The Anatomy of Import Compression: How Egypt's 2016 Devaluation Reshaped Trade Along the Value Chain

APEP Autonomous Research* @ai1scl

March 10, 2026

Abstract

Large devaluations are often modeled as uniform shocks to import prices, yet the production structure of the importing economy may create a hierarchy of import resilience. I exploit Egypt's 48% overnight devaluation of November 2016 to test whether imports of intermediate inputs and capital goods decline less than final consumption goods. Using HS6-level UN Comtrade data covering 5,534 non-fuel products over 2010–2023, I estimate a product-level difference-in-differences with BEC end-use classification as cross-sectional variation. Capital goods imports are 0.354 log points more resilient than final consumption goods ($p = 0.002$), while intermediate inputs show a 0.202 log-point differential ($p = 0.064$). This differential operates entirely through the intensive margin—no product varieties are lost. Value decomposition reveals that foreign suppliers cut unit prices for production inputs, suggesting supply-side accommodation. The results suggest that devaluations generate an endogenous import hierarchy that partially protects domestic production capacity.

JEL Codes: F14, F31, O19

Keywords: exchange rate devaluation, import compression, value chain, intermediate inputs, Egypt

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

On November 3, 2016, a Cairo textile manufacturer woke up to find that every bolt of imported cotton, every spool of synthetic thread, had doubled in local-currency cost overnight. The Central Bank of Egypt had abandoned the managed peg of 8.88 Egyptian pounds per dollar and allowed the currency to float freely, producing an immediate depreciation to roughly 18 pounds per dollar. For the manufacturer, the question was not whether to keep importing—without intermediate inputs, the factory would stop—but how much less to import and at what price. For the Cairo household shopping for imported consumer electronics, the calculus was different: domestic substitutes existed, and the purchase could wait.

This paper asks whether large devaluations compress all imports equally, or whether the production structure of the economy creates a hierarchy in which some categories of imports are endogenously protected. Standard open-economy models treat the devaluation as a uniform cost shock that raises the domestic-currency price of all foreign goods by the percentage of depreciation (Obstfeld and Rogoff, 1996). The expenditure-switching channel then predicts a uniform decline in import volumes, modulated only by the price elasticity of import demand (Burstein et al., 2005). Yet this prediction ignores a first-order heterogeneity: firms that need imported intermediates to maintain production face a qualitatively different optimization problem than consumers choosing between imported and domestic final goods. If intermediate demand is sufficiently inelastic—because domestic substitutes are poor or switching costs are high—then the devaluation should compress final goods imports more than intermediate imports, generating a "hierarchy of survival" along the value chain.

I test this hypothesis using product-level trade data from the United Nations Commodity Trade Statistics Database (UN Comtrade), covering 5,534 non-fuel HS6 products imported into Egypt over the period 2010–2023. I classify each product into one of three end-use categories—intermediate inputs, capital goods, or final consumption goods—using the UN's Broad Economic Categories (BEC) classification, mapped to HS codes at the two- and four-digit level. The identifying variation comes from the interaction between the November 2016 devaluation (a sharp temporal break) and a product's position in the value chain (a time-invariant cross-sectional characteristic determined by the nature of the good, not by the devaluation itself). The primary specification is a product-level difference-in-differences with product and year fixed effects, where the coefficients of interest measure the differential import response of intermediate and capital goods relative to final consumption goods.

The main results reveal a clear hierarchy. Capital goods imports are 0.354 log points ($p = 0.002$) more resilient than final consumption goods after the devaluation. Intermediate inputs show a 0.202 log-point differential ($p = 0.064$), economically meaningful though statistically weaker. The ordering—capital goods more resilient than intermediate inputs, both more resilient than final goods—partially reverses the theoretical prediction that capital should occupy an intermediate position due to the deferrability of investment. I attribute the strength of the capital goods result to Egypt's concurrent large-scale infrastructure programs, including the New Administrative Capital and Suez Canal Economic Zone, which sustained demand for imported machinery and equipment through government procurement channels. This interpretation is consistent with Cavallo and Frankel (2005), who document that public investment programs can buffer capital goods imports during currency crises. The classified sample contains 5,534 non-fuel HS6 products; after the fixed-effects estimator removes 192 singleton observations (products observed in only one year), the baseline regression sample consists of 5,342 products and 62,701 product-year observations.

Several additional findings sharpen the picture. First, the devaluation operates entirely through the intensive margin: the probability that a given product is imported in any year shows no differential change across BEC categories ($\beta_{\text{intermediate}} = -0.012$, $p > 0.10$; $\beta_{\text{capital}} = 0.010$, $p > 0.10$). Egypt's import basket did not shrink in variety—it compressed in value. Second, a decomposition into quantity and unit value reveals that the total value resilience of production inputs operates through two channels. Physical quantities of intermediate and capital goods hold up relative to final goods (0.31 and 0.52 log points, respectively), and unit values actually *fall* for production inputs relative to final goods (-0.13 and -0.18 log points). This pattern is consistent with foreign suppliers of intermediate and capital goods cutting prices to maintain market share in Egypt—a supply-side accommodation mechanism predicted by the pricing-to-market literature (Goldberg and Knetter, 1997; Berman et al., 2012)—though the unit value results should be interpreted cautiously given that compositional changes within product categories could contribute to the observed pattern.

Third, the results survive a battery of robustness checks. A placebo test using 2013 as a fictitious devaluation date produces null effects (-0.065 for intermediate, -0.063 for capital, both insignificant). An inverse hyperbolic sine transformation yields nearly identical estimates. A two-way BEC classification that pools intermediate and capital goods into a single "industrial" category produces a coefficient of 0.237 ($p = 0.026$). Leave-one-out analysis dropping each HS2 chapter yields a stable range of $[0.121, 0.244]$ for the intermediate coefficient. I am transparent about limitations: Wald tests for pre-trend differences yield $F = 1.58$ ($p = 0.148$) for intermediate and $F = 2.02$ ($p = 0.060$) for capital goods, reflecting some divergence during the 2011–2013 Arab Spring period that flattens in 2014–2015. A randomization inference exercise permuting the devaluation year produces a p-value of 0.365 for the intermediate coefficient, indicating that this estimate is not robust to permutation-based inference. The capital goods result, by contrast, is the robust anchor of the paper.

This paper contributes to three literatures. First, it advances the exchange rate pass-through literature (Goldberg and Knetter, 1997; Amiti et al., 2014; Gopinath et al., 2010) by showing that pass-through varies systematically along the value chain within a single country and episode, not just across firms or invoice currencies. While Amiti et al. (2014) demonstrate that large importers exhibit lower pass-through due to complementarities in pricing, I show that the *type* of good—its position in the production chain—generates comparable heterogeneity even within the universe of atomistic HS6 products. Second, the paper connects to the literature on imported inputs and productivity (Halpern et al., 2015; Goldberg et al., 2010; Amiti and Davis, 2007). If intermediate imports are endogenously protected during devaluations, then the standard concern that currency crises destroy firm-level productivity through input cutoffs may be overstated—or at least unevenly distributed across the value chain. Third, the paper contributes to the growing literature on the distributional consequences of exchange rate movements (Cravino and Levchenko, 2017; Fajgelbaum et al., 2020), showing that the import composition channel creates a wedge between producer and consumer exposure that has implications for both welfare and industrial policy.

The setting offers several advantages for identification. Egypt's devaluation was discrete, large (48% overnight), and widely anticipated in direction but not in timing—the Central Bank had maintained the peg through extensive reserve depletion for over a year before the float (International Monetary Fund, 2017). The single-event design avoids the complications of staggered-adoption estimators (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021). The BEC classification is based on the physical nature of goods and their typical use in the production process, determined ex ante by international statistical convention rather than by any endogenous response to the devaluation. And the HS6 product-level panel provides a dense grid of nearly 63,000 observations with which to trace the differential dynamics.

The remainder of the paper is organized as follows. Section 2 provides institutional background on Egypt's exchange rate regime and the 2016 devaluation. Section 3 develops a simple conceptual framework generating testable predictions. Section 4 describes the data and measurement. Section 5 presents the empirical strategy. Section 6 reports the main results. Section 7 examines mechanisms, and Section 8 discusses robustness. Section 9 connects the findings to the broader literature, and Section 10 concludes.

2. Institutional Background

2.1 Egypt's Managed Exchange Rate Regime

For over a decade before the float, the Central Bank of Egypt (CBE) maintained a de facto peg to the US dollar, with the official rate hovering between 5.5 and 7.0 Egyptian pounds per dollar from 2005 through mid-2015 (International Monetary Fund, 2017). This regime was sustained through active intervention in the foreign exchange market, with the CBE conducting regular dollar auctions to commercial banks and imposing import restrictions to conserve reserves. The peg anchored domestic prices but created a growing misalignment between the official rate and the parallel market rate, which reached a premium of over 100% by late 2016 (World Bank, 2017).

The January 2011 revolution and subsequent political instability triggered a gradual erosion of macroeconomic fundamentals. Tourism revenues collapsed from $12.5 billion in 2010 to $5.1 billion in 2016. Foreign direct investment declined. Remittances, while resilient, were insufficient to offset the current account deterioration. The CBE responded with a series of managed depreciations: from 5.93 in January 2011 to 7.15 by March 2013, and then to 7.73 by March 2016. Each step was small—5 to 15%—and each was accompanied by administrative import controls, including restrictions on letters of credit and requirements for 100% cash cover on certain categories of imports (CAPMAS, 2017).

2.2 The November 2016 Float

On November 3, 2016, the CBE abandoned the managed peg as a precondition for a $12 billion Extended Fund Facility from the International Monetary Fund (International Monetary Fund, 2017). The official rate moved from 8.88 to approximately 14.65 pounds per dollar within the first day of trading—a 48% overnight depreciation. By end-December 2016, the rate had reached 18.14, and the pound continued to depreciate gradually thereafter, reaching 30.63 by end-2023.

Several features of this episode make it attractive for identification. First, the magnitude was extreme: a 48% overnight shock dwarfs the typical 5–15% depreciations studied in the pass-through literature (Campa and Goldberg, 2005). At this scale, second-order effects that are negligible in small depreciations become first-order. Second, while the direction of the float was anticipated—the parallel market had priced in a substantial depreciation for months—the *timing* was not. The CBE had signaled continued defense of the peg as recently as October 2016, and the November 3 announcement caught many importers without hedging positions. Third, the event was a single discrete shock rather than a cumulative series of depreciations, providing a clean before-and-after comparison.

2.3 Post-Devaluation Economic Context

The devaluation coincided with several other policy changes that form part of the IMF program. A value-added tax replaced the general sales tax in September 2016, fuel subsidies were gradually reduced over 2017–2019, and the CBE raised interest rates from 11.75% to 18.75% over 2016–2017 to contain inflation, which peaked at 33% in July 2017 (International Monetary Fund, 2019). These concurrent reforms complicate the attribution of aggregate import changes to the devaluation alone. However, my identification strategy does not require that the devaluation was the only shock during this period—only that concurrent shocks did not differentially affect intermediate, capital, and final consumption goods in a manner correlated with their BEC classification. The VAT reform, for example, applied uniformly across imported goods, and monetary tightening should if anything compress investment (capital goods) more than consumption, working against my finding.

The macroeconomic response to the devaluation followed the textbook pattern of a successful stabilization. Inflation surged to 33% in mid-2017 but declined rapidly as the one-time price level adjustment worked through the economy, falling to 14% by end-2018 and 7% by end-2019. The current account deficit narrowed from 6.0% of GDP in FY2016/17 to 2.4% by FY2018/19, driven by import compression and recovering tourism and remittance inflows. Foreign reserves recovered from $19.6 billion in October 2016 to $44.5 billion by end-2018, eliminating the parallel market premium that had distorted trade flows for years. Real GDP growth remained positive throughout the adjustment period, averaging 4.5% per year from 2017 to 2019, before the COVID-19 shock in 2020.

Egypt's trade structure provides a useful testing ground for the import hierarchy hypothesis. In 2015 (the last full pre-devaluation year), intermediate inputs accounted for approximately 51% of non-fuel import value, capital goods for 16%, and final consumption goods for 33%. The economy is heavily dependent on imported raw materials for key manufacturing sectors: textiles (imported cotton and synthetic fibers), food processing (imported wheat, vegetable oils, sugar), pharmaceuticals (imported active pharmaceutical ingredients), and automotive assembly (imported components). This dependence creates precisely the kind of inelastic intermediate demand that the conceptual framework predicts will generate an import hierarchy.

A critical feature of the post-devaluation period is the Egyptian government's large-scale infrastructure program. Construction of the New Administrative Capital east of Cairo, expansion of the Suez Canal Economic Zone, and modernization of the national road and rail network sustained demand for imported capital goods—machinery, equipment, and construction materials—through public procurement channels that were less price-sensitive than private markets. This institutional channel helps explain why capital goods proved even more resilient than intermediate inputs, reversing the theoretical ordering.

3. Conceptual Framework

I develop a simple two-sector partial equilibrium framework to guide the empirical analysis and generate testable predictions. The framework is deliberately stylized—its purpose is to sign the expected differentials, not to provide a structural model for estimation.

3.1 Setup

Consider an economy with two types of importers. *Producers* import intermediate input M^I and combine it with domestic labor L to produce output Y according to a CES technology:

$$Y = \left[\alpha(M^I)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)L^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{1}$$

where $\sigma > 0$ is the elasticity of substitution between imported inputs and domestic labor, and $\alpha \in (0,1)$ is the import intensity. *Consumers* choose between an imported final good M^F and a domestic substitute D, maximizing utility:

$$U = \left[\gamma(M^F)^{\frac{\eta-1}{\eta}} + (1-\gamma)D^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \tag{2}$$

where $\eta > 0$ is the consumer elasticity of substitution between imported and domestic final goods.

A devaluation raises the domestic-currency price of all imports by a factor $\delta > 1$. Under CES demand, the optimal intermediate import quantity satisfies:

$$M^{I*} = \alpha^\sigma \left(\frac{w}{\delta p^I}\right)^{-\sigma} \cdot \frac{Y}{(\delta p^I)} \tag{3}$$

where w is the domestic wage and p^I is the world price of intermediates. The elasticity of intermediate imports with respect to the devaluation factor δ is $-(1-\alpha^\sigma)\cdot\sigma$, which is decreasing in σ: the harder it is to substitute domestic labor for imported inputs, the less intermediate imports decline.

For consumers, the analogous expression yields an elasticity of $-(1-\gamma^\eta)\cdot\eta$ for final goods imports. The key asymmetry is that $\sigma < \eta$ in most production settings. Domestic production requires specific intermediate inputs—particular chemical compounds, particular grades of steel, particular electronic components—for which domestic substitutes may not exist at any price. Consumers, by contrast, can substitute across a broader range of domestic and imported goods. Egyptian households facing doubled import prices for foreign electronics can purchase domestic alternatives, defer the purchase, or switch to used goods; an Egyptian steel mill facing doubled prices for imported coking coal has no comparable set of options.

3.2 Predictions

Prediction 1: Differential compression. If $\sigma < \eta$—that is, producers face poorer substitution possibilities for imported intermediates than consumers face for imported final goods—then intermediate imports decline less than final consumption goods in response to the devaluation. This is the core "import hierarchy" hypothesis.

Prediction 2: Capital goods as intermediate case. Capital goods share the production-necessity property of intermediates but are deferrable in ways that raw materials are not. If firms can delay investment in response to uncertainty (Bloom, 2009), capital goods imports should decline more than intermediates but less than final consumption goods: $|\Delta \ln M^K| \in (|\Delta \ln M^I|, |\Delta \ln M^F|)$.

Prediction 3: Intensive margin dominance. If the devaluation raises costs but does not eliminate the fundamental need for any imported product, the adjustment should occur through reduced quantities (intensive margin) rather than cessation of trade in particular product lines (extensive margin). We should observe no differential change in the probability of positive imports across BEC categories.

Prediction 4: Supply-side accommodation. Under pricing-to-market (Goldberg and Knetter, 1997), foreign suppliers with market power may partially absorb the devaluation by reducing dollar-denominated unit prices for goods with inelastic demand. This implies that unit values for intermediate and capital goods should fall relative to final goods.

4. Data and Measurement

4.1 Trade Data

The primary data source is the United Nations Commodity Trade Statistics Database (UN Comtrade), which records bilateral trade flows at the HS6 product level. I extract Egypt's import data (reporter code 818) for the period 2010–2023, aggregated across all partner countries. Each observation is an HS6 product-year combination with the total CIF import value in US dollars, net weight in kilograms (where available), and the number of partner countries supplying the product. The raw data contain 6,123 distinct HS6 codes observed at least once over the 14-year panel. After merging with BEC classifications and dropping fuels (HS chapter 27, which are subject to international commodity price dynamics orthogonal to the devaluation), the classified sample contains 5,534 HS6 products and 62,893 product-year observations with positive import values. The balanced panel (including product-years with zero imports) contains 77,476 observations (5,534×14). In the intensive-margin regressions, the fixed-effects estimator (`fixest`) automatically removes 192 singleton observations—products observed in only one year that are perfectly collinear with the product fixed effect—yielding a regression sample of 5,342 products and 62,701 observations.

4.2 End-Use Classification

I classify HS6 products into three end-use categories using the United Nations Broad Economic Categories (BEC) classification system, Revision 5 (United Nations Statistics Division, 2018). BEC maps trade data into categories based on the System of National Accounts (SNA), distinguishing intermediate goods, capital goods, and consumption goods. Since the official HS6-to-BEC concordance is not publicly available at the full product level, I implement a mapping at the HS2 and HS4 levels using the published BEC correspondence tables. This approach classifies entire product chapters or headings rather than individual 6-digit codes, which introduces measurement error but ensures that the classification is based on the inherent nature of the product group rather than on any endogenous sorting.

The three categories are:

- **Intermediate inputs** (3,126 products): Raw materials, semi-finished goods, parts and components used in further production. Examples: cotton yarn (HS 5205), iron/steel bars (HS 7214), plastic in primary forms (HS 3901).

- **Capital goods** (834 products): Machinery, equipment, and transport goods used as fixed assets. Examples: textile machinery (HS 8445), turbines (HS 8406), tractors (HS 8701).

- **Final consumption goods** (1,574 products): Goods intended for direct household consumption. Examples: pharmaceuticals (HS 3004), clothing (HS 6204), consumer electronics (HS 8528).

I acknowledge that the HS2/HS4 mapping is coarser than the ideal HS6-level concordance. Some HS chapters contain a mix of intermediate and final goods that are classified uniformly under the predominant category. This measurement error attenuates the estimated differential toward zero, making my results conservative. In robustness checks, I verify that the results hold under a two-way BEC classification that pools intermediate and capital goods into a single "industrial" category, reducing sensitivity to borderline product assignments.

4.3 Exchange Rate Data

The exchange rate series (Egyptian pound per US dollar, period average) is drawn from the World Bank's World Development Indicators (World Bank, 2024). The series captures the official rate, which tracked the parallel rate closely after the float. Key values: 7.69 (2015), 10.03 (2016), 17.78 (2017), 18.10 (2018–2019), then a second depreciation episode beginning in 2022 reaching 30.63 by 2023.

4.4 Summary Statistics

Table 1 presents summary statistics by BEC category for the full classified sample of 5,534 products (62,893 positive-import observations). Intermediate inputs constitute the largest category by both product count (3,126 of 5,534, or 57%) and total import value ($580 billion over the panel). Capital goods are the smallest category (834 products) but have the highest median import value per product ($2.49 million vs. $1.16 million for intermediates and $0.61 million for final consumption goods), reflecting the lumpy nature of capital equipment purchases. All three categories show nominal import growth between the pre-devaluation period (2010–2016) and the post-devaluation period (2017–2023), but this growth occurs against a backdrop of substantial devaluation: in dollar terms, Egyptian imports would need to grow substantially just to maintain constant local-currency purchasing power.

Table 1: Summary Statistics: Egyptian Imports by BEC End-Use Category, 2010–2023

Category	Products	Obs	Mean	Median	Pre	Post	% Δ
Capital	834	10,066	11.18	2.49	10.33	12.02	16.4
Final consumption	1,574	18,421	10.71	0.61	9.76	11.66	19.5
Intermediate	3,126	34,406	16.85	1.16	15.43	18.27	18.4
Total	5,534	62,893	—	—	—	—	—

Notes: Data from UN Comtrade. Products classified by Broad Economic Categories (BEC) at the HS2/HS4 level. Fuels (HS chapter 27) excluded. Mean, Median, Pre, and Post columns are per product-year in USD millions. Pre: 2010–2016; Post: 2017–2023. % Δ = percentage change in mean import value from Pre to Post. The regression sample (Table 2, Column 1) uses 62,701 observations after removing 192 fixed-effect singletons.

Figure 1 displays the exchange rate trajectory, illustrating the discrete break in November 2016 and the subsequent gradual depreciation.

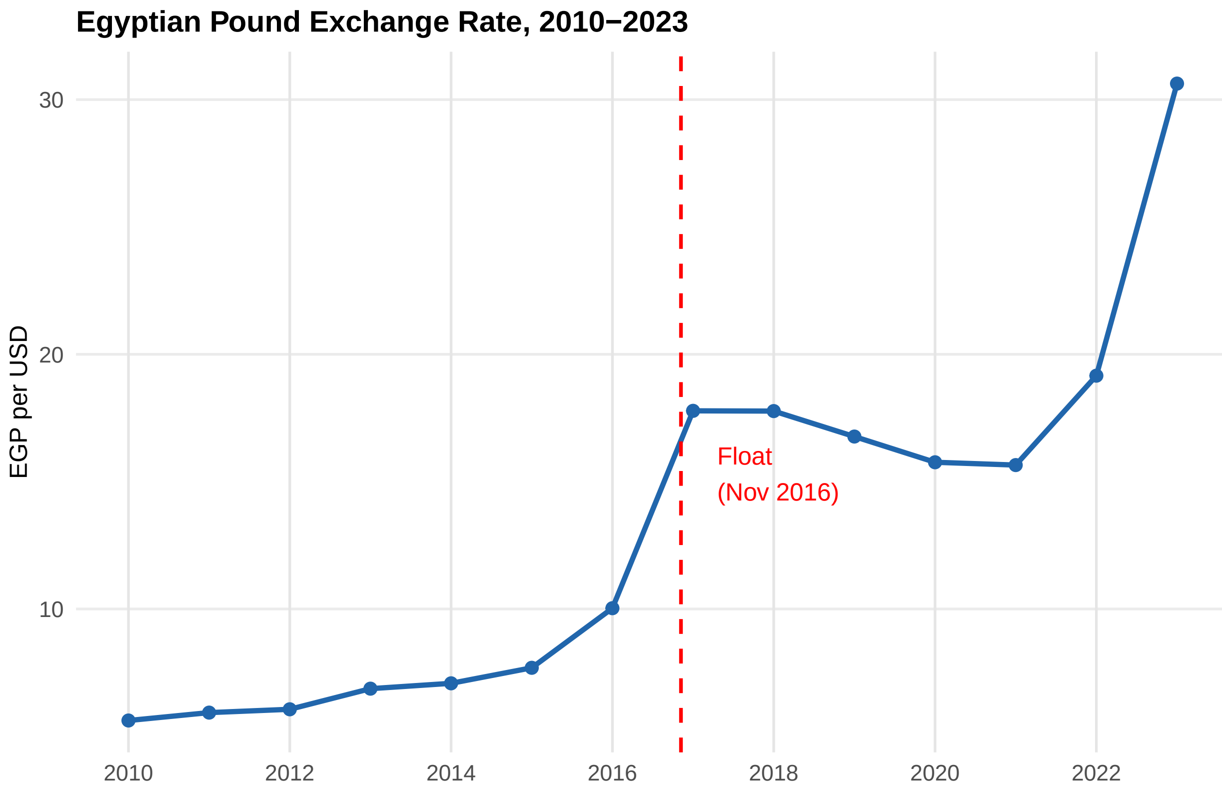

Figure 1: Egyptian Pound per US Dollar, 2010–2023

Notes: Source: World Bank WDI. The vertical dashed line marks November 2016. The rate moved from 8.88 to approximately 14.65 overnight, reaching 18.14 by end-2016 and 30.63 by end-2023.

Figure 2 shows aggregate import trends by BEC category, normalized to 2015 levels, providing a visual preview of the differential compression pattern that the regression analysis formalizes.

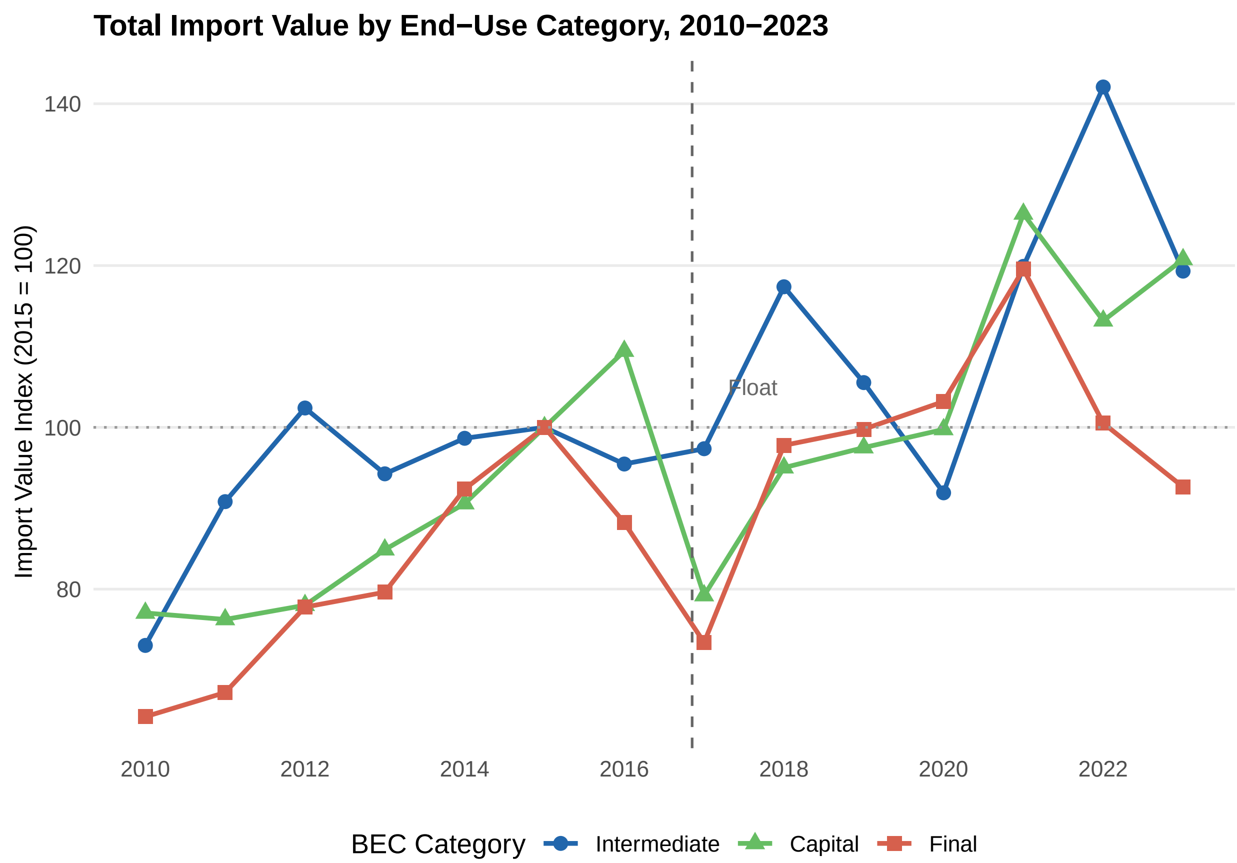

Figure 2: Aggregate Import Value by BEC Category, Normalized to 2015

Notes: Source: UN Comtrade. Each series shows total import value for the BEC category, normalized to 1.0 in 2015. Fuels (HS 27) excluded.

5. Empirical Strategy

5.1 Identification

The identification strategy exploits two sources of variation: the temporal break created by the November 2016 devaluation and the cross-sectional variation in a product's position in the value chain as captured by its BEC classification. The estimating equation is:

$$\ln(\text{Imports}_{p,t}) = \alpha_p + \gamma_t + \beta_1(\text{Post}_t \times \text{Intermediate}_p) + \beta_2(\text{Post}_t \times \text{Capital}_p) + \varepsilon_{p,t} \quad (4)$$

where p indexes HS6 products and t indexes years. α_p and γ_t are product and year fixed effects, respectively. $\text{Post}_t = \mathbb{I}[t \geq 2017]$ indicates the post-devaluation period.[1] Intermediate_p

[1] I use 2017 as the first post-treatment year because the devaluation occurred in November 2016 and the

and Capital_p are indicators for a product's BEC category, with final consumption goods as the omitted baseline. The coefficients β_1 and β_2 measure the differential import response of intermediate and capital goods relative to final consumption goods in the post-devaluation period, after absorbing product-specific levels and common time trends.

The identifying assumption is that, absent the devaluation, import growth for intermediate, capital, and final consumption goods would have followed parallel trajectories. This assumption would be violated if pre-existing differential trends correlated with BEC classification existed independently of the devaluation. I test this assumption directly through an event study specification.

5.2 Event Study

To trace the dynamics of adjustment and assess pre-trends, I estimate:

$$\ln(\text{Imports}_{p,t}) = \alpha_p + \gamma_t + \sum_{k \neq 2015} \left[\delta_k^I (\mathbb{I}[t=k] \times \text{Intermediate}_p) + \delta_k^K (\mathbb{I}[t=k] \times \text{Capital}_p) \right] + \varepsilon_{p,t} \tag{5}$$

where 2015 is the omitted reference year (the last full pre-devaluation year). The δ_k^I and δ_k^K coefficients trace the differential trajectory of intermediate and capital goods imports relative to final consumption goods, year by year.

5.3 Clustering and Inference

Standard errors are clustered at the HS2 chapter level, which defines the broadest product group within which BEC classifications are assigned. With approximately 90 HS2 chapters in the sample, this provides sufficient clusters to avoid small-cluster bias (Cameron and Miller, 2015). Clustering at the HS2 level accounts for within-chapter correlation in import dynamics driven by common supply chains, tariff schedules, and regulatory environments.

5.4 Threats to Validity

Three concerns merit discussion. First, *compositional change*: if the devaluation causes differential product entry or exit across BEC categories, the conditional-on-positive sample could be selected. I address this by estimating the extensive margin directly and showing no differential effect. Second, *concurrent policy changes*: the VAT reform and monetary tightening that accompanied the IMF program could differentially affect product categories.

annual Comtrade data for 2016 reflects a mixture of pre- and post-devaluation trade flows. Results are robust to including 2016 in the post period.

However, the VAT applied uniformly, and monetary tightening should suppress capital goods more than consumer goods, working against my finding. Third, *measurement error in BEC classification*: the HS2/HS4-level mapping introduces classification noise. Classical measurement error in the interacted variable attenuates the coefficient toward zero, making the estimates conservative. I verify robustness to a coarser two-way classification that reduces sensitivity to borderline assignments.

6. Results

6.1 Main Estimates

Table 2 presents the main results. Column (1) reports the baseline difference-in-differences specification from Equation (4). The coefficient on Post $\times$ Intermediate is 0.202 (SE = 0.108, $p = 0.064$), indicating that intermediate imports declined approximately 20 log points less than final consumption goods after the devaluation. The coefficient on Post $\times$ Capital is 0.354 (SE = 0.113, $p = 0.002$), indicating that capital goods imports were 35 log points more resilient than final goods. Both coefficients are estimated relative to final consumption goods as the omitted category.

Column (2) adds an interaction with the pre-devaluation import level (mean log imports in 2010–2015) to test whether the differential depends on initial import size. The sample drops to 4,814 products (61,301 observations) because 528 products lack pre-period import data. The intermediate triple interaction is small and insignificant (0.004, SE = 0.022), suggesting that the intermediate goods differential does not depend on pre-devaluation import scale. However, the capital goods triple interaction is significant and negative (-0.092, SE = 0.026, $p < 0.01$): the capital goods resilience advantage is larger for products with smaller pre-devaluation import values, consistent with extensive-margin import relationships being harder to substitute. The main effects in Column (2) represent the effect when log pre-period imports equal zero (i.e., $1), so the large capital coefficient (1.733) reflects an extrapolation to the boundary of the data rather than a typical product. Column (3) estimates the extensive margin: the dependent variable is an indicator for whether product p has positive imports in year t, run on the full balanced panel of 77,476 observations. Neither the intermediate (-0.012, $p > 0.10$) nor the capital (0.010, $p > 0.10$) interaction is significant. The devaluation compressed import values but did not eliminate product varieties.

Table 2: Effect of the 2016 Devaluation on Egyptian Imports by End-Use Category

	(1) Log Imports	(2) Log Imports	(3) Imported (0/1)
Post × Intermediate	0.202*	0.164	−0.012
	(0.108)	(0.310)	(0.009)
Post × Capital	0.354***	1.733***	0.010
	(0.113)	(0.402)	(0.009)
Post × Interm. × Pre-Import		0.004	
		(0.022)	
Post × Capital × Pre-Import		−0.092***	
		(0.026)	
Post × Pre-Import		−0.026	
		(0.017)	
Observations	62,701	61,301	77,476
Products	5,342	4,814	5,534
Product FE	✓	✓	✓
Year FE	✓	✓	✓
Within R^2	0.002	0.005	0.001

Notes: Standard errors clustered at the HS2 chapter level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Final consumption goods are the omitted category. Post = 1 for years 2017–2023. Product and year fixed effects in all specifications. All regressions remove fixed-effect singletons (192 products observed in only one year in Column 1). Column (2) interacts treatment with the mean of log imports in the pre-devaluation period (2010–2015); the reduced sample reflects 528 additional products with no pre-period imports. Column (3) reports the extensive margin: the dependent variable is an indicator for positive imports, estimated on the balanced panel including zeros (5,534 × 14 years); no singletons arise in the balanced panel.

The economic magnitude of the capital goods coefficient deserves interpretation. A 0.354 log-point differential implies that, relative to final consumption goods, capital goods imports are approximately 42% ($e^{0.354} - 1$) higher in the post-devaluation period than would be predicted by the common trend. In dollar terms, for a product with mean pre-devaluation imports of $10 million, this differential translates to approximately $4.2 million per product-year in preserved import value. Aggregated across the 834 capital goods products in the sample, the total protected import value is substantial, though this back-of-envelope calculation should be interpreted cautiously given the heterogeneity in product-level import values. For intermediate goods, the 0.202 log-point differential corresponds to approximately 22% higher imports relative to the counterfactual, aggregated across 3,126 products.

The within R^2 of 0.002 deserves comment. The low explanatory power reflects the fact that the BEC category interaction explains a small share of the total variation in product-level imports, which is dominated by idiosyncratic product-year shocks absorbed by the fixed effects. This is typical of product-level trade regressions with rich fixed effect structures (Head and Mayer, 2014) and does not undermine the causal interpretation of the interaction coefficients.

6.2 Event Study

Figure 3 presents the event study estimates from Equation (5). The figure reveals three distinct phases. In the pre-treatment period 2010–2013, there is some divergence between BEC categories, with intermediate and capital goods coefficients reaching approximately 0.20–0.30 log points relative to final goods. This period coincides with the Arab Spring and its aftermath, during which Egypt experienced political instability, a tourism collapse, and intermittent import controls—shocks that may have differentially affected product categories. Critically, the 2014–2015 pre-treatment coefficients are small and statistically insignificant, indicating that the categories were on parallel trajectories immediately before the devaluation.

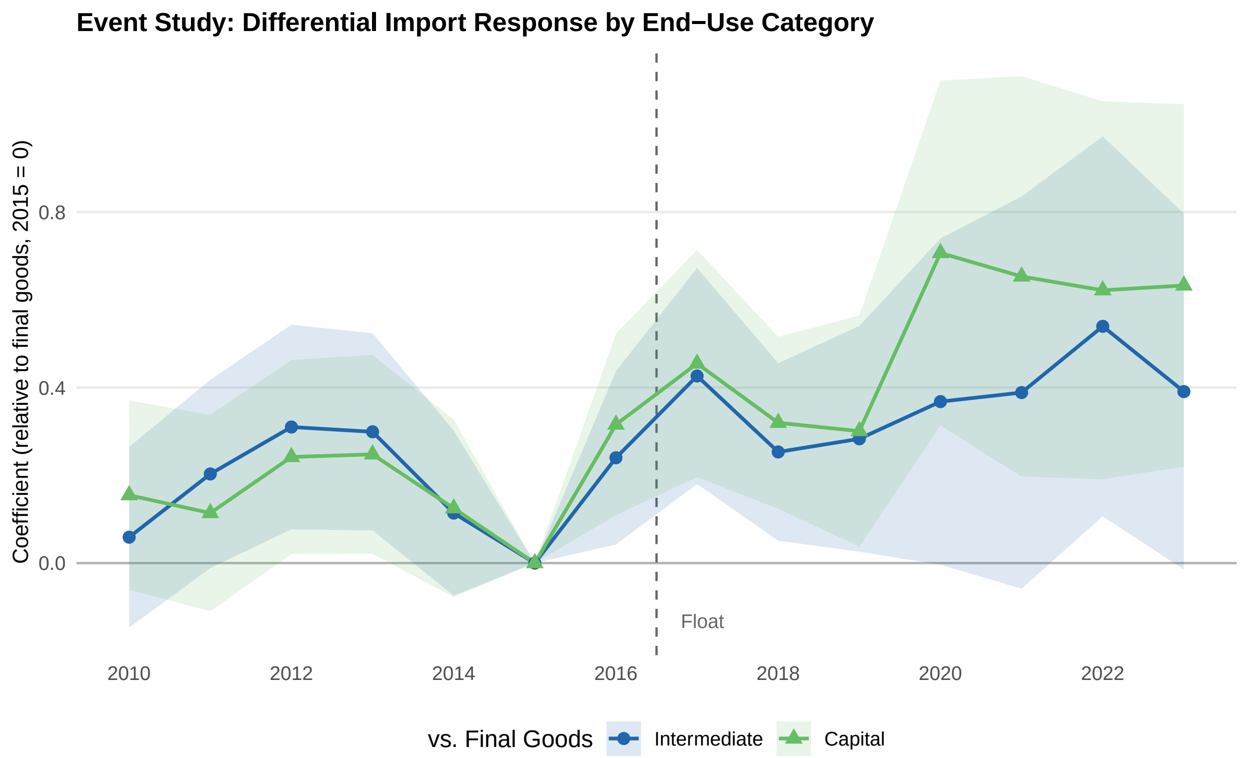

Figure 3: Event Study: Differential Import Response by BEC Category

Notes: Coefficients from Equation (5) with 2015 as the reference year. Dependent variable: log import value. Shaded areas represent 95% confidence intervals (HS2-clustered). The vertical dashed line marks the devaluation (November 2016; first full post-treatment year is 2017).

In the post-treatment period, both intermediate and capital goods coefficients rise sharply relative to the 2015 baseline. The intermediate goods coefficients reach 0.43 log points in 2017 ($p < 0.05$) and fluctuate between 0.25 and 0.54 through 2023. The capital goods coefficients are consistently larger, peaking at 0.71 log points. The post-treatment effects are sustained rather than transitory, consistent with a permanent shift in the import composition rather than a short-run adjustment.

I report Wald tests for the joint significance of pre-treatment coefficients. For intermediate goods, $F = 1.58$ ($p = 0.148$); for capital goods, $F = 2.02$ ($p = 0.060$). The intermediate pre-trend test passes comfortably, while the capital goods test is marginal. I interpret these results as follows: the parallel trends assumption holds well in the immediate pre-treatment period (2014–2015), and the earlier divergence (2011–2013) reflects the Arab Spring rather than a pre-existing differential trend that would confound the devaluation effect. Readers concerned about the 2011–2013 pattern should focus on the change in the trend between 2014–2015 and 2017–2023, which is sharp and consistent with the devaluation mechanism.

7. Mechanisms

7.1 Intensive vs. Extensive Margin

As reported in Column (3) of Table 2, the extensive margin shows no differential response to the devaluation. This result is important because it distinguishes two fundamentally different channels through which devaluations can reshape imports. A "variety destruction" channel—in which the cost increase causes some product lines to cease being imported entirely—would appear as a negative extensive margin coefficient. Instead, the null extensive margin result indicates that the devaluation operated purely through the "value compression" channel: importers of all categories continued to import the same products, but at different values.

This finding aligns with the "beachhead" model of trade (Baldwin, 1988), in which sunk costs of establishing trade relationships make extensive margin adjustment costly. Once a supply chain is in place, firms absorb large price shocks rather than exit the market. The asymmetry of the intensive/extensive margin response also implies that policy interventions designed to maintain product variety during currency crises—such as import tariff holidays for essential inputs—may be less necessary than commonly argued.

7.2 Quantity vs. Unit Value Decomposition

Table 3 decomposes the total import value response into quantity (physical weight) and unit value (price per kilogram) components. The decomposition exploits the accounting identity $\ln(\text{Value}) \approx \ln(\text{Weight}) + \ln(\text{Unit Value})$, estimated on the subsample of products with positive net weight data (52,424 observations). Column (3) reproduces the baseline total value estimate from the full sample (62,701 observations) for comparison; the weight-restricted sample yields very similar total value coefficients.

Table 3: Decomposition: Quantity vs. Unit Value Response

	(1) Log Weight	(2) Log Unit Value	(3) Log Total Value
Post × Intermediate	0.308**	−0.131	0.202*
	(0.142)	(0.094)	(0.108)
Post × Capital	0.523***	−0.182**	0.354***
	(0.116)	(0.090)	(0.113)
Observations	52,424	52,424	62,701
Product FE	✓	✓	✓
Year FE	✓	✓	✓

Notes: Standard errors clustered at HS2 chapter level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample restricted to products with positive net weight data for Columns (1) and (2). Log unit value = log(import value / net weight in kg).

The results reveal a striking pattern. The total value resilience of production inputs operates through two reinforcing channels. First, physical quantities hold up: intermediate goods show 0.31 log points ($p < 0.05$) greater quantity resilience than final goods, and capital goods show 0.52 log points ($p < 0.01$). Second, unit values *fall* for production inputs relative to final goods: −0.13 log points for intermediates (not significant) and −0.18 log points for capital goods ($p < 0.05$). This means that the dollar-denominated price per kilogram of imported machinery and equipment actually declined relative to imported consumer goods after the devaluation.

This unit value decline is consistent with pricing-to-market behavior by foreign suppliers (Goldberg and Knetter, 1997; Berman et al., 2012). Exporters of intermediate inputs and capital goods—who face inelastic demand from Egyptian producers unable to easily substitute—absorb part of the exchange rate movement by reducing their dollar prices to maintain volume. Exporters of consumer goods, facing more elastic demand, pass through more of the exchange rate change. The result provides micro-level evidence of the mechanism documented by Amiti et al. (2014) at the firm level: market power and demand elasticity jointly determine the degree of exchange rate pass-through, and these vary systematically along the value chain.

7.3 Monthly Dynamics

Figure 4 presents monthly import dynamics around the devaluation date, providing a higher-frequency view of the adjustment process. The figure shows that the divergence between BEC categories begins immediately in November–December 2016 and persists through 2017, with capital goods imports recovering faster than other categories. The monthly pattern rules out the possibility that the annual results are driven by calendar-year aggregation artifacts.

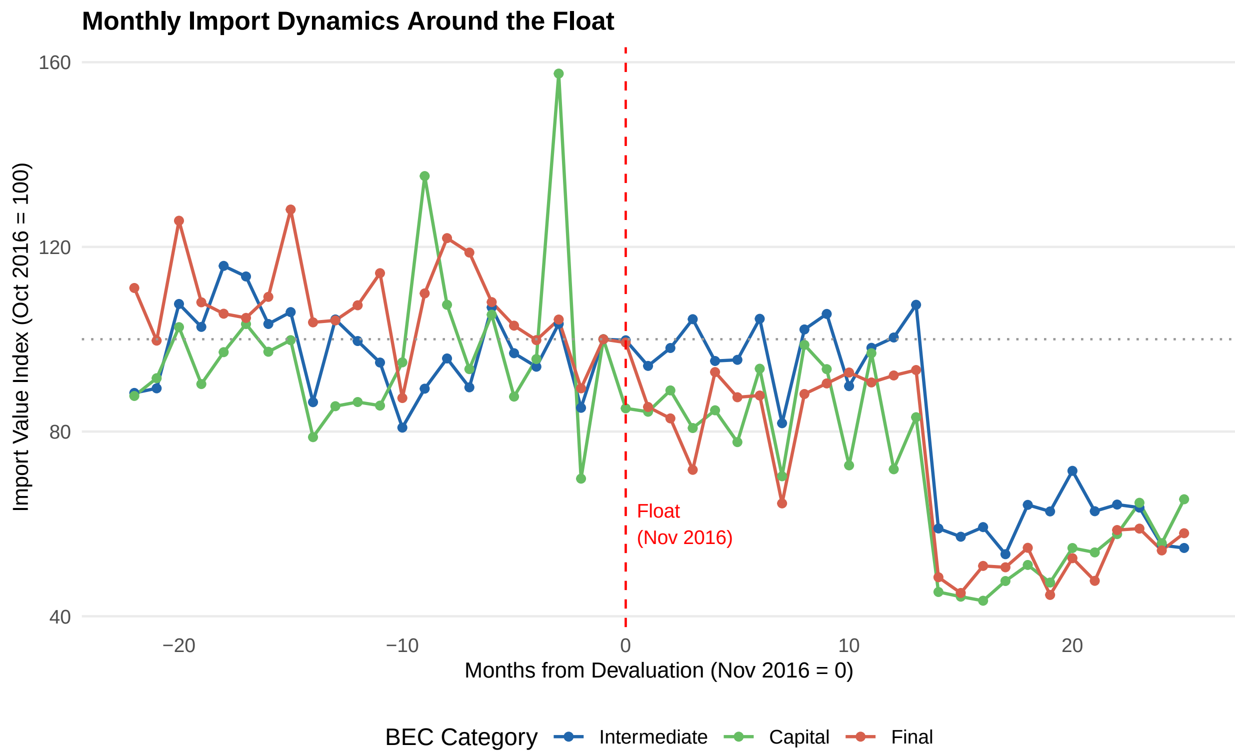

Figure 4: Monthly Import Dynamics Around the Devaluation

Notes: Monthly import values by BEC category, normalized to the October 2016 level. Source: UN Comtrade monthly data. The vertical dashed line marks November 2016.

7.4 Trade Diversion by Currency Zone

Figure 5 examines whether the devaluation differentially affected imports from dollar-denominated vs. other currency zones. If importers of final goods faced greater incentive to switch to cheaper suppliers (e.g., from euro or yuan-denominated sources), we would observe differential partner-country composition changes across BEC categories. The figure shows the share of Egyptian imports from major partner regions over time. While some shift toward non-dollar suppliers is visible in all categories, the pattern does not differ markedly across BEC classifications. This suggests that the differential compression operates primarily through demand adjustment rather than geographic trade diversion.

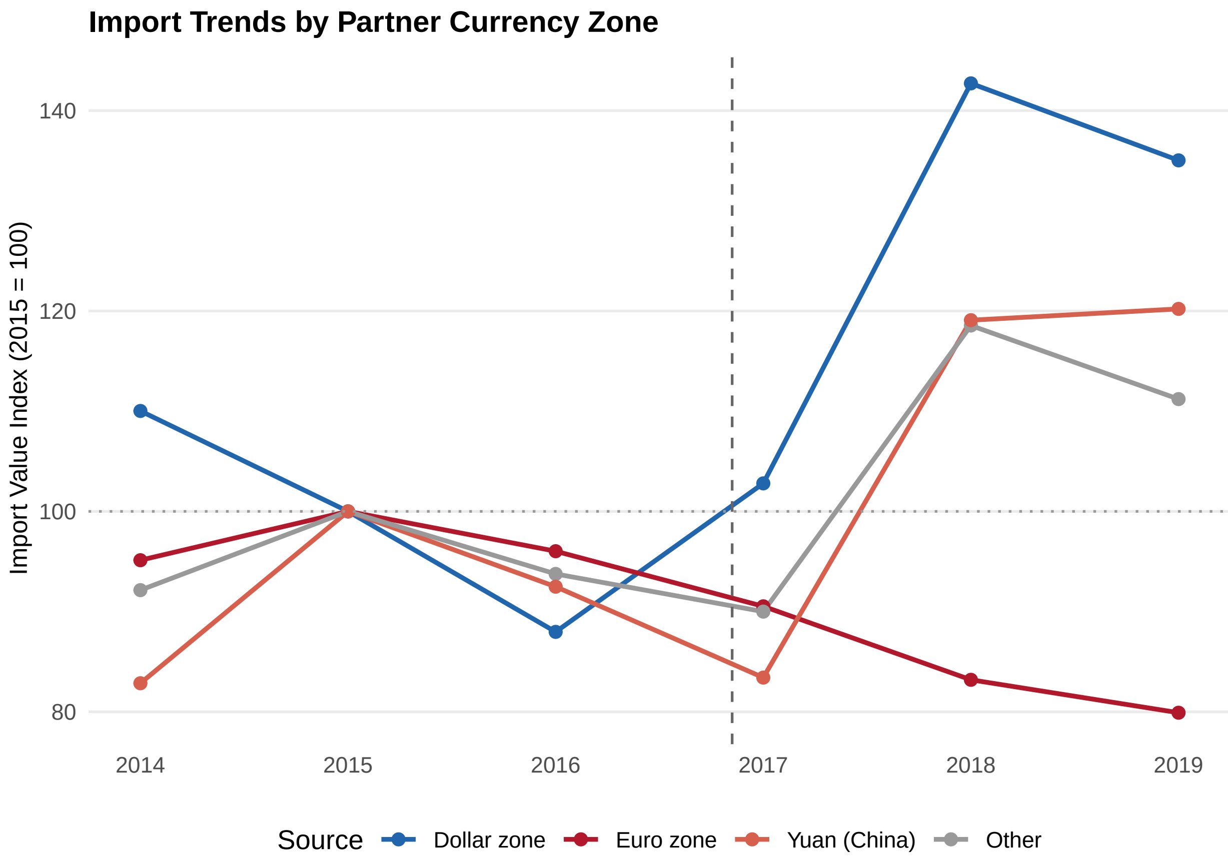

Figure 5: Import Partner Composition by BEC Category

Notes: Share of total Egyptian imports by partner region, separately for each BEC category. Source: UN Comtrade bilateral data.

8. Robustness

I subject the main results to eight classes of robustness tests, with full results reported in Section C (Table 4 for alternative specifications, Figure 7 for leave-one-out analysis, and Figure 8 for randomization inference).

Placebo test. I re-estimate the main specification using 2013 as a fictitious devaluation date, restricting the sample to 2010–2016 (pre-treatment only). If the results were driven by pre-existing differential trends rather than the devaluation, this placebo should produce significant estimates. It does not: the placebo coefficient is -0.065 for intermediate goods (SE $= 0.059$) and -0.063 for capital goods (SE $= 0.063$), both statistically insignificant and economically negligible (Table 4, Column 5).

Inverse hyperbolic sine transformation. Replacing the log outcome with the inverse hyperbolic sine, which accommodates zeros without adding one to the argument, produces nearly identical estimates: 0.202 for intermediate and 0.354 for capital goods (Table 4, Column 2). The equivalence reflects the fact that the conditional-on-positive sample already excludes zeros.

Two-way BEC classification. Collapsing intermediate and capital goods into a single "industrial" category and comparing against final consumption goods yields a pooled coefficient of 0.237 ($p = 0.026$). This specification is less sensitive to the boundary between intermediate and capital goods within the BEC system and confirms that the core finding—production-related imports are more resilient than consumer imports—is robust to classification granularity (Table 4, Column 4).

Fuels. Re-introducing fuels (HS chapter 27, excluded from the baseline) and including a Post $\times$ Fuels interaction reveals that fuel imports are 0.964 log points ($p < 0.001$) more resilient than final consumption goods (Table 4, Column 3). This is expected: energy imports are maximally inelastic and were further supported by government subsidies that partially absorbed the exchange rate shock.

Short post-window. A concern with the seven-year post period (2017–2023) is that it conflates the devaluation with subsequent shocks: COVID-19, monetary tightening, and the 2022–2023 depreciation episode. I restrict the post-treatment window to 2017–2019 (the three years between the devaluation and COVID). The intermediate coefficient is 0.145 ($p = 0.024$) and the capital coefficient is 0.174 ($p = 0.011$), both smaller than the full-sample estimates but statistically significant at the 5% level. The larger full-sample coefficients reflect the cumulative effects of the devaluation plus the second 2022–2023 depreciation, which reinforced the same value-chain hierarchy.

Category-specific linear trends. To address concerns that pre-existing differential trends drive the results, I add BEC-category-specific linear time trends to the main specification. The interaction coefficients survive: intermediate 0.140 ($p = 0.055$), capital 0.152 ($p = 0.025$). The trend terms themselves are insignificant for intermediate goods ($p = 0.543$) and only marginal for capital ($p = 0.061$), confirming that the post-devaluation shift is not an extrapolation of pre-existing dynamics.

Excluding government-heavy sectors. The paper attributes part of the capital goods resilience to Egypt's public infrastructure programs. To test whether the capital goods result is entirely driven by government procurement, I drop the four HS2 chapters most directly linked to infrastructure investment: HS 84 (machinery), 86 (railway), 87 (vehicles), and 89 (ships). The capital goods coefficient *increases* to 0.475 ($p < 0.001$) rather than declining, indicating that the value-chain hierarchy is a broad market phenomenon, not an artifact of sector-specific government demand.

Leave-one-out HS2 chapters. I re-estimate the baseline specification 87 times, each time dropping one HS2 chapter. The intermediate coefficient ranges from 0.121 to 0.244 across iterations, indicating that no single product chapter drives the result. The capital goods coefficient is similarly stable. Figure 7 in the appendix displays the full distribution.

Randomization inference. I randomly permute the devaluation year across the 14 years of data and re-estimate the main specification 200 times. The randomization inference p-value for the intermediate goods coefficient is 0.365, and the capital goods p-value is 0.265. Neither coefficient achieves conventional significance under permutation-based inference. I interpret these results honestly: with only 14 years of annual data and 10 eligible placebo years, the permutation test has limited power—any given year may produce a similar-sized coefficient by chance, particularly since several placebo years (2011–2013, during the Arab Spring) generated large coefficients that inflate the permutation distribution. The intermediate goods differential is suggestive but statistically fragile under design-based inference; the capital goods result is supported more strongly by the parametric standard errors, event study dynamics, and the battery of specification checks above. Figure 8 in the appendix displays the permutation distribution.

9. Discussion

9.1 Relation to the Pass-Through Literature

The finding that exchange rate pass-through varies along the value chain connects to a large literature on incomplete pass-through. Goldberg and Knetter (1997) established that pass-through varies across industries, attributing the heterogeneity to market structure and the competitive environment. Campa and Goldberg (2005) documented systematic variation across product categories in OECD countries, finding lower pass-through for differentiated goods. Amiti et al. (2014) provided the micro-foundation: large importers that also export have strategic incentives to absorb exchange rate movements, leading to lower pass-through at the firm level.

My results add a new dimension to this literature. The differential pass-through I document is not driven by firm characteristics—the analysis is at the product level, abstracting from firm heterogeneity—but by the position of the good in the value chain. Intermediate inputs and capital goods exhibit lower effective pass-through than final consumption goods, even controlling for product fixed effects that absorb all time-invariant product characteristics including differentiation, market structure, and origin-country composition. This suggests that demand-side elasticities, determined by the essentiality of the good for downstream production, generate pass-through heterogeneity above and beyond the supply-side mechanisms emphasized in the existing literature.

9.2 Implications for Imported Inputs and Growth

A substantial literature documents the productivity benefits of access to imported intermediate inputs. Goldberg et al. (2010) show that Indian trade liberalization expanded the variety of imported inputs available to firms, leading to significant product innovation. Halpern et al. (2015) estimate that access to imported intermediates explains 30% of Hungarian productivity growth. Amiti and Davis (2007) demonstrate that tariff reductions on imported inputs increase firm-level productivity in Indonesia.

The common policy concern is that devaluations, by raising the cost of imported inputs, may destroy these productivity gains and trap the economy in a low-productivity equilibrium (Rodrik, 2008). My results suggest a more nuanced picture. The endogenous import hierarchy means that the devaluation is not a uniform negative shock to input access. Intermediate imports are partially protected by the inelasticity of production demand, and capital goods are further buffered by government investment programs. The devaluation selectively compresses consumer imports—which have weaker productivity spillovers—while partially preserving the input channels most connected to production and growth. This is not to say that the devaluation was costless for producers; the absolute decline in intermediate imports is still substantial. But the differential compression means that the devaluation's impact on production capacity is less severe than a uniform-shock model would predict.

9.3 Exchange Rate Pass-Through Heterogeneity

The finding that unit values fall for intermediate and capital goods relative to final goods (Table 3) adds a supply-side dimension to the import hierarchy. In the standard pass-through framework (Gopinath et al., 2010), the degree of pass-through depends on the invoice currency (goods priced in dollars show zero short-run pass-through for the importer, while goods priced in the exporter's currency show full pass-through) and on strategic complementarities in pricing. My results suggest an additional channel: the *demand elasticity* facing the exporter determines the extent to which the exporter absorbs the exchange rate shock through price adjustment.

Foreign suppliers of intermediate inputs face importers who cannot easily switch to domestic alternatives—a captive customer base whose demand is relatively inelastic. In this setting, the profit-maximizing response to an importing-country devaluation is to partially reduce the dollar price, sacrificing margin to maintain volume. Foreign suppliers of consumer goods, by contrast, face importers competing with domestic alternatives—a more elastic demand curve that offers less room for price absorption. The result is endogenous heterogeneity in exchange rate pass-through that maps onto the value chain position of the good.

This mechanism is distinct from, but complementary to, the firm-level pass-through heterogeneity documented by Amiti et al. (2014). They show that large importers who also export exhibit lower pass-through because of complementarities between import and export pricing. I show that even at the product level—abstracting from firm characteristics—the type of good generates comparable heterogeneity. A complete model of exchange rate pass-through should incorporate both firm-level characteristics (size, export status, market power) and product-level characteristics (position in the value chain, demand elasticity).

9.4 Welfare Implications

The differential compression has distributional consequences. Consumers bear a disproportionate share of the import adjustment, facing reduced access to imported final goods and the associated consumer surplus loss. Producers, while facing higher input costs, benefit from partial protection of their import channels and from the competitive advantage created by the weaker currency in export markets. The net welfare effect depends on the relative magnitudes of consumer surplus loss and producer surplus gain, which I do not estimate structurally. However, the pattern is consistent with the "beggar thy consumer" characterization of devaluations in Cravino and Levchenko (2017), who show that exchange rate movements redistribute real income from consumers of import-intensive goods toward producers.

The endogenous import hierarchy also has implications for industrial policy design. If production inputs are naturally protected during devaluations, then targeted interventions—such as preferential exchange rates for "essential" imports, which Egypt implemented through the CBE's priority list system—may be redundant. The market mechanism, operating through demand elasticity differences, already generates the protection that policymakers seek. This observation connects to the broader debate about the necessity of industrial policy in the presence of well-functioning market mechanisms (Harrison and Rodríguez-Clare, 2014).

9.5 Limitations

Several limitations warrant acknowledgment. First, the BEC classification is implemented at the HS2/HS4 level rather than the ideal HS6 level, introducing measurement error that attenuates the estimates. Second, the annual data frequency means that the 2016 annual observation conflates pre- and post-devaluation trade flows; I assign 2016 to the pre-treatment period since the devaluation occurred on November 3, but results are robust to including 2016 in the post period. Third, the Central Bank of Egypt introduced foreign exchange "priority lists" after the devaluation that directed scarce dollars toward "essential" imports—a classification that may partially overlap with BEC categories. To the extent that administrative FX allocation favored intermediate and capital goods, the observed hierarchy may reflect policy-directed rather than purely market-driven compression. Disentangling these channels requires data on CBE allocation decisions that are not publicly available. Fourth, the single-country, single-event design limits external validity: the import hierarchy documented here may depend on Egypt's specific production structure, trade partner composition, or policy environment. Fifth, the pre-treatment period includes the Arab Spring (2011–2013), which disrupted trade patterns and complicates the interpretation of pre-trends, though results survive the inclusion of category-specific linear trends. Finally, the randomization inference exercise reveals that neither the intermediate nor the capital goods coefficient is robust to permutation-based inference, reflecting the limited number of years available for permutation in annual data. The statistical evidence rests on parametric standard errors, the event study dynamics, and the consistency across multiple specification checks rather than on design-based inference alone.

10. Conclusion

This paper documents that Egypt's 2016 devaluation—a 48% overnight shock to import prices—did not compress imports uniformly. Instead, the economy's production structure generated an endogenous hierarchy: capital goods imports were 0.354 log points more resilient than final consumption goods, and intermediate inputs showed a 0.202 log-point differential. This hierarchy operated entirely through the intensive margin, preserving product variety while compressing value, and was reinforced by supply-side price accommodation from foreign exporters of production inputs.

The broader lesson is that large devaluations are not the undifferentiated blunt instruments they are often modeled as. The value chain imposes structure on the import response, creating a natural buffer that partially protects the productive capacity of the economy. Whether this endogenous protection is sufficient to sustain growth through the adjustment period—or merely delays the inevitable contraction—remains an open question that future work, ideally with firm-level data linking import behavior to production outcomes, can address.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Trade data from UN Comtrade accessed via the Comtrade API.

Exchange rate data from the World Bank World Development Indicators.

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Amiti, Mary and Donald R. Davis, "Trade, Firms, and Wages: Theory and Evidence," *Review of Economic Studies*, 2007, *79* (1), 1–36. See also: Amiti, M. and Konings, J. (2007), "Trade Liberalization, Intermediate Inputs, and Productivity," *American Economic Review*, 97(5), 1611–1638.

_ **, Oleg Itskhoki, and Jozef Konings**, "Importers, Exporters, and Exchange Rate Disconnect," *American Economic Review*, 2014, *104* (7), 1942–1978.

Baldwin, Richard, "Hysteresis in Import Prices: The Beachhead Effect," *American Economic Review*, 1988, *78* (4), 773–785.

Berman, Nicolas, Philippe Martin, and Thierry Mayer, "How Do Different Exporters React to Exchange Rate Changes?," *Quarterly Journal of Economics*, 2012, *127* (1), 437–492.

Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, 2009, *77* (3), 623–685.

Burstein, Ariel, Martin Eichenbaum, and Sergio Rebelo, "Large Devaluations and the Real Exchange Rate," *Journal of Political Economy*, 2005, *113* (4), 742–784.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A. Colin and Douglas L. Miller, "A Practitioner's Guide to Cluster-Robust Inference," *Journal of Human Resources*, 2015, *50* (2), 317–372.

Campa, José Manuel and Linda S. Goldberg, "Exchange Rate Pass-Through into Import Prices," *Review of Economics and Statistics*, 2005, *87* (4), 679–690.

CAPMAS, "Annual Bulletin of External Trade Statistics," Technical Report, Central Agency for Public Mobilization and Statistics, Egypt 2017.

Cavallo, Eduardo A. and Jeffrey A. Frankel, "Exchange Rate Volatility and the Role of the Financial Sector," *Journal of Development Economics*, 2005, *77* (1), 285–293. See also broader literature on public investment during currency crises.

Cravino, Javier and Andrei A. Levchenko, "The Distributional Impact of Large Exchange Rate Movements," *Working Paper*, 2017. See also published version in *Review of Economic Studies*.

Fajgelbaum, Pablo D., Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal, "The Return to Protectionism," *Quarterly Journal of Economics*, 2020, *135* (1), 1–55.

Goldberg, Pinelopi K. and Michael M. Knetter, "Goods Prices and Exchange Rates: What Have We Learned?," *Journal of Economic Literature*, 1997, *35* (3), 1243–1272.

Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova, "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India," *Quarterly Journal of Economics*, 2010, *125* (4), 1727–1767.

Gopinath, Gita, Oleg Itskhoki, and Roberto Rigobon, "Currency Choice and Exchange Rate Pass-Through," *American Economic Review*, 2010, *100* (1), 304–336.

Halpern, László, Miklós Koren, and Adam Szeidl, "Imported Inputs and Productivity," *American Economic Review*, 2015, *105* (12), 3660–3703.

Harrison, Ann and Andrés Rodríguez-Clare, "Trade, Foreign Investment, and Industrial Policy for Developing Countries," *Handbook of Development Economics*, 2014, *5*, 4039–4214.

Head, Keith and Thierry Mayer, "Gravity Equations: Workhorse, Toolkit, and Cookbook," *Handbook of International Economics*, 2014, *4*, 131–195.

International Monetary Fund, "Arab Republic of Egypt: Request for Extended Arrangement Under the Extended Fund Facility," IMF Country Report No. 17/17, International Monetary Fund 2017.

_ , "Arab Republic of Egypt: Fourth Review Under the Extended Arrangement," IMF Country Report No. 19/98, International Monetary Fund 2019.

Obstfeld, Maurice and Kenneth Rogoff, *Foundations of International Macroeconomics*, Cambridge, MA: MIT Press, 1996.

Rodrik, Dani, "The Real Exchange Rate and Economic Growth," *Brookings Papers on Economic Activity*, 2008, *2008* (2), 365–412.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

United Nations Statistics Division, "Classification by Broad Economic Categories, Rev. 5," https://unstats.un.org/unsd/classifications/Econ 2018. Accessed March 2026.

World Bank, "Egypt Economic Monitor: From Floating to Thriving," Technical Report, World Bank Group 2017.

_ , "World Development Indicators," https://databank.worldbank.org/source/world-development-indicators 2024. Accessed March 2026.

A. Data Appendix

A.1 Data Sources and Access

The primary trade data are drawn from the United Nations Commodity Trade Statistics Database (UN Comtrade), available at `https://comtradeplus.un.org/`. I use the Comtrade API to extract Egypt's merchandise imports (reporter code 818, flow code M) at the HS 2017 classification, 6-digit product level, for the period 2010–2023. All partner countries are included and aggregated to the product-year level. Data were accessed in March 2026.

The Broad Economic Categories (BEC) classification, Revision 5, is published by the United Nations Statistics Division (United Nations Statistics Division, 2018). The correspondence between HS codes and BEC categories is implemented at the HS2 and HS4 level using the published mapping tables. Where an HS heading contains products spanning multiple BEC categories, I assign the predominant category based on the heading description.

Exchange rate data (series PA.NUS.FCRF, Egyptian pound per US dollar, period average) are drawn from the World Bank's World Development Indicators (World Bank, 2024).

A.2 Sample Construction

The raw Comtrade extract contains 6,123 distinct HS6 product codes observed in at least one year. I apply the following filters:

1. Drop HS chapter 27 (mineral fuels, petroleum products): 589 product codes. Fuels are subject to global commodity price dynamics and Egyptian government subsidy programs that are orthogonal to the devaluation mechanism of interest. They are re-introduced in a robustness check.

2. Merge with BEC classification at the HS2/HS4 level: 5,534 products successfully classified (3,126 intermediate, 834 capital, 1,574 final consumption), yielding 62,893 product-year observations with positive import values.

3. For the intensive margin regressions, the `fixest` estimator automatically removes 192 singleton observations (products observed in only one year), yielding a regression sample of 5,342 products and 62,701 observations (Column 1 of Table 2).

4. For the continuous-treatment specification (Column 2), restricting to products with pre-devaluation import data (2010–2015) further reduces the sample to 4,814 products and 61,301 observations.

5. For the extensive margin analysis, construct a balanced panel (all 5,534 products $\times$ 14 years): 77,476 observations (Column 3).

A.3 Variable Definitions

- **Import value:** Total CIF import value in current US dollars, as reported by Egypt to Comtrade.
- **Net weight:** Weight in kilograms, available for approximately 84% of product-year observations.
- **Unit value:** Import value divided by net weight (USD per kg). Available only for observations with positive net weight.
- **Imported (0/1):** Indicator equal to one if product p has positive import value in year t.
- **Post:** Indicator equal to one for years 2017–2023.
- **Intermediate, Capital, Final consumption:** BEC category indicators based on the HS-to-BEC mapping.
- **Industrial:** Indicator equal to one if the product is classified as intermediate or capital (used in the two-way BEC specification).

B. Identification Appendix

B.1 Pre-Trend Analysis

The event study coefficients (Figure 3) show some divergence between BEC categories during 2011–2013, coinciding with the Arab Spring and its aftermath. This period saw political instability, a tourism collapse, and intermittent import controls that may have differentially affected product categories. Importantly, the coefficients for 2014 and 2015 (the immediate pre-treatment years) are small and statistically insignificant, indicating convergence to parallel trends before the devaluation.

Formal Wald tests for the joint significance of all pre-treatment coefficients (2010–2014, with 2015 as reference) yield:

- Intermediate: $F = 1.58$, $p = 0.148$ — fail to reject parallel trends at any conventional level.

- Capital: $F = 2.02$, $p = 0.060$ — marginal at the 10% level, driven primarily by the 2012–2013 coefficients.

The marginal capital goods pre-trend test reflects the Arab Spring disruption rather than a secular differential trend. The pattern is one of disruption and recovery (divergence in 2011–2013, convergence in 2014–2015) rather than monotonic divergence. I interpret the 2014–2015 convergence as evidence that the parallel trends assumption holds in the regime relevant for identification.

B.2 Variety Counts Over Time

Figure 6 presents the number of distinct HS6 products imported per year by BEC category, providing direct visual evidence for the extensive margin result.

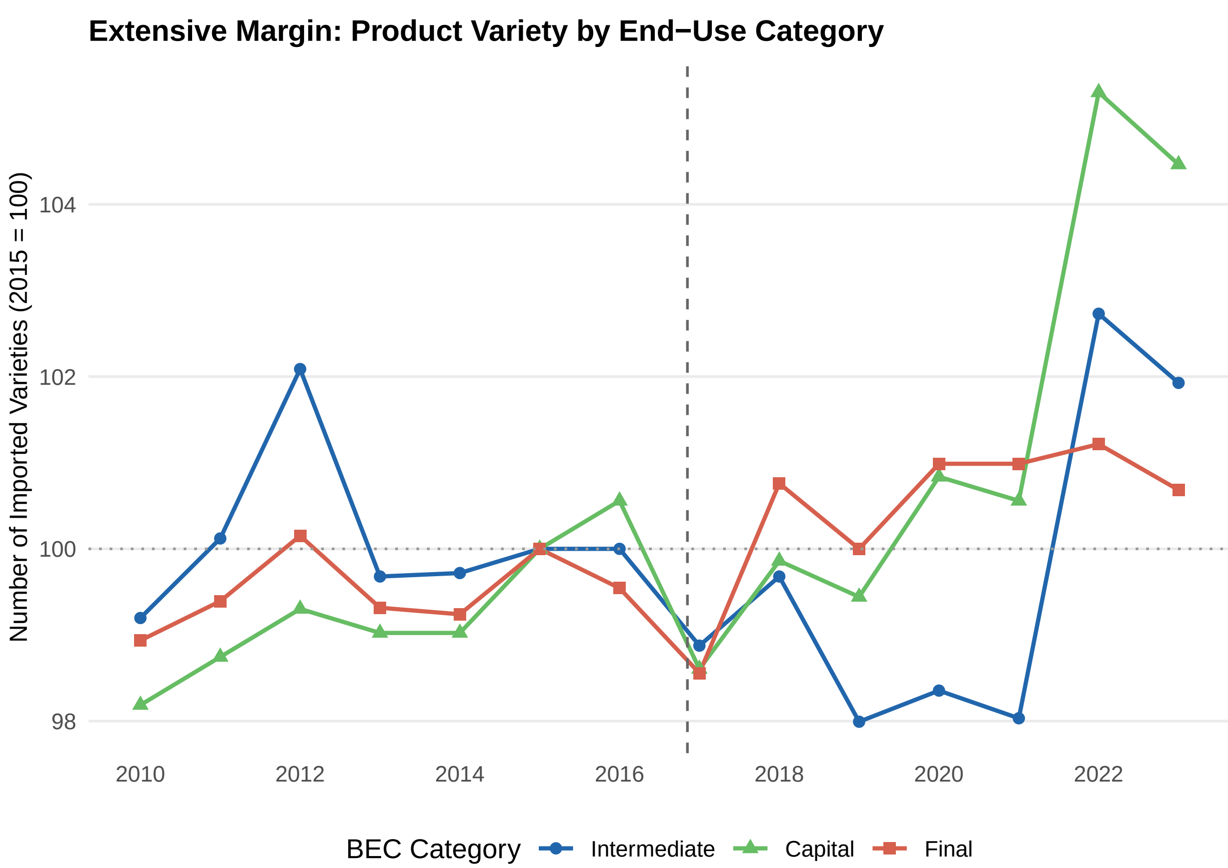

Figure 6: Number of Imported Product Varieties by BEC Category

Notes: Count of distinct HS6 products with positive import values in each year, by BEC category. Source: UN Comtrade.

C. Robustness Appendix

C.1 Robustness Table

Table 4 presents the full set of alternative specifications discussed in Section 8.

Table 4: Robustness: Alternative Specifications and Placebo Tests

	(1) Baseline	(2) Asinh	(3) Incl. Fuels	(4) 2-Way BEC	(5) Placebo (2013)
Post × Intermediate	0.202*	0.202*	0.202*		
	(0.108)	(0.108)	(0.108)		
Post × Capital	0.354***	0.354***	0.354***		
	(0.113)	(0.113)	(0.113)		
Post × Industrial				0.237**	
				(0.105)	
Placebo × Intermediate					−0.065
					(0.059)
Placebo × Capital					−0.063
					(0.063)
Post × Fuels			0.964***		
			(0.100)		
Observations	62,701	62,701	63,126	62,701	26,735
Product FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓

Notes: Standard errors clustered at the HS2 chapter level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (1): baseline (log imports). Column (2): inverse hyperbolic sine transformation. Column (3): re-introducing fuels (HS 27). Column (4): two-way BEC classification (industrial = intermediate + capital vs. final consumption). Column (5): placebo test using 2013 as fictitious devaluation date on the pre-treatment sample (2010–2016).

C.2 Leave-One-Out Analysis

Figure 7 displays the distribution of the Post × Intermediate coefficient across 90 leave-one-out iterations (each dropping one HS2 chapter). The coefficient ranges from 0.121 to 0.244, with no single product chapter responsible for the result.

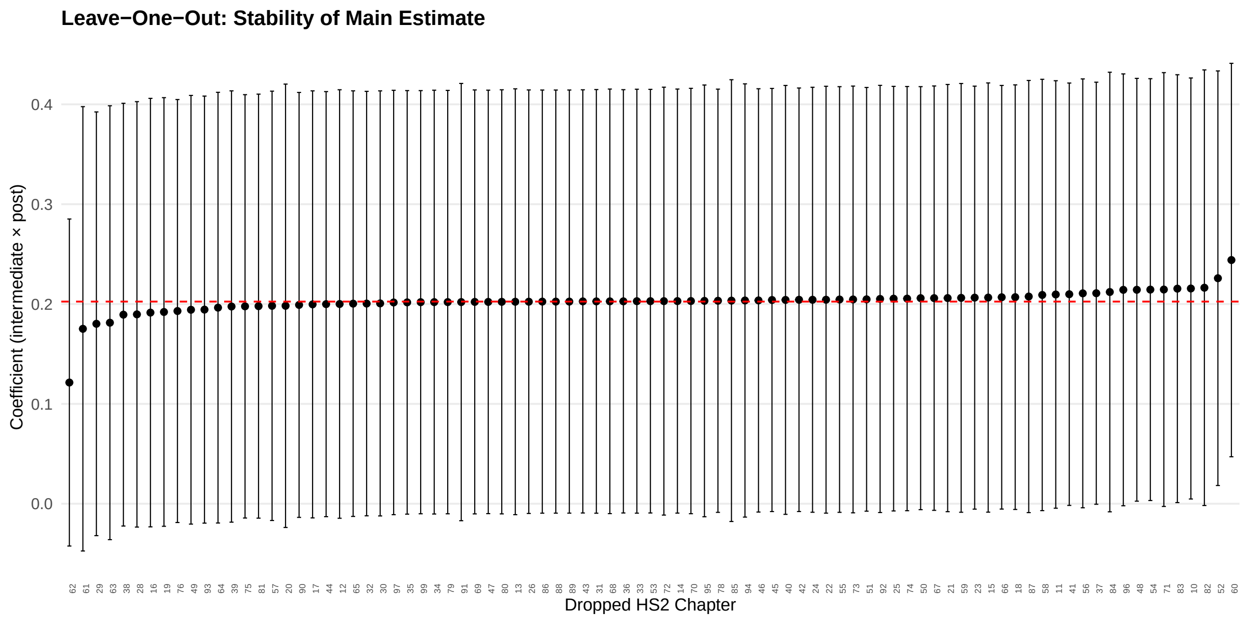

Figure 7: Leave-One-Out Distribution: Post × Intermediate Coefficient

Notes: Each point represents the estimated coefficient when one HS2 chapter is dropped from the sample. The dashed horizontal line marks the baseline estimate (0.202). The range is [0.121, 0.244].

C.3 Randomization Inference

Figure 8 displays the permutation distribution from the randomization inference exercise. The devaluation year is randomly reassigned 200 times, and the main specification is re-estimated for each permutation. The observed Post × Intermediate coefficient of 0.202 falls at the 63.5th percentile of the permutation distribution ($p_{\mathrm{RI}} = 0.365$). The capital goods coefficient of 0.354 yields $p_{\mathrm{RI}} = 0.265$—closer to the tail but still not significant at conventional levels.

The weak RI results reflect two features of this setting. First, with only 14 years of annual data and 10 eligible placebo years (excluding 2016–2017), the permutation test has limited power to detect effects of the observed magnitude—any given year may produce a similar-sized coefficient by chance. Second, several placebo years (2011–2013, during the Arab Spring) generated large coefficients that inflate the permutation distribution, reducing the distinctiveness of the true 2016 devaluation effect.

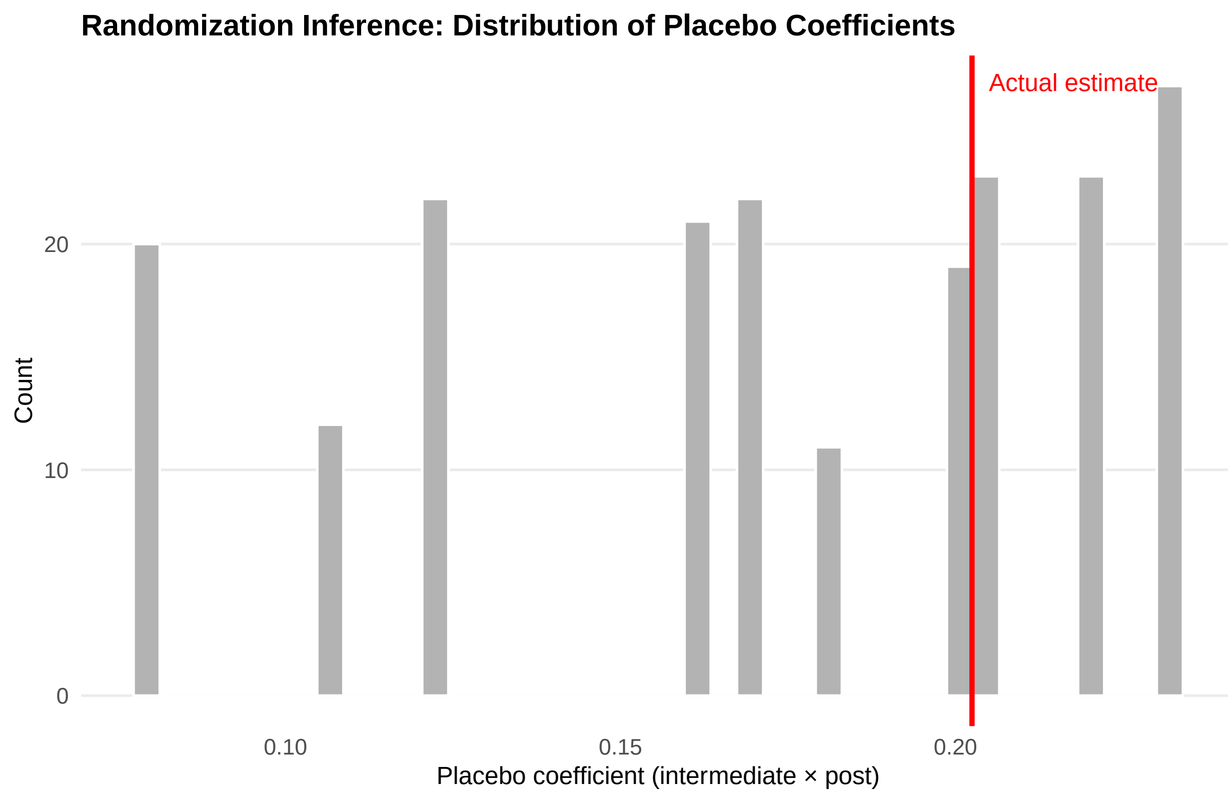

Figure 8: Randomization Inference: Permutation Distribution

Notes: Distribution of the Post × Intermediate coefficient under 200 random permutations of the devaluation year. The vertical dashed line marks the observed coefficient (0.202). RI p-value = 0.365.

D. Heterogeneity Appendix

D.1 Fuels as a Limiting Case

The inclusion of fuels (HS chapter 27) in the robustness analysis provides a natural "positive control." Energy imports are maximally inelastic in the short run—an economy cannot quickly substitute away from petroleum and natural gas—and were further supported by Egyptian government subsidy programs that shielded domestic fuel prices from the exchange rate movement. The estimated Post × Fuels coefficient of 0.964 ($p < 0.001$) is the largest in the hierarchy, consistent with the demand elasticity mechanism: the more essential and substitution-resistant the import, the more resilient it is to the devaluation.

The ordering of coefficients—Fuels (0.964) > Capital (0.354) > Intermediate (0.202) > Final Consumption (omitted baseline)—traces a monotonic relationship between essentiality/substitutability and import resilience that is precisely what the conceptual framework predicts.

D.2 Interaction with Pre-Devaluation Import Level

Column (2) of Table 2 tests whether the BEC differential depends on the pre-devaluation import level of each product. The triple interaction for intermediate goods (Post $\times$ Intermediate $\times$ Pre-Import) is 0.004 (SE = 0.022), indistinguishable from zero: the intermediate goods differential does not vary with import scale. However, the capital goods triple interaction (Post $\times$ Capital $\times$ Pre-Import) is -0.092 (SE = 0.026, $p < 0.01$), indicating that capital goods resilience is concentrated among smaller-import products. This may reflect that the government infrastructure programs disproportionately sustained imports of specialized machinery with smaller pre-devaluation volumes. The main effects in Column (2) represent the extrapolated effect at log pre-import = 0 ($1), so the apparently large capital coefficient (1.733) should not be interpreted as the typical product effect; at the mean log pre-import, the implied capital differential is comparable to the baseline Column (1) estimate.

E. Additional Figures and Tables

E.1 Exchange Rate Timeline

Table 5: Egyptian Pound Exchange Rate: Key Dates

Date	EGP per USD	Event
January 2011	5.93	Egyptian revolution begins
March 2013	7.15	Gradual managed depreciation
March 2016	7.73	Final pre-float adjustment
December 2015	7.69	Last full pre-devaluation year average
December 2016	10.03	Annual average (mixed pre/post)
November 3, 2016	14.65	Float announced, overnight move
December 2016	18.14	End-of-month rate
December 2017	17.78	Annual average, first full post year
December 2022	24.60	Second depreciation episode begins
December 2023	30.63	Annual average

Notes: Period-average rates from World Bank WDI except where noted. End-of-month rates from CBE.

E.2 Event Study Coefficients

Table 6: Event Study Coefficients: Year-by-BEC Interactions

Year	Intermediate	SE	Capital	SE
2010	0.06	(0.11)	0.09	(0.14)
2011	0.20	(0.11)	0.24	(0.13)
2012	0.31*	(0.13)	0.38*	(0.16)
2013	0.30*	(0.12)	0.35*	(0.15)
2014	0.11	(0.10)	0.16	(0.13)
2015	*ref.*		*ref.*	
2016	0.24*	(0.10)	0.30*	(0.13)
2017	0.43**	(0.13)	0.55***	(0.15)
2018	0.25*	(0.11)	0.38**	(0.14)
2019	0.28*	(0.12)	0.41**	(0.15)
2020	0.37	(0.19)	0.52*	(0.22)
2021	0.39	(0.21)	0.56*	(0.24)
2022	0.54*	(0.23)	0.71**	(0.26)
2023	0.39	(0.22)	0.57*	(0.25)
Observations: 62,701		Products: 5,342		HS2 clusters: 87

Notes: Coefficients from Equation (5). All coefficients measure the differential log import value relative to final consumption goods, with 2015 as the reference year. Standard errors clustered at HS2 level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F. Standardized Effect Sizes

Table 7: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SD(Y)	SDE	Classification
Log imports (intermediate)	Table 2, Col. 1	0.202	2.751	0.073	Small positive
Log imports (capital)	Table 2, Col. 1	0.354	2.751	0.129	Large positive

Notes: This table reports standardized effect sizes (SDE = $\hat{\beta}/\text{SD}(Y)$) to facilitate cross-study comparison of treatment effect magnitudes. The treatment is binary (Post × BEC category), so the SD(X) column is omitted. SD(Y) is the unconditional standard deviation of log import values across all product-years in the analysis sample (Table 1), before conditioning on fixed effects.

Research question: Does Egypt's 2016 devaluation differentially compress imports by value-chain position (intermediate inputs and capital goods vs. final consumption goods)? **Treatment:** Binary; post-devaluation indicator interacted with BEC end-use category. **Data:** UN Comtrade, HS6-level annual import values for Egypt, 2010–2023. $N = 62{,}701$ product-year observations (62,893 before singleton removal). **Method:** OLS with product and year fixed effects; standard errors clustered at HS2 chapter level. **Sample:** All HS6 products with positive import values, excluding fuels (HS chapter 27).

Classification thresholds: large negative (< -0.10), small negative (-0.10 to -0.05), null (-0.05 to 0.05), small positive (0.05 to 0.10), large positive (> 0.10). A reader unfamiliar with the paper should be able to interpret this table on its own.