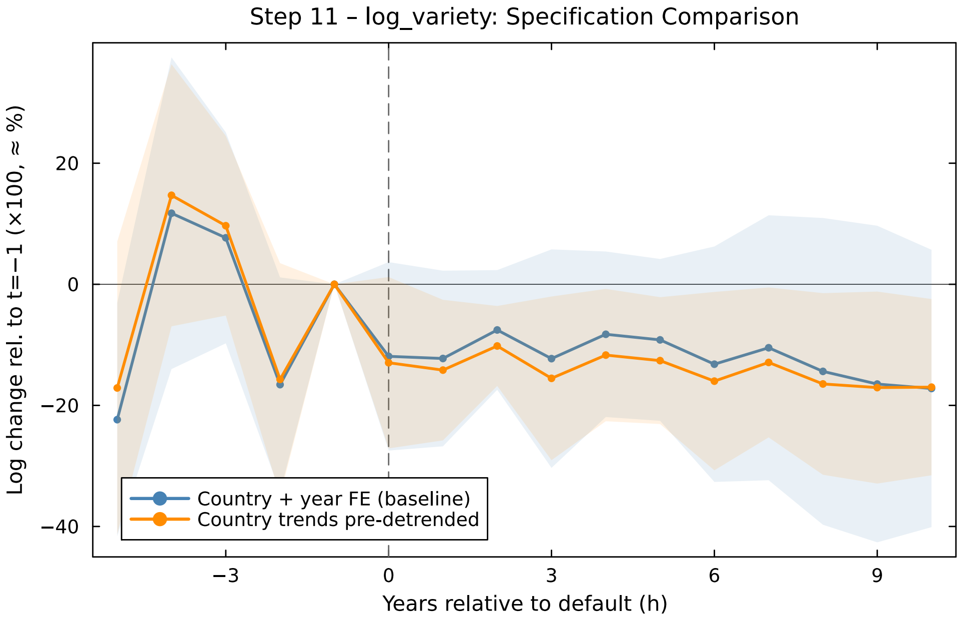

D.3 Including Fuels and Minerals

D.4 Bilateral Variety

D.5 Higher Value Threshold

D.6 Dropping Argentina

D.7 Preemptive Restructurings (Placebo)

D.8 Relationship Dynamics: Heterogeneity

D.9 Additional Relationship Dynamics

D.10 Network Anatomy: Additional Evidence

E Import-Side Evidence

The import-side results provide limited support for an output-cost channel through imported intermediate inputs or capital goods.

F Output Cost of Trade Scarring (Suggestive Evidence)

These exercises are suggestive. The model in Section 4 provides the structural quantification.

G Computational Appendix

[TODO: Computational details]

Figure 14: Variety and volume impulse responses including HS chapters 26–27 (fuels and minerals).

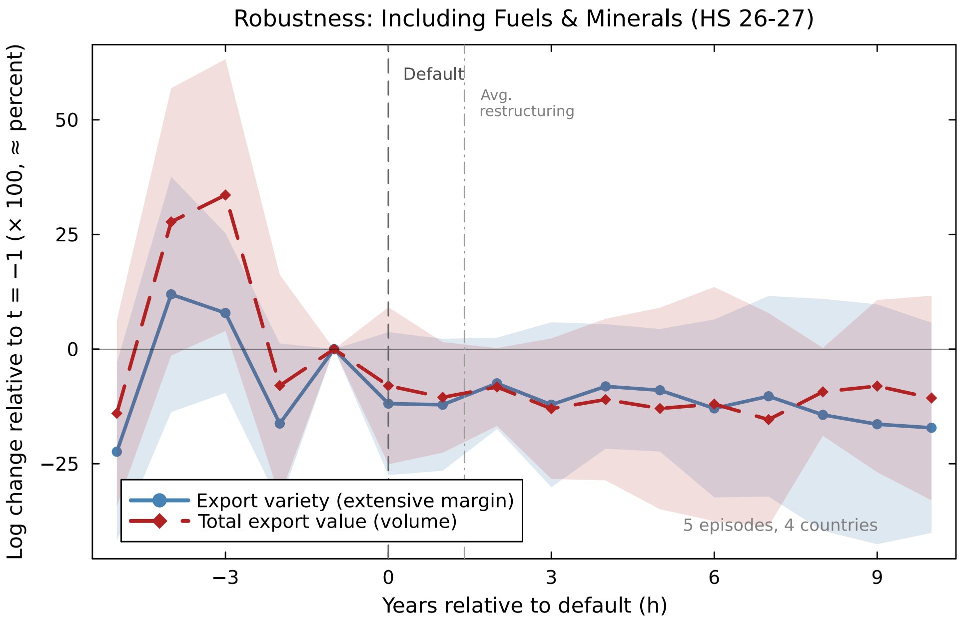

Figure 15: Impulse responses using bilateral variety (product × destination pairs) instead of product count.

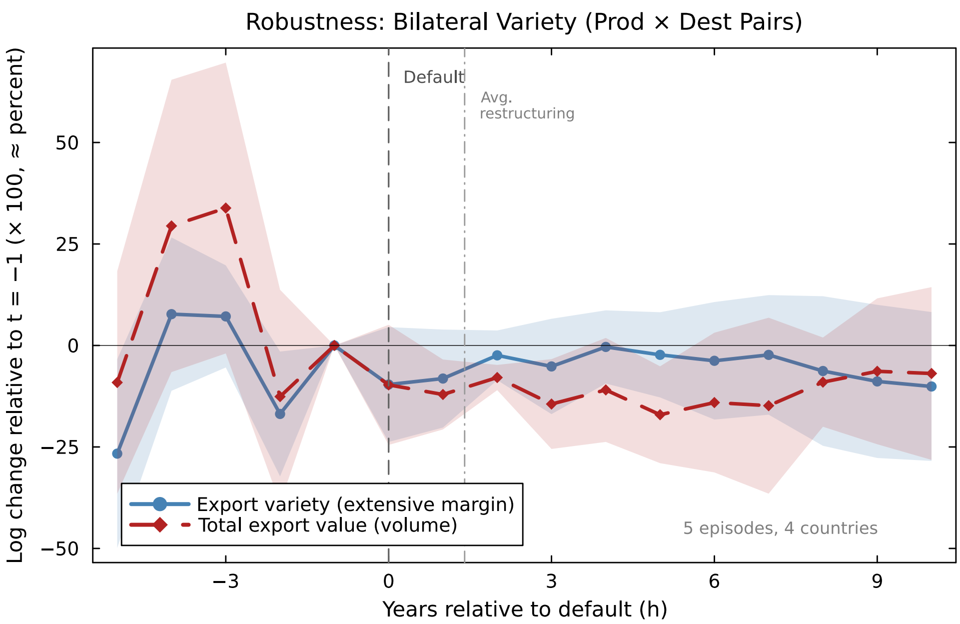

Figure 16: Impulse responses with USD 50,000 minimum threshold for variety counts.

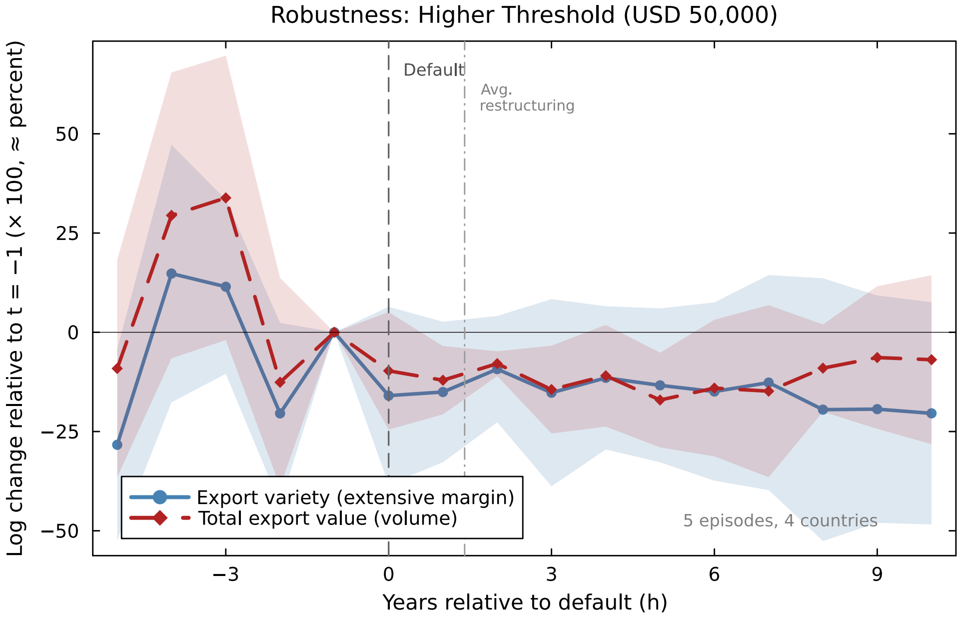

Figure 17: Impulse responses excluding Argentina from the treated sample.

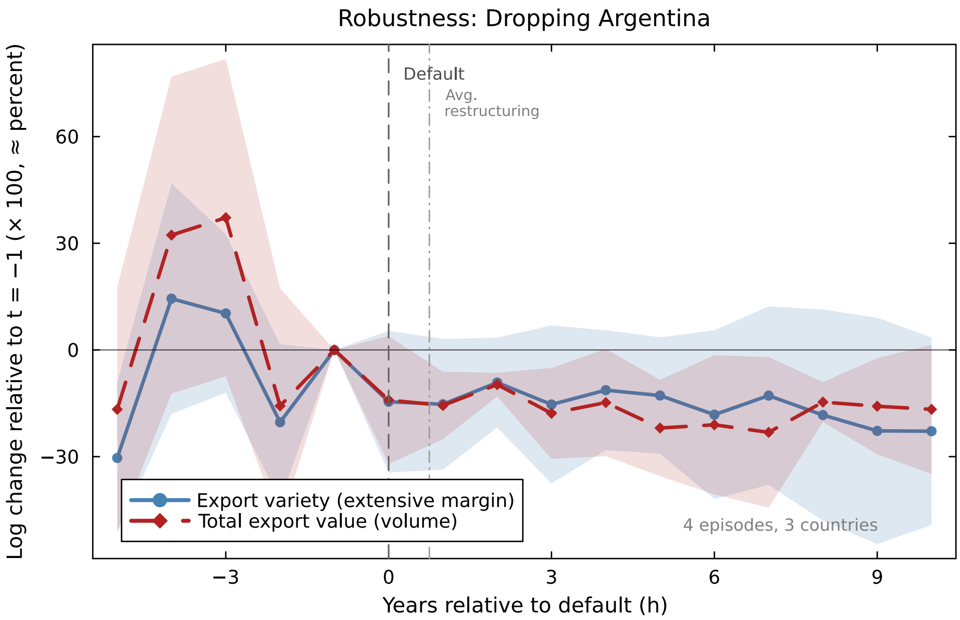

Figure 18: Placebo test: impulse responses for preemptive restructuring episodes. Variety does not fall, consistent with the mechanism being specific to hard default.

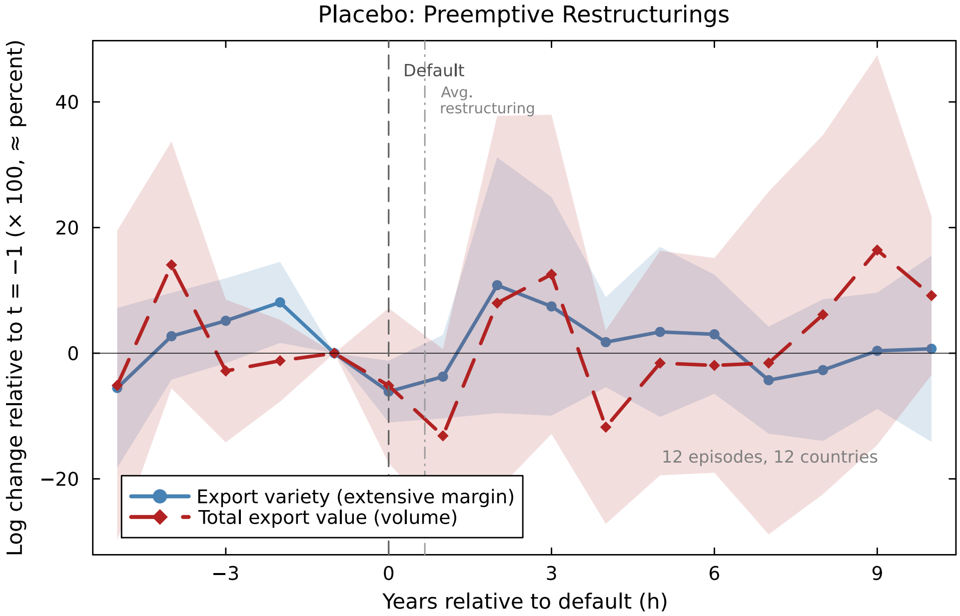

Figure 19: Survival of pre-default export links by pre-default value quartile. Larger relationships are substantially more resilient after default.

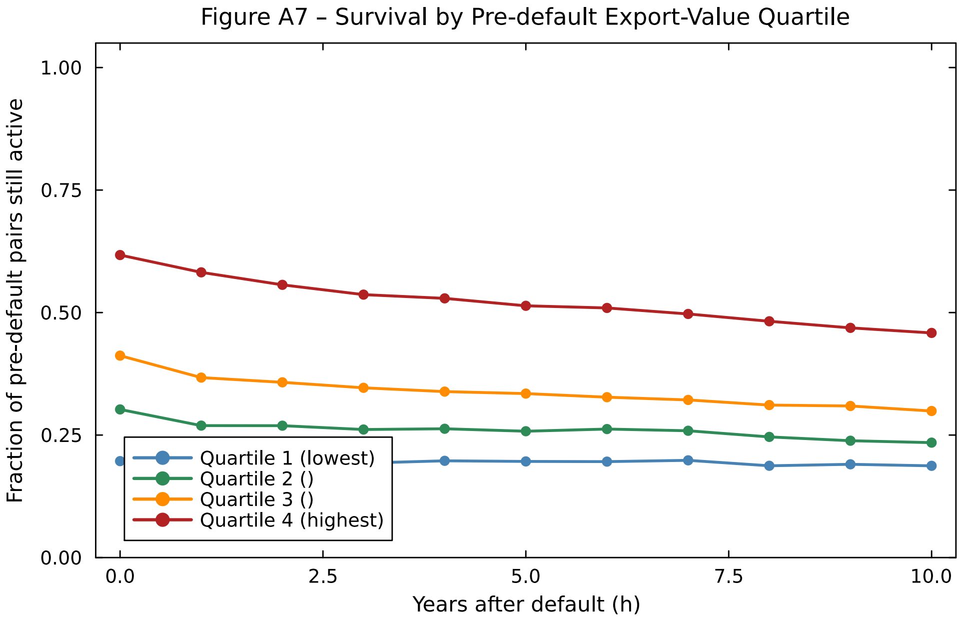

Figure 20: Survival of pre-default export links by link age. Older relationships are more persistent.

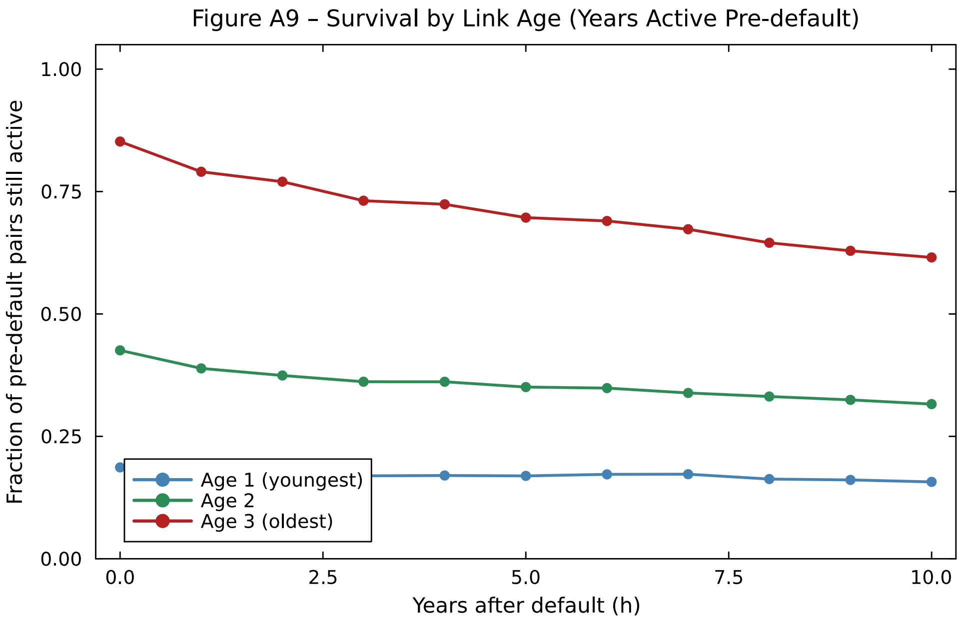

Figure 21: Cumulative re-entry of exited export links by pre-default value quartile. Conditional on exit, re-entry heterogeneity by initial size is weaker than survival heterogeneity.

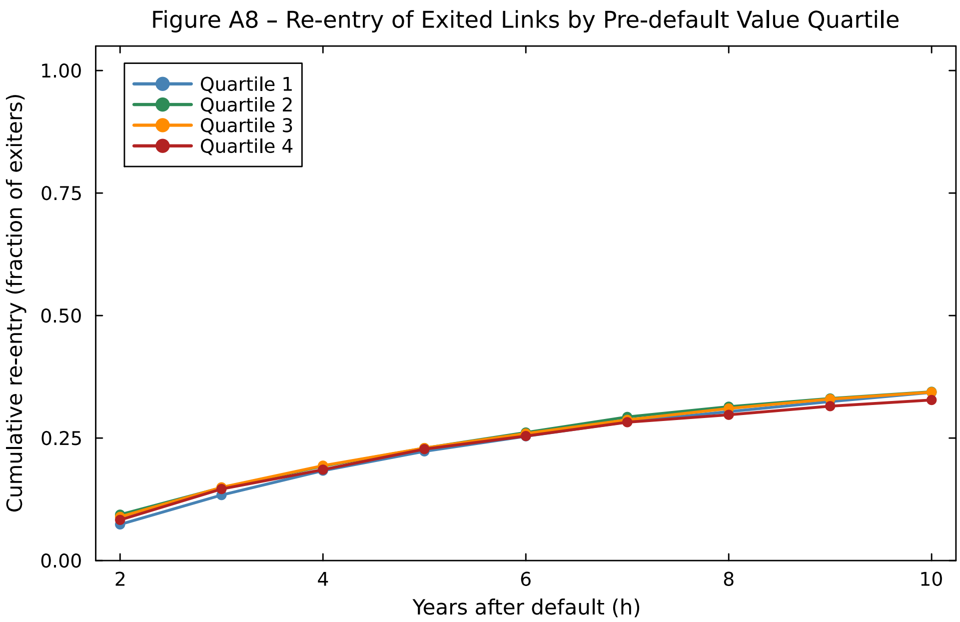

Figure 22: Hazard of exit of pre-default export relationships: treated vs matched controls.

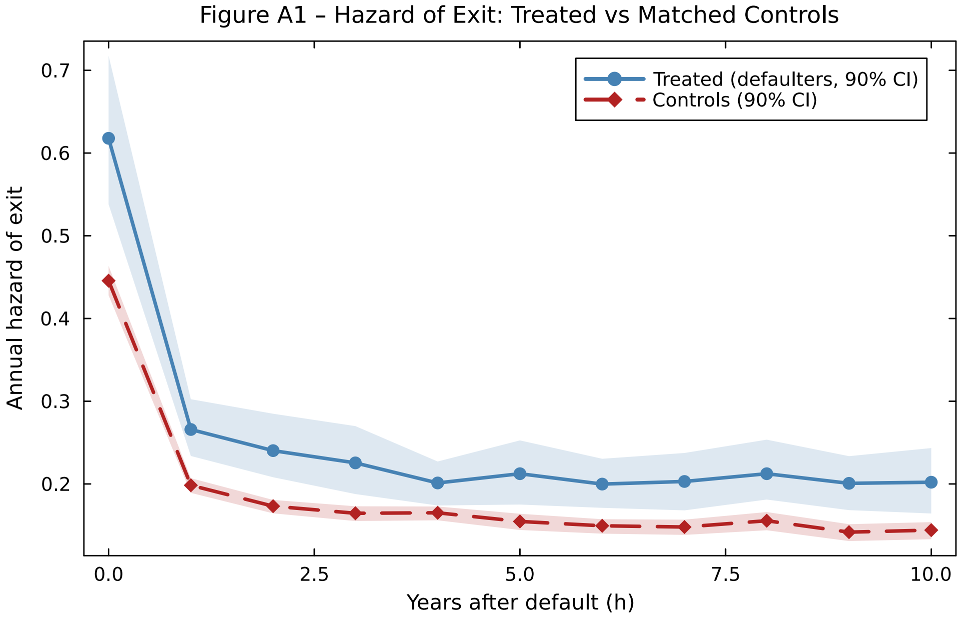

Figure 23: Annual re-entry hazard of dormant export links: treated vs controls.

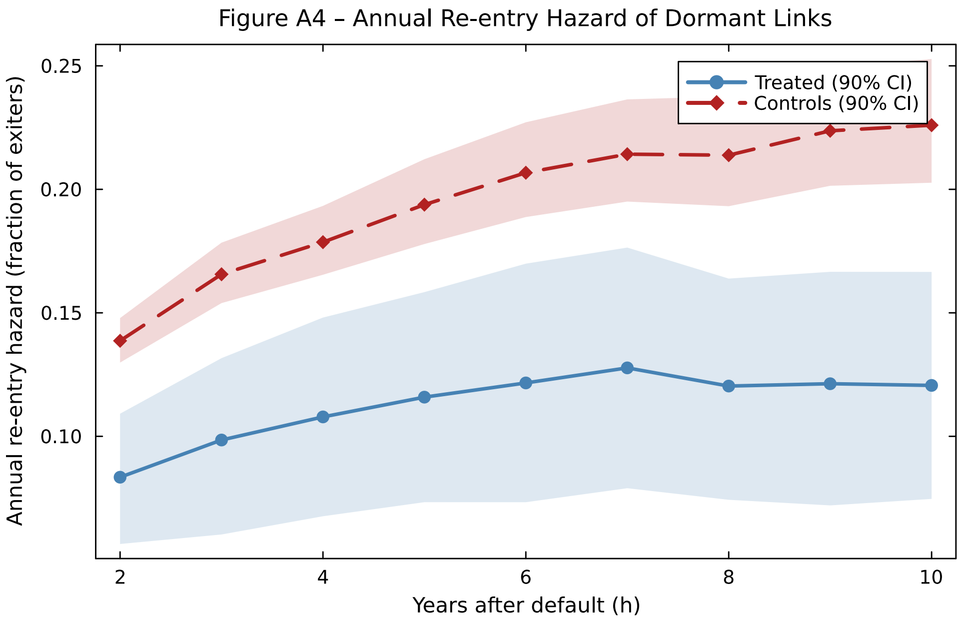

Figure 24: Composition of active export pairs after default in counts. New pairs appear in meaningful numbers even though they do not dominate recovery in value terms.

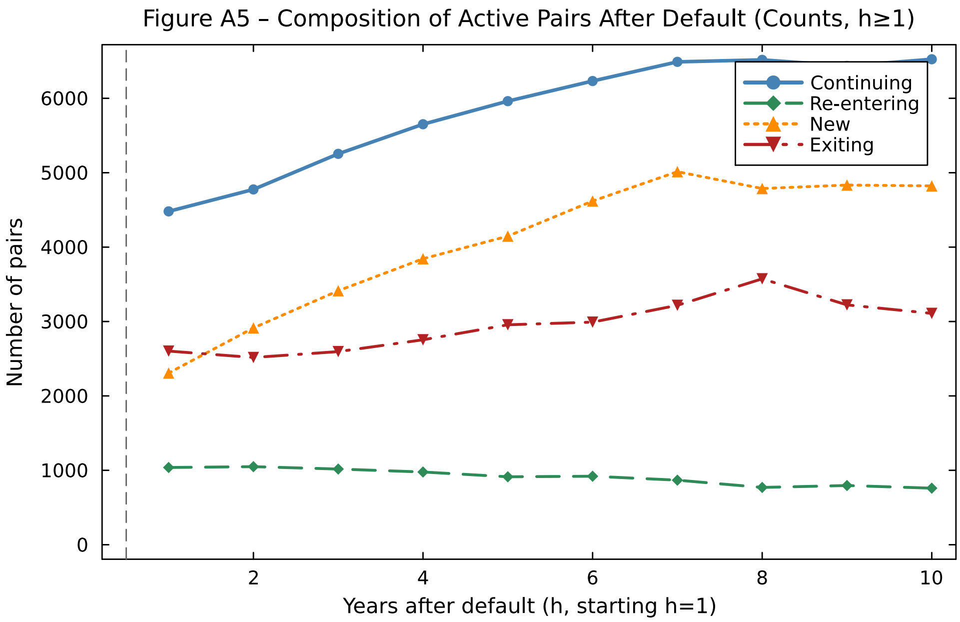

Figure 25: Decomposition of the pair decline into destination breadth and within-destination depth.

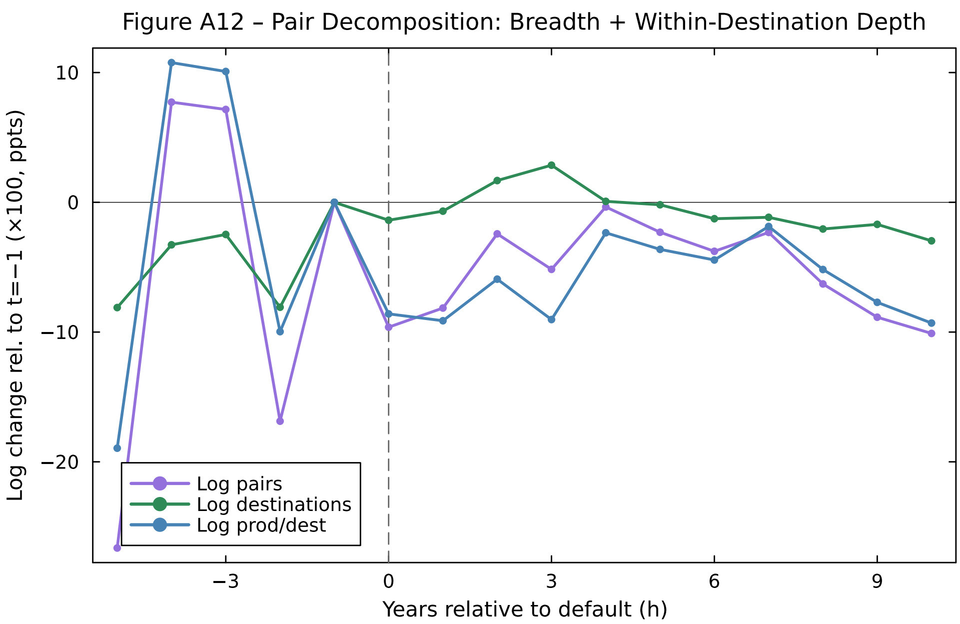

Figure 26: Kernel density of pre-default log export value for links that survive vs. exit by $h = 5$. Exiters are concentrated at the left tail of the value distribution in both treated and control episodes.

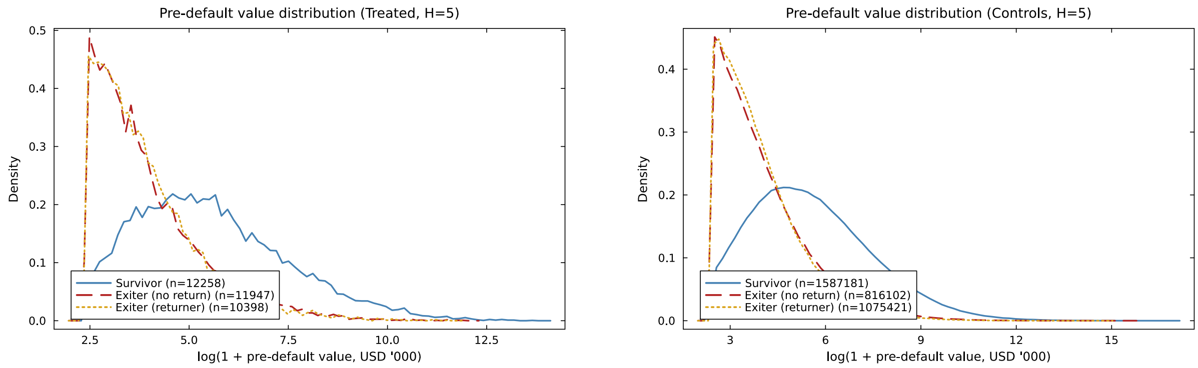

(a) Treated episodes

(b) Control episodes

Figure 27: $X = A \times \bar{x}$ decomposition using product–destination pairs (instead of HS6 product counts). Active pair count drives the persistent decline; per-pair revenue recovers.

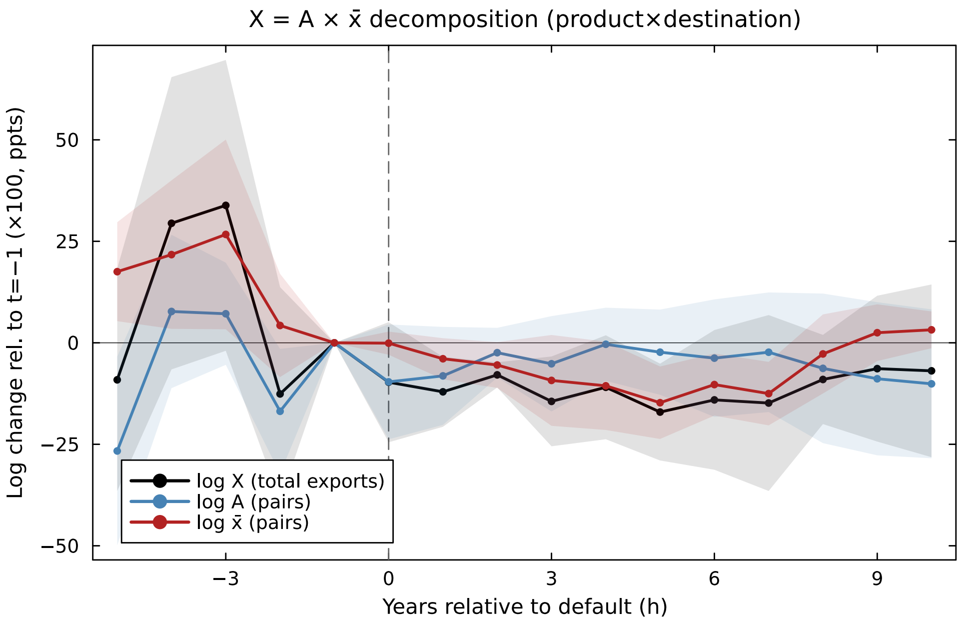

Figure 28: Value shares of surviving incumbents, returners, and new entrants over time. Survivors dominate the value composition throughout the post-default period.

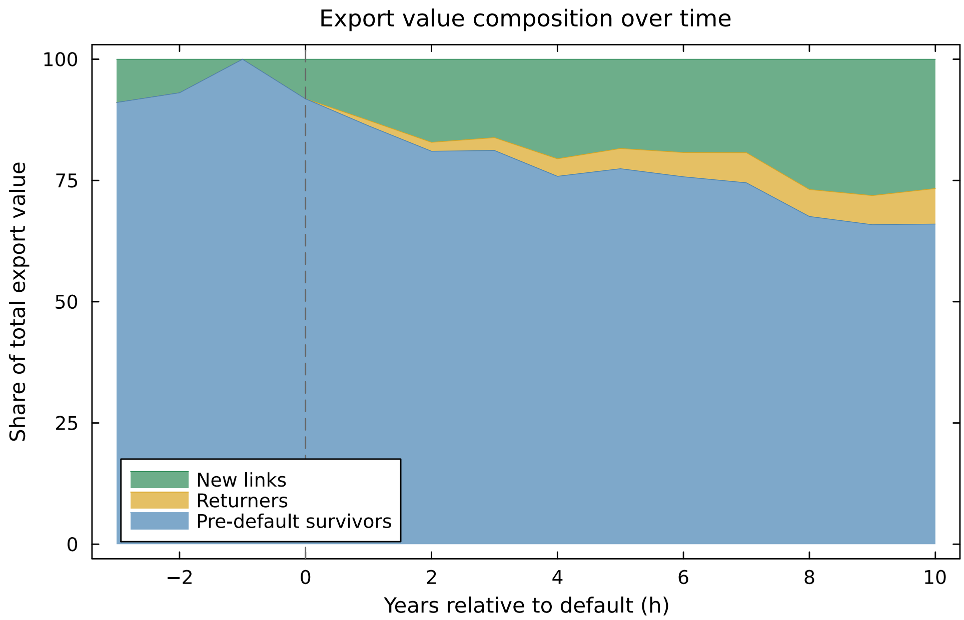

Figure 29: Revenue path of surviving incumbents by pre-default size quartile. Larger incumbents (Q4) recover faster, but the treated–control gap is present across all quartiles.

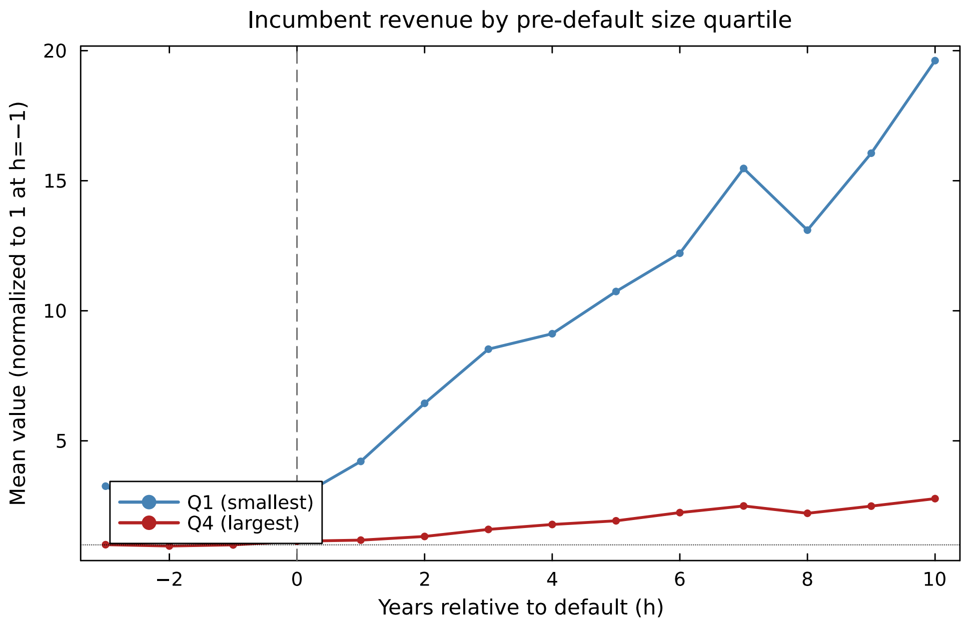

Figure 30: Survival of pre-default import links. Some deterioration relative to controls, but the effect is more modest than on the export side.

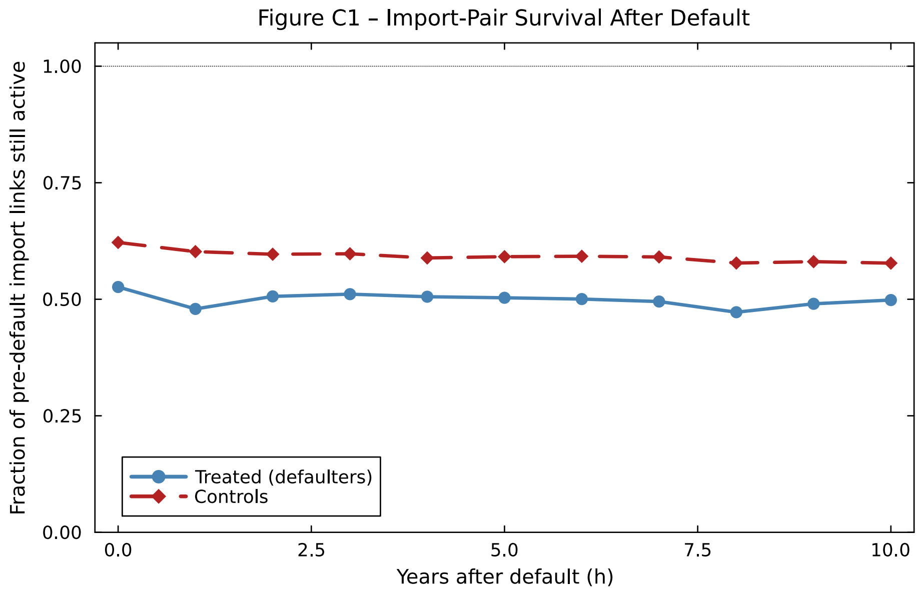

Figure 31: Cumulative re-entry of exited import links.

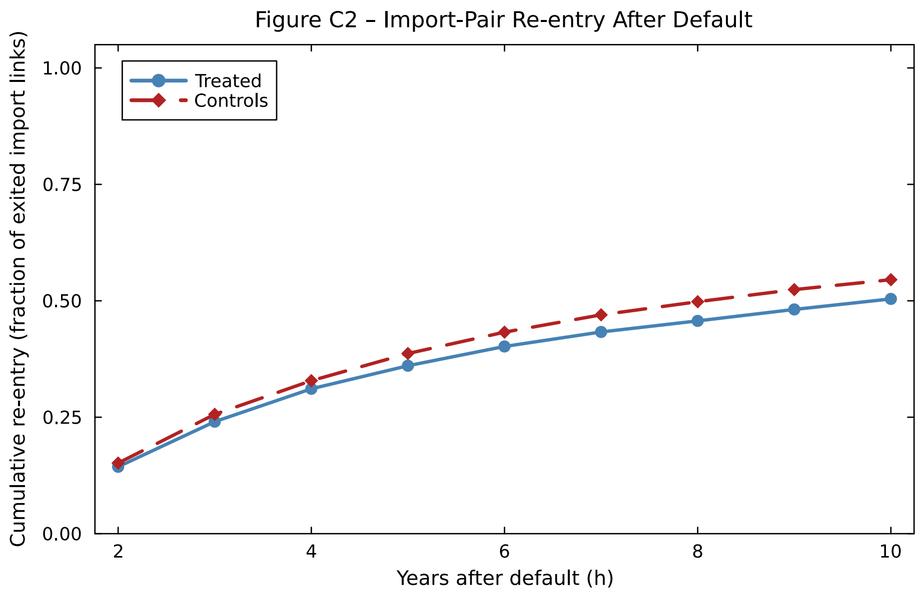

Figure 32: Imported intermediate-good variety after default.

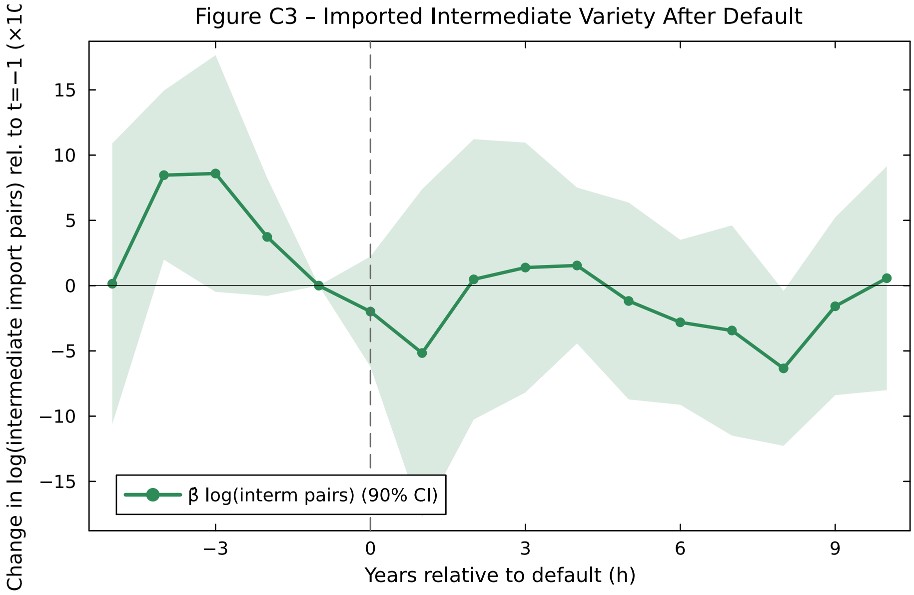

Figure 33: Imported capital-good variety after default.

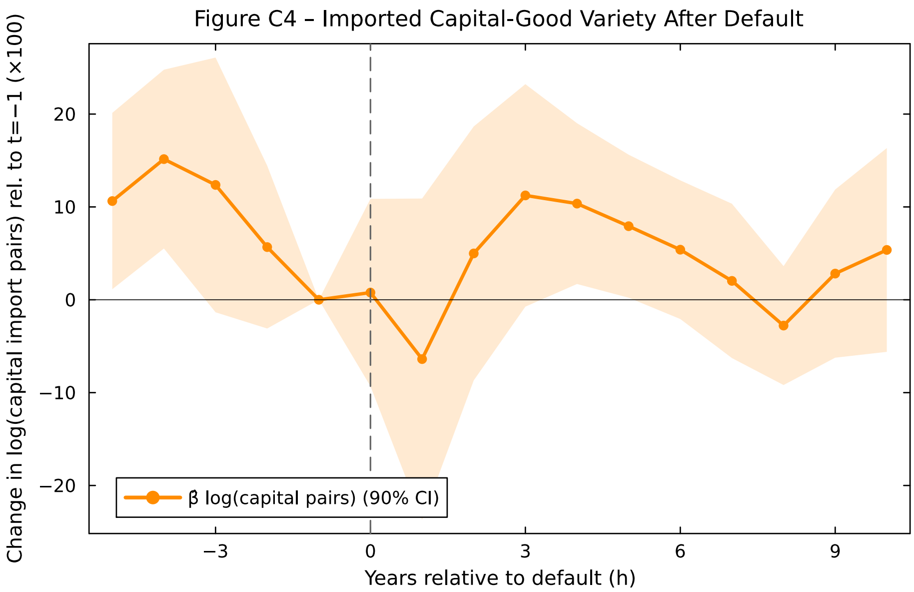

Figure 34: Revenue cost of variety loss (E1).

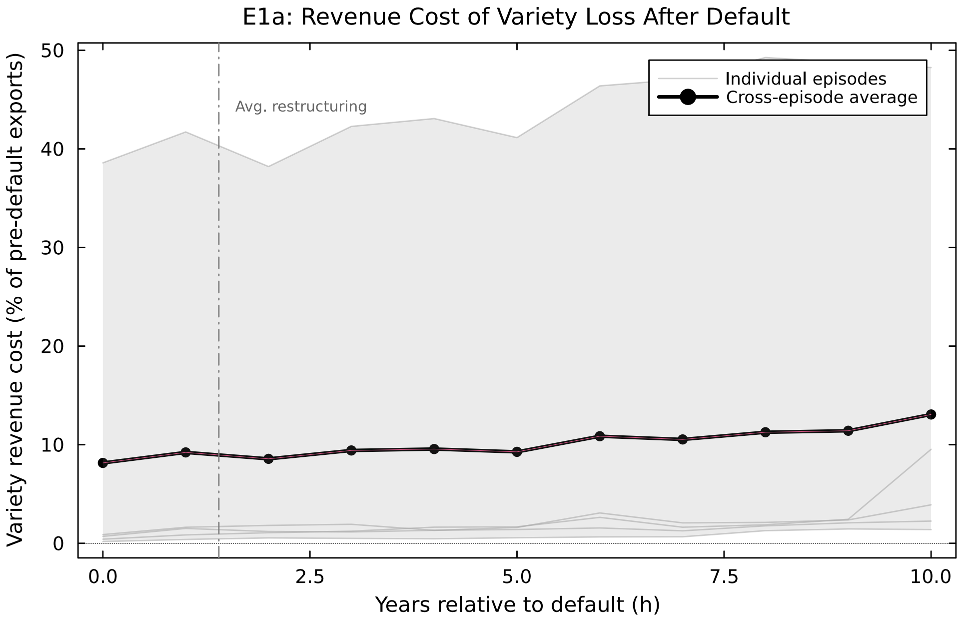

Figure 35: Extensive versus intensive margin decomposition of export recovery (E2).

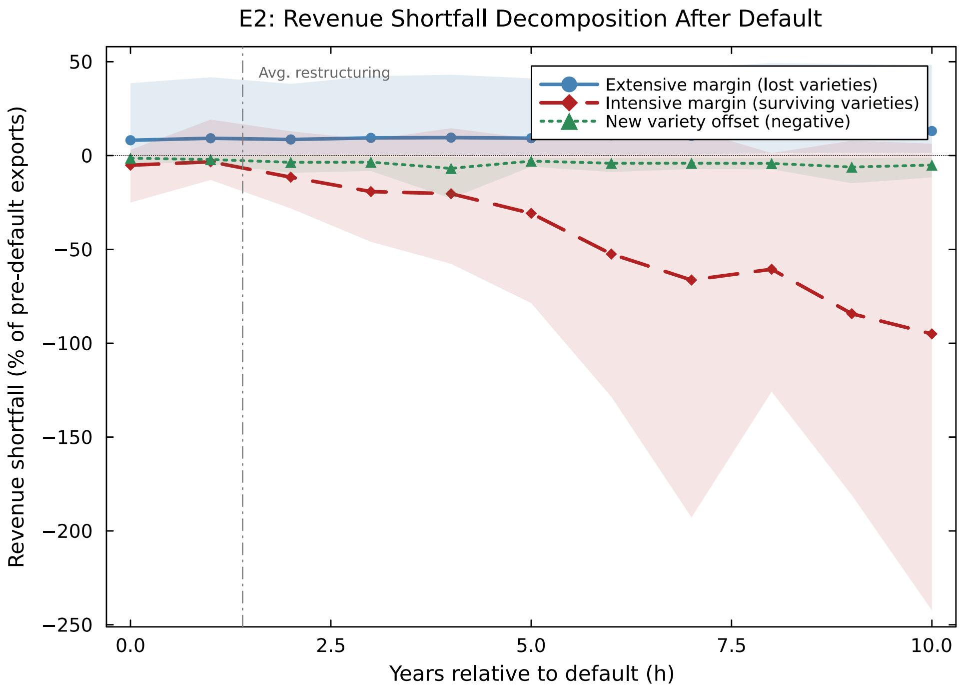

Figure 36: PDV output cost with hysteresis share (E3).

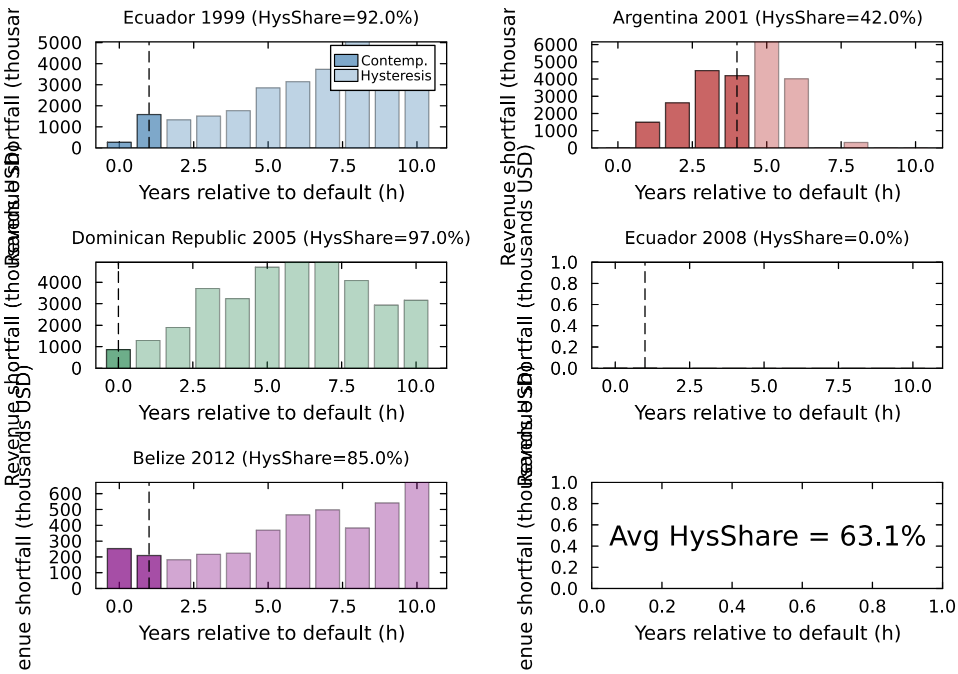

H Additional Model Results

[TODO: Additional model tables and figures]