

When Does Uncertainty Help or Hurt Trade? A Two-Channel Mechanism and Cross-Sector Elasticity Heterogeneity

Junkai Wang

May 2026

(Preliminary Draft with Potential Typos)

Abstract

This paper studies how trade-related uncertainty affects bilateral trade. Building on Baley et al. (2020), I develop a dynamic extension of the Melitz model in which firms face uncertainty over future bilateral trade costs. The model identifies two opposing channels. A pessimistic shift in expected trade costs raises the export cutoff and reduces trade, while a mean-preserving increase in dispersion can raise expected profits through convex firm payoffs and increase trade. The dispersion channel becomes relatively more important compared with the pessimistic channel in high-elasticity sectors, although its absolute magnitude is hump-shaped under the Pareto specification. Using UN Comtrade data, Economic Policy Uncertainty indices, and sectoral elasticities from Broda and Weinstein (2006), I find that the response of trade to uncertainty rises with the substitution elasticity σ —positive on both the exporter and importer sides and robust on the importer side, which serves as the primary anchor—and is attenuated by higher tariffs. A calibrated reality check against the 2025 U.S. Trump tariff cycle shows the model’s trade-flow and tariff-revenue implications to be quantitatively consistent with the realised data. These findings clarify the mechanism through which uncertainty shocks affect international trade.

1 Introduction

A long-standing view in international trade is that uncertainty impedes bilateral trade and therefore reduces the gains from trade. Recent work, however, has shown that the effect of uncertainty on trade is not necessarily monotonic. Baley et al. (2020), for example, develop a framework in which uncertainty can either increase or decrease trade depending on the preference structure. This raises an important question: through what mechanism does uncertainty affect trade flows, and what does the empirical evidence suggest?

International trade is increasingly shaped by uncertainty arising from trade policy, geopolitical tensions, and macroeconomic fluctuations. Since the 2010s, a series of major events—including Brexit, the Russia–Ukraine conflict, the U.S.–China trade war, the COVID-19 pandemic, and renewed tariff uncertainty—have highlighted how rapidly changing expectations can reshape trade patterns. A growing empirical literature has examined the effects of these specific shocks on trade outcomes and firm behaviour, but it remains organised around three largely separate strands.

A first strand studies discrete trade-policy shocks. On the Brexit side, Hassan et al. (2024) document the global financial-market and trade impact of Brexit-related uncertainty, and Graziano et al. (2018) trace how anticipation of Brexit alone reduced U.K. trade through firm-level disintegration even before tariffs were realised. Steinberg (2019) provides a complementary quantitative-trade calibration in which the option value of waiting under Brexit-induced uncertainty accounts for a substantial share of the documented trade decline. On the U.S.–China side, Fajgelbaum and Khandelwal (2022) survey the welfare consequences of the 2018–2019 tariff escalation and decompose them into terms-of-trade and adjustment channels; Benguria et al. (2022) and Itakura (2020) use Chinese firm-level data and CGE methods, respectively, to estimate the additional cost of *uncertainty* in tariff escalations beyond the realised tariff increases themselves. Handley and Limão (2017) provide the canonical structural framework for trade-policy uncertainty and welfare that this strand builds on. Novy and Taylor (2020) take a different angle, showing that uncertainty depresses trade primarily through firm-level inventory adjustments rather than entry decisions; their mechanism is complementary to the export-cutoff channel emphasised in this paper.

A second strand studies geopolitical frictions and broader political environments. Steinbach (2023) documents large reallocations of global trade following the Russia–Ukraine war, with sanctions reshaping bilateral patterns within months. Góes and Bekkers (2022) use a quantitative trade model to estimate the welfare costs of geopolitical decoupling, and Liu and Yang (2025) study the role of global hegemony in shaping the long-run pattern of trade. These contributions emphasise the role of state-level conflict and geopolitical alignment, but they remain episode-specific and do not nest naturally into firm-level un-

certainty mechanisms.

A third strand uses broad indices of policy uncertainty — most prominently the Economic Policy Uncertainty (EPU) family of indices — to study the macroeconomic effects of uncertainty. The macroeconomic literature on uncertainty shocks can be traced back to Bloom (2009), who shows that uncertainty surges sharply and contracts aggregate investment through real-options effects on lumpy investment. Altig et al. (2020) document the surge in measured uncertainty during COVID-19 across multiple indicators, and Al-Thaqeb (2022) links the pandemic-era spike in policy uncertainty to firm-level investment behaviour and macro outcomes. While these papers establish the macroeconomic relevance of policy uncertainty, they typically aggregate across sectors and do not exploit the cross-sectoral heterogeneity implied by the Melitz-style model.

Across these three strands, the literature has produced compelling event-specific evidence but no unified mechanism for how a common uncertainty shock translates into bilateral trade flows. A more general question, therefore, remains open: when a common uncertainty shock hits international trade, what is the mechanism, and is its effect heterogeneous across industries?

A parallel theoretical literature has begun to address the firm-level mechanism. Handley (2014) and Handley and Limão (2015) develop sunk-cost export-entry models in which firms condition entry decisions on the expected distribution of future trade policy, generating an option-value channel through which trade-policy uncertainty depresses trade. Baley et al. (2020) generalise this insight by showing that uncertainty can either raise or lower expected trade depending on the curvature of preferences over future trade costs. These contributions provide the theoretical scaffolding for the paper, but they typically focus on a single mechanism (option value or convexity) rather than on the interaction between competing channels and on the cross-sectoral heterogeneity that this interaction generates.

To answer the open question, I build on the canonical static trade frameworks of Armington (1969), Krugman (1980), Eaton and Kortum (2002), and Melitz (2003). Arkolakis et al. (2012) show that, despite their different microfoundations, the welfare implications of these frameworks can be expressed in terms of two sufficient statistics—the trade elasticity and the share of expenditure on domestic goods—making them quantitatively equivalent for welfare analysis. These frameworks are static, however, and are not well-suited to studying uncertainty in trade. Recent work has begun to address this limitation by developing dynamic extensions (e.g., Boehm et al., 2024; Chen et al., 2024) and embedding them in DSGE environments. These developments motivate the multiple-period Melitz model developed in this paper.

This paper develops a dynamic extension of the Melitz (2003) model and identifies two distinct channels through which uncertainty operates on bilateral

trade. A *pessimistic channel*, driven by a shift in the expected distribution of future trade costs toward worse outcomes, raises the export cutoff and contracts trade. A *dispersion channel* (or *mean-preserving spread, MPS*), in which a mean-preserving increase in the variance of future trade costs raises expected profits through the convexity of firm payoffs in $\tau^{1-\sigma}$, lowers the cutoff and expands trade. The two channels operate simultaneously in general, so the net effect of uncertainty on trade is theoretically ambiguous.

A key implication of the model is that the overall effect depends critically on the elasticity of substitution. In sectors with low substitution elasticities, the pessimistic channel dominates, leading to a reduction in trade flows. In sectors with high substitution elasticities, the convexity channel becomes relatively more important and can offset—or even overturn—the negative effect of uncertainty. Under Pareto productivity, the relative magnitudes admit closed-form expressions for both pessimistic and the dispersion channels, generating a sharp testable prediction that the response of trade flows to uncertainty varies systematically with substitution elasticities.

To test the model’s prediction, I construct a bilateral panel of trade flows using UN Comtrade data at the HS2 level over 2000–2024, combined with the Economic Policy Uncertainty index from Baker et al. (2016) and bilateral applied tariffs from MAcMap–HS6 (Fontagné et al., 2022). Substitution elasticities are taken from Broda and Weinstein (2006). The empirical strategy exploits time variation in uncertainty and cross-sector variation in substitution elasticities. I estimate a multi-interaction PPML specification that includes $\text{EPU} \times \sigma$, $\text{EPU} \times \text{Tariff}$, and the triple interaction $\text{EPU} \times \sigma \times \text{Tariff}$ as primary objects of interest, controlling for high-dimensional exporter-time and importer-time fixed effects that absorb all country-time-varying shocks.

The empirical results are consistent with the theoretical mechanism. The interaction between uncertainty and substitution elasticity is positive and statistically significant on both the exporter and the importer side, and the triple interaction with bilateral tariffs is negative and significant on the importer side. These two patterns mirror the closed-form predictions: high-elasticity sectors are systematically less adversely affected by uncertainty, and this advantage is dampened in environments with higher trade costs. The pattern survives a comprehensive group of robustness exercises, including a progressive fixed-effects structure, business-cycle controls, and lagged-uncertainty specifications addressing reverse causality.

A complementary calibration exercise on Broda–Weinstein sectoral elasticities reproduces the sign and the unit-free cross-tariff attenuation pattern of the reduced-form coefficients, providing a structural cross-check on the proposed mechanism. The welfare implications under endogenous tariff revenue are derived in Appendix A and assessed in Section 4 as an illustrative quantitative benchmark rather than as point estimates. Using Trump’s 2025 tariff policy, the

calibration shows a good fit with the U.S. macroeconomic data, revealing that the closed-form mechanism matches the realised contraction in bilateral trade flows.

This paper makes two main contributions. First, it offers a unified closed-form mechanism that reconciles the negative effects of uncertainty documented in event-specific empirical studies with a positive convexity channel arising from convex firm profits. Second, it sharpens the heterogeneous-firm trade literature by highlighting the role of substitution elasticities in shaping firm responses to uncertainty and deriving closed-form expressions for the σ -heterogeneity that map directly onto reduced-form interaction coefficients estimated via PPML with high-dimensional fixed effects, following the gravity-equation conventions of Head and Mayer (2014).

The rest of the paper is organised as follows. Section 2 presents the model. Section 3 describes the data and empirical strategy and reports the baseline regressions and robustness exercises. Section 4 calibrates the closed-form expressions and discusses the σ - and $\bar{\tau}$ -heterogeneity of the predictions. Section 5 concludes. Appendix A extends the model to characterise welfare under endogenous tariff revenue.

2 A Dynamic Melitz Model with Trade-Policy Uncertainty

Consider a finite set of countries indexed by $i, j \in \mathcal{C}$ and a finite set of sectors indexed by $m \in \{1, \dots, M\}$. Time is discrete and infinite, $t = 0, 1, 2, \dots$. Labor is the only factor of production.

2.1 Households

In country j , the representative household chooses consumption and bond holdings to maximize expected lifetime utility:

$$\max_{\{C_{jt}, B_{j,t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{C_{jt}^{1-\gamma}}{1-\gamma}, \quad (1)$$

where $\beta \in (0, 1)$ is the subjective discount factor and $\gamma > 0$ is the coefficient of relative risk aversion.

The sequential budget constraint is

$$P_{jt}C_{jt} + B_{j,t+1} = w_{jt}L_j + \Pi_{jt} + R_{jt}^g + (1 + r_t^b)B_{jt}, \quad \forall t, \quad (2)$$

where P_{jt} is the aggregate price index, w_{jt} is the wage, L_j is the exogenous labor endowment, Π_{jt} denotes aggregate firm profits rebated to households, and R_{jt}^g denotes tariff revenue rebated lump-sum.

The Euler equation is

$$\frac{C_{jt}^{-\gamma}}{P_{jt}} = \beta \mathbb{E}_t \left[(1 + r_{t+1}^b) \frac{C_{j,t+1}^{-\gamma}}{P_{j,t+1}} \right]. \quad (3)$$

2.2 Sectoral Demand

Aggregate consumption is Cobb–Douglas across sectors:

$$C_{jt} = \prod_{m=1}^M C_{jmt}^{\alpha_m}, \quad \alpha_m > 0, \quad \sum_{m=1}^M \alpha_m = 1. \quad (4)$$

Let aggregate expenditure be

$$E_{jt} = P_{jt} C_{jt}. \quad (5)$$

Sector- m expenditure is therefore

$$E_{jmt} = \alpha_m E_{jt}. \quad (6)$$

Within sector m , consumption is CES across varieties:

$$C_{jmt} = \left[\sum_{i \in \mathcal{C}} \int_{\Omega_{ijmt}} q_{ijmt}(\omega)^{\frac{\sigma_m - 1}{\sigma_m}} d\omega \right]^{\frac{\sigma_m}{\sigma_m - 1}}, \quad (7)$$

where $\sigma_m > 1$ is the sector-specific elasticity of substitution.

Demand for variety ω is

$$q_{ijmt}(\omega) = E_{jmt} P_{jmt}^{\sigma_m - 1} p_{ijmt}(\omega)^{-\sigma_m}, \quad (8)$$

where

$$P_{jmt} = \left[\sum_{i \in \mathcal{C}} \int_{\Omega_{ijmt}} p_{ijmt}(\omega)^{1 - \sigma_m} d\omega \right]^{\frac{1}{1 - \sigma_m}}. \quad (9)$$

2.3 Policy History and the Recursive State

Let $\mathcal{H}_{ij,t}$ denote the full publicly observed bilateral trade-policy history up to date t . It may contain past tariff realizations, tariff announcements, negotiations, sanctions, political events, and other public information relevant for future bilateral trade policy.

The growing history is summarized by a recursive sufficient state:

$$\boxed{s_{ij,t} = \Psi(\mathcal{H}_{ij,t})}. \quad (10)$$

New information $y_{ij,t+1}$ updates this state according to

$$s_{ij,t+1} = \Gamma(s_{ij,t}, y_{ij,t+1}). \quad (11)$$

For analytical tractability, suppose that the recursive state takes values in a finite state space

$$s_{ij,t} \in \mathcal{S}_{ij} = \{1, \dots, K\}. \quad (12)$$

Its law of motion is summarized by the transition matrix

$$Q_{ij} = [Q_{ij,kk'}]_{k,k'=1}^K, \quad (13)$$

where

$$Q_{ij,kk'} = \Pr(s_{ij,t+1} = k' \mid s_{ij,t} = k), \quad \sum_{k'=1}^K Q_{ij,kk'} = 1. \quad (14)$$

The recursive state is sufficient in the sense that

$$\Pr(s_{ij,t+1} = k' \mid \mathcal{H}_{ij,t}) = \Pr(s_{ij,t+1} = k' \mid s_{ij,t}). \quad (15)$$

2.4 Sector-Specific Tariffs

The gross tariff factor in sector m is a function of the recursive policy state:

$$\boxed{\tau_{ijmt} = \mathcal{T}_{ijm}(s_{ij,t}) \geq 1.} \quad (16)$$

Thus a common bilateral policy history can generate different tariff consequences across sectors:

$$\mathcal{T}_{ijm}(s) \neq \mathcal{T}_{ijm'}(s), \quad m \neq m'. \quad (17)$$

2.5 Firms

A firm in country i and sector m is characterized by a time-invariant productivity ϕ .

The delivered marginal cost of serving destination j is

$$MC_{ijmt}(\phi, s) = \frac{w_{it}}{\phi} \mathcal{T}_{ijm}(s). \quad (18)$$

Under CES demand and monopolistic competition, the optimal price is

$$p_{ijmt}(\phi, s) = \frac{\sigma_m}{\sigma_m - 1} \frac{w_{it}}{\phi} \mathcal{T}_{ijm}(s). \quad (19)$$

Revenue is

$$x_{ijmt}(\phi, s) = E_{jmt} P_{jmt}^{\sigma_m - 1} \left[\frac{\sigma_m}{\sigma_m - 1} \frac{w_{it}}{\phi} \mathcal{T}_{ijm}(s) \right]^{1 - \sigma_m}. \quad (20)$$

Operating profits are

$$\pi_{ijmt}(\phi, s) = \frac{1}{\sigma_m} E_{jmt} P_{jmt}^{\sigma_m - 1} \left[\frac{\sigma_m w_{it} \mathcal{T}_{ijm}(s)}{(\sigma_m - 1) \phi} \right]^{1 - \sigma_m}. \quad (21)$$

Define

$$A_{ijmt} \equiv \frac{1}{\sigma_m} E_{ijmt} P_{ijmt}^{\sigma_m - 1} \left(\frac{\sigma_m w_{it}}{\sigma_m - 1} \right)^{1 - \sigma_m}. \quad (22)$$

Then

$$\pi_{ijmt}(\phi, s) = A_{ijmt} \phi^{\sigma_m - 1} \mathcal{T}_{ijm}(s)^{1 - \sigma_m}. \quad (23)$$

The tariff component is decreasing and strictly convex:

$$\frac{\partial}{\partial \tau} \tau^{1 - \sigma_m} = (1 - \sigma_m) \tau^{-\sigma_m} < 0, \quad (24)$$

$$\frac{\partial^2}{\partial \tau^2} \tau^{1 - \sigma_m} = \sigma_m (\sigma_m - 1) \tau^{-\sigma_m - 1} > 0. \quad (25)$$

2.6 Export Status and Dynamic Entry

For each destination j , a firm can be either an incumbent exporter or a potential exporter:

$$e_{ijm,t} \in \{0, 1\}. \quad (26)$$

If $e_{ijm,t} = 1$, the firm has already paid the sunk entry cost. If $e_{ijm,t} = 0$, the firm has not yet entered and retains the option to do so.

Let

$$F_{ijm} > 0 \quad (27)$$

be the one-time sunk entry cost.

Firms survive to the next period with probability $1 - \delta$. Define

$$\varrho \equiv \beta(1 - \delta) \in (0, 1). \quad (28)$$

2.7 Incumbent Exporter Value

For the propositions below, hold aggregate objects fixed at their stationary partial-equilibrium values. Define

$$k_m \equiv \sigma_m - 1 > 0. \quad (29)$$

For state k , define the flow-profit coefficient

$$b_{ijm,k} \equiv A_{ijm} \mathcal{T}_{ijm}(k)^{1 - \sigma_m}. \quad (30)$$

Collect these coefficients in

$$b_{ijm} = \begin{pmatrix} b_{ijm,1} \\ \vdots \\ b_{ijm,K} \end{pmatrix}. \quad (31)$$

The incumbent value vector satisfies

$$V_{ijm}^1(\phi) = \phi^{k_m} b_{ijm} + \varrho Q_{ij} V_{ijm}^1(\phi). \quad (32)$$

Therefore

$$\boxed{V_{ijm}^1(\phi) = \phi^{k_m} (I - \varrho Q_{ij})^{-1} b_{ijm}.} \quad (33)$$

Define the resolvent

$$\boxed{R_{ij} \equiv (I - \varrho Q_{ij})^{-1} = \sum_{h=0}^{\infty} \varrho^h Q_{ij}^h.} \quad (34)$$

Then

$$\boxed{V_{ijm}^1(\phi) = \phi^{k_m} R_{ij} b_{ijm}.} \quad (35)$$

Equivalently, state by state,

$$V_{ijm}^1(\phi, s) = A_{ijm} \phi^{\sigma_m - 1} \mathbb{E}_s \left[\sum_{h=0}^{\infty} \varrho^h \tau_{ijm,t+h}^{1-\sigma_m} \right]. \quad (36)$$

2.8 Potential Exporters

A potential exporter chooses between immediate entry and waiting.

Immediate entry yields

$$V_{ijm}^E(\phi, s) = V_{ijm}^1(\phi, s) - F_{ijm}. \quad (37)$$

Waiting yields

$$V_{ijm}^W(\phi, s) = \varrho \mathbb{E}_s[V_{ijm}^0(\phi, s')]. \quad (38)$$

Hence

$$\boxed{V_{ijm}^0(\phi, s) = \max \{V_{ijm}^1(\phi, s) - F_{ijm}, \varrho \mathbb{E}_s[V_{ijm}^0(\phi, s')]\}.} \quad (39)$$

Equivalently,

$$V_{ijm}^0(\phi, s) = \sup_{T \geq 0} \mathbb{E}_s [\varrho^T (V_{ijm}^1(\phi, s_T) - F_{ijm})]. \quad (40)$$

Define the option value of waiting:

$$\boxed{\Omega_{ijm}(\phi, s) = V_{ijm}^0(\phi, s) - [V_{ijm}^1(\phi, s) - F_{ijm}].} \quad (41)$$

Then

$$\boxed{\Omega_{ijm}(\phi, s) = \max \{0, (1 - \varrho)F_{ijm} - \pi_{ijm}(\phi, s) + \varrho \mathbb{E}_s[\Omega_{ijm}(\phi, s')]\}.} \quad (42)$$

Immediate entry occurs iff

$$\boxed{\pi_{ijm}(\phi, s) \geq (1 - \varrho)F_{ijm} + \varrho \mathbb{E}_s[\Omega_{ijm}(\phi, s')].} \quad (43)$$

2.9 Proposition 1: Dynamic Pessimism

Proposition 1 (Dynamic pessimism lowers incumbent value). *Fix (i, j, m, ϕ) . Consider two recursive history states s^A and s^B that generate the same current tariff:*

$$\mathcal{T}_{ijm}(s^A) = \mathcal{T}_{ijm}(s^B). \quad (44)$$

Suppose that for every horizon $h \geq 1$,

$$\tau_{ijm,t+h} \mid s^A \quad (45)$$

first-order stochastically dominates

$$\tau_{ijm,t+h} \mid s^B \quad (46)$$

toward higher tariffs.

Then

$$\boxed{V_{ijm}^1(\phi, s^A) \leq V_{ijm}^1(\phi, s^B)}. \quad (47)$$

The inequality is strict if dominance is strict at at least one horizon on a set of positive probability.

Proof. Let

$$g_m(\tau) = \tau^{1-\sigma_m}. \quad (48)$$

Since $\sigma_m > 1$,

$$g'_m(\tau) = (1 - \sigma_m)\tau^{-\sigma_m} < 0. \quad (49)$$

The closed-form incumbent value is

$$V_{ijm}^1(\phi, s) = A_{ijm}\phi^{\sigma_m-1} \sum_{h=0}^{\infty} \varrho^h \mathbb{E}_s[g_m(\tau_{ijm,t+h})]. \quad (50)$$

The $h = 0$ terms are identical because current tariffs are equal.

First-order stochastic dominance toward higher tariffs implies, for every decreasing function g ,

$$\mathbb{E}[g(\tau_{t+h}) \mid s^A] \leq \mathbb{E}[g(\tau_{t+h}) \mid s^B]. \quad (51)$$

Applying this to g_m gives

$$\mathbb{E}[\tau_{t+h}^{1-\sigma_m} \mid s^A] \leq \mathbb{E}[\tau_{t+h}^{1-\sigma_m} \mid s^B], \quad h \geq 1. \quad (52)$$

Since $\varrho^h > 0$ and $A_{ijm}\phi^{\sigma_m-1} > 0$, summing across horizons yields

$$V_{ijm}^1(\phi, s^A) \leq V_{ijm}^1(\phi, s^B). \quad (53)$$

Strict dominance at any positively weighted horizon implies strict inequality. $\square$

2.10 Proposition 2: Dynamic Dispersion

Proposition 2 (Dynamic tariff dispersion raises incumbent value). *Fix (i, j, m, ϕ) . Consider two recursive history states s^A and s^B with identical current tariffs.*

Suppose that for every horizon $h \geq 1$,

$$\mathbb{E}[\tau_{ijm,t+h} \mid s^A] = \mathbb{E}[\tau_{ijm,t+h} \mid s^B], \quad (54)$$

and

$$\tau_{ijm,t+h} \mid s^A \succeq_{cx} \tau_{ijm,t+h} \mid s^B, \quad (55)$$

where $\succeq_{cx}$ denotes convex order.

Then

$$\boxed{V_{ijm}^1(\phi, s^A) \geq V_{ijm}^1(\phi, s^B)}. \quad (56)$$

The inequality is strict whenever convex-order dominance is strict at some future horizon.

Proof. Again define

$$g_m(\tau) = \tau^{1-\sigma_m}. \quad (57)$$

Its second derivative is

$$g_m''(\tau) = \sigma_m(\sigma_m - 1)\tau^{-\sigma_m-1} > 0. \quad (58)$$

Thus g_m is strictly convex.

Convex-order dominance implies

$$\mathbb{E}[g_m(\tau_{t+h}) \mid s^A] \geq \mathbb{E}[g_m(\tau_{t+h}) \mid s^B] \quad (59)$$

for every $h \geq 1$.

Using the closed-form representation

$$V_{ijm}^1(\phi, s) = A_{ijm}\phi^{\sigma_m-1} \sum_{h=0}^{\infty} \varrho^h \mathbb{E}_s[g_m(\tau_{t+h})], \quad (60)$$

and noting that current tariffs are identical, gives

$$V_{ijm}^1(\phi, s^A) \geq V_{ijm}^1(\phi, s^B). \quad (61)$$

Strict convex-order dominance at any horizon implies strict inequality because g_m is strictly convex. $\square$

The result depends only on discounted horizon-specific marginal distributions. Cross-horizon covariance does not enter incumbent value directly because flow profits are additively separable over time.

2.11 Proposition 3: Dynamic Productivity Cutoff

Define

$$G_{ijm} \equiv R_{ij} b_{ijm}. \quad (62)$$

Then

$$V_{ijm}^1(\phi) = \phi^{k_m} G_{ijm}. \quad (63)$$

Define normalized sunk cost

$$y \equiv \frac{F_{ijm}}{\phi^{k_m}} \quad (64)$$

and normalized potential-exporter value

$$v_{ijm}(s; y) \equiv \frac{V_{ijm}^0(\phi, s)}{\phi^{k_m}}. \quad (65)$$

The normalized Bellman equation becomes

$$v_{ijm}(s; y) = \max \{ G_{ijm}(s) - y, \varrho \mathbb{E}_s[v_{ijm}(s'; y)] \}. \quad (66)$$

Proposition 3 (Existence and uniqueness of a dynamic productivity cutoff). *For every recursive state s , there exists a unique threshold $y_{ijm}^*(s) > 0$, and therefore a unique productivity cutoff*

$$\phi_{ijm}^*(s) = \left[\frac{F_{ijm}}{y_{ijm}^*(s)} \right]^{\frac{1}{\sigma_m - 1}}, \quad (67)$$

such that

$$\phi \geq \phi_{ijm}^*(s) \iff \text{the firm enters destination } j. \quad (68)$$

Proof. Define the normalized entry surplus:

$$\Delta(s; y) = G(s) - y - \varrho \mathbb{E}_s[v(s'; y)]. \quad (69)$$

Immediate entry is optimal iff $\Delta(s; y) \geq 0$.

The optimal-stopping representation is

$$v(s; y) = \max \left\{ 0, \sup_{T < \infty} \mathbb{E}_s [\varrho^T (G(s_T) - y)] \right\}. \quad (70)$$

For any fixed stopping time T , the payoff is affine in y :

$$\mathbb{E}_s[\varrho^T G(s_T)] - y \mathbb{E}_s[\varrho^T], \quad (71)$$

with slope in $[-1, 0]$.

Therefore $v(s; y)$ is non-increasing, convex, and satisfies the Lipschitz bound

$$-(y_2 - y_1) \leq v(s; y_2) - v(s; y_1) \leq 0 \quad (72)$$

for $y_2 > y_1$.
Hence

$$\Delta(s; y_2) - \Delta(s; y_1) = -(y_2 - y_1) \quad (73)$$

$$- \varrho \mathbb{E}_s [v(s'; y_2) - v(s'; y_1)]. \quad (74)$$

Using the Lipschitz bound,

$$-(y_2 - y_1) \leq \Delta(s; y_2) - \Delta(s; y_1) \leq -(1 - \varrho)(y_2 - y_1) < 0. \quad (75)$$

Thus $\Delta(s; y)$ is strictly decreasing in y .

At $y = 0$, entry is costless and

$$v(s; 0) = G(s). \quad (76)$$

Since the incumbent coefficient satisfies

$$G(s) = b(s) + \varrho \mathbb{E}_s [G(s')], \quad (77)$$

we obtain

$$\Delta(s; 0) = b(s) > 0. \quad (78)$$

As $y \rightarrow \infty$, optimal entry disappears and $v(s; y) \rightarrow 0$, so

$$\Delta(s; y) \rightarrow -\infty. \quad (79)$$

Continuity and strict monotonicity therefore imply the existence of a unique $y^*(s) > 0$ satisfying

$$\Delta(s; y^*(s)) = 0. \quad (80)$$

Since

$$y = \frac{F}{\phi^{\sigma_m - 1}} \quad (81)$$

is strictly decreasing in ϕ , the unique normalized-cost threshold maps one-to-one into

$$\phi^*(s) = \left[\frac{F}{y^*(s)} \right]^{1/(\sigma_m - 1)}. \quad (82)$$

□

2.12 Piecewise Closed Form for the Waiting Option

For a given productivity ϕ , let $\mathcal{W}(\phi)$ denote the set of states in which the firm waits and $\mathcal{E}(\phi)$ the set of states in which it enters.

Partition the transition matrix accordingly:

$$Q = \begin{pmatrix} Q_{WW} & Q_{WE} \\ Q_{EW} & Q_{EE} \end{pmatrix}. \quad (83)$$

Since

$$\Omega_E(\phi) = 0, \quad (84)$$

the waiting-state Bellman system is

$$\Omega_W(\phi) = (1 - \varrho)F1_W - \phi^{k_m}b_W + \varrho Q_{WW}\Omega_W(\phi). \quad (85)$$

Therefore, conditional on the waiting region,

$$\boxed{\Omega_W(\phi) = (I - \varrho Q_{WW})^{-1} [(1 - \varrho)F1_W - \phi^{k_m}b_W]}. \quad (86)$$

Thus the optimal-stopping problem is piecewise linear in ϕ^{σ_m-1} .

For a state s , conditional on continuation waiting set $\mathcal{W}$, the indifference condition can be written

$$0 = (1 - \varrho)F - \phi^{k_m}b_s + \varrho Q_{sW}\Omega_W(\phi). \quad (87)$$

Let

$$R_W \equiv (I - \varrho Q_{WW})^{-1}. \quad (88)$$

Substituting the closed form for Ω_W gives

$$\phi^{k_m} [b_s + \varrho Q_{sW}R_W b_W] = (1 - \varrho)F [1 + \varrho Q_{sW}R_W 1_W]. \quad (89)$$

Hence, conditional on the candidate waiting region,

$$\boxed{\phi_s^* = \left\{ \frac{(1 - \varrho)F [1 + \varrho Q_{sW}R_W 1_W]}{b_s + \varrho Q_{sW}R_W b_W} \right\}^{1/k_m}}. \quad (90)$$

The candidate solution is valid whenever the implied waiting and entry inequalities are mutually consistent.

2.13 Proposition 4: Tariff Uncertainty and Dynamic Entry

For a given firm, define flow profitability

$$z(\tau) = B\tau^{1-\sigma_m}, \quad B \equiv A_{ijm}\phi^{\sigma_m-1} > 0. \quad (91)$$

Then

$$z_\tau = -(\sigma_m - 1)\frac{z}{\tau}, \quad (92)$$

$$z_{\tau\tau} = \sigma_m(\sigma_m - 1)\frac{z}{\tau^2}. \quad (93)$$

Suppose the next-period tariff is locally represented by

$$\tau' = \bar{\tau} + \varepsilon, \quad (94)$$

with

$$\mathbb{E}[\varepsilon] = 0, \quad \mathbb{E}[\varepsilon^2] = v, \quad (95)$$

where v indexes tariff uncertainty.

Let

$$H(v) = (1 - \varrho)F + \varrho \mathbb{E}[\Omega(z(\tau'))] \quad (96)$$

denote the dynamic entry hurdle.

Proposition 4 (Two curvature channels of tariff uncertainty). *Suppose the waiting-option value is locally decreasing and convex in profitability:*

$$\Omega_z < 0, \quad \Omega_{zz} \geq 0. \quad (97)$$

For a small mean-preserving increase in tariff variance v ,

$$\boxed{\frac{\partial H}{\partial v} \approx \frac{\varrho}{2} [\Omega_z z_{\tau\tau} + \Omega_{zz} z_\tau^2]}. \quad (98)$$

The first term is negative and represents the operating-profit convexity channel. The second term is positive and represents the real-option curvature channel.

Tariff uncertainty delays entry iff

$$\boxed{\Omega_{zz} z_\tau^2 > -\Omega_z z_{\tau\tau}}. \quad (99)$$

Equivalently, defining

$$\boxed{\kappa_\Omega(z) \equiv \frac{z\Omega_{zz}(z)}{-\Omega_z(z)}}, \quad (100)$$

uncertainty delays entry iff

$$\boxed{\kappa_\Omega(z) > \frac{\sigma_m}{\sigma_m - 1}}. \quad (101)$$

Proof. A second-order Taylor expansion of profitability around $\bar{\tau}$ gives

$$z(\tau') \approx z(\bar{\tau}) + z_\tau(\bar{\tau})\varepsilon + \frac{1}{2} z_{\tau\tau}(\bar{\tau})\varepsilon^2. \quad (102)$$

Hence

$$\mathbb{E}[z'] - z(\bar{\tau}) \approx \frac{1}{2} z_{\tau\tau} v \quad (103)$$

and

$$\text{Var}(z') \approx z_\tau^2 v. \quad (104)$$

A second-order expansion of the waiting-option value gives

$$\mathbb{E}[\Omega(z')] \approx \Omega(\bar{z}) + \Omega_z (\mathbb{E}[z'] - \bar{z}) + \frac{1}{2} \Omega_{zz} \text{Var}(z'). \quad (105)$$

Therefore

$$\frac{\partial}{\partial v} \mathbb{E}[\Omega(z')] \approx \frac{1}{2} [\Omega_z z_{\tau\tau} + \Omega_{zz} z_\tau^2]. \quad (106)$$

Since

$$H(v) = (1 - \varrho)F + \varrho \mathbb{E}[\Omega(z')], \quad (107)$$

it follows that

$$\frac{\partial H}{\partial v} \approx \frac{\varrho}{2} [\Omega_z z_{\tau\tau} + \Omega_{zz} z_\tau^2]. \quad (108)$$

Because

$$\Omega_z < 0, \quad z_{\tau\tau} > 0, \quad (109)$$

the first term is negative.

Because

$$\Omega_{zz} \geq 0, \quad z_\tau^2 > 0, \quad (110)$$

the second term is non-negative.

Thus uncertainty raises the entry hurdle iff

$$\Omega_{zz} z_\tau^2 > -\Omega_z z_{\tau\tau}. \quad (111)$$

Finally,

$$\frac{z_\tau^2}{z_{\tau\tau}} = \frac{\sigma_m - 1}{\sigma_m} z. \quad (112)$$

Therefore

$$\Omega_{zz} \frac{\sigma_m - 1}{\sigma_m} z > -\Omega_z \quad (113)$$

iff

$$\frac{z\Omega_{zz}}{-\Omega_z} > \frac{\sigma_m}{\sigma_m - 1}. \quad (114)$$

This yields the stated condition. $\square$

[Closed-form decomposition] The local uncertainty derivative can equivalently be written as

$$\boxed{\frac{\partial H}{\partial v} \approx \frac{\varrho}{2\tau^2} (\sigma_m - 1) z (-\Omega_z) [(\sigma_m - 1) \kappa_\Omega(z) - \sigma_m]}. \quad (115)$$

Hence all terms outside the square bracket are positive, and the sign is entirely determined by

$$(\sigma_m - 1) \kappa_\Omega(z) - \sigma_m. \quad (116)$$

2.14 Closed-Form Comparative Statics of the Incumbent Block

Suppose a scalar uncertainty parameter η affects the transition matrix:

$$Q = Q(\eta). \quad (117)$$

The incumbent value is

$$V^1 = (I - \varrho Q)^{-1} \pi. \quad (118)$$

If current flow profits do not depend directly on η , then

$$\boxed{\frac{\partial V^1}{\partial \eta} = \varrho (I - \varrho Q)^{-1} Q_\eta V^1}. \quad (119)$$

If Q is twice differentiable,

$$\boxed{\frac{\partial^2 V^1}{\partial \eta^2} = \varrho R Q_{\eta\eta} V^1 + 2\varrho^2 R Q_{\eta} R Q_{\eta} V^1,} \quad (120)$$

where

$$R = (I - \varrho Q)^{-1}. \quad (121)$$

For a linear perturbation

$$Q(\eta) = Q_0 + \eta H, \quad (122)$$

this simplifies to

$$\boxed{\frac{\partial^2 V^1}{\partial \eta^2} = 2\varrho^2 R H R H V^1.} \quad (123)$$

2.15 Firm Population Dynamics

Let

$$\mu_{ijm,t}^1(d\phi) \quad (124)$$

denote the measure of incumbent exporters at the beginning of period t , and let

$$\mu_{ijm,t}^0(d\phi) \quad (125)$$

denote the measure of potential exporters.

Given state $s_{ij,t}$, define

$$d_{ijm,t}(\phi) = \mathbf{1} \{ \phi \geq \phi_{ijm}^*(s_{ij,t}) \}. \quad (126)$$

New entrants are

$$\mu_{ijm,t}^E(d\phi) = d_{ijm,t}(\phi) \mu_{ijm,t}^0(d\phi). \quad (127)$$

The incumbent distribution evolves according to

$$\boxed{\mu_{ijm,t+1}^1(d\phi) = (1 - \delta) [\mu_{ijm,t}^1(d\phi) + d_{ijm,t}(\phi) \mu_{ijm,t}^0(d\phi)] .} \quad (128)$$

Let $B_{im,t+1}$ denote the mass of newborn potential firms, with productivity distribution G_{im} . Then

$$\boxed{\mu_{ijm,t+1}^0(d\phi) = (1 - \delta) [1 - d_{ijm,t}(\phi)] \mu_{ijm,t}^0(d\phi) + B_{im,t+1} dG_{im}(\phi).} \quad (129)$$

2.16 Aggregate Bilateral Trade

Individual revenue can be written

$$x_{ijmt}(\phi, s_t) = \mathcal{B}_{ijmt}(s_t) \phi^{\sigma_m - 1}, \quad (130)$$

where

$$\mathcal{B}_{ijmt}(s_t) = E_{jmt} P_{jmt}^{\sigma_m - 1} \left(\frac{\sigma_m}{\sigma_m - 1} w_{it} \right)^{1 - \sigma_m} \mathcal{T}_{ijm}(s_t)^{1 - \sigma_m}. \quad (131)$$

Define

$$K_{ijm,t}^1 = \int \phi^{\sigma_m - 1} \mu_{ijm,t}^1(d\phi) \quad (132)$$

and

$$K_{ijm,t}^E = \int_{\phi_{ijm}^*(s_t)}^{\infty} \phi^{\sigma_m - 1} \mu_{ijm,t}^0(d\phi). \quad (133)$$

Then aggregate bilateral trade is

$$X_{ijmt} = \mathcal{B}_{ijmt}(s_t) (K_{ijm,t}^1 + K_{ijm,t}^E). \quad (134)$$

The incumbent productivity moment evolves as

$$K_{ijm,t+1}^1 = (1 - \delta) [K_{ijm,t}^1 + K_{ijm,t}^E]. \quad (135)$$

Thus bilateral trade admits the decomposition

$$X_{ijmt} = \underbrace{\mathcal{B}_{ijmt}(s_t)}_{\text{current tariff / intensive margin}} \times \underbrace{(K_{ijm,t}^1 + K_{ijm,t}^E)}_{\text{dynamic exporter-stock / extensive margin}}. \quad (136)$$

2.17 Pareto Productivity and Cohort Hysteresis

Suppose newborn productivity follows

$$G_{im}(\phi) = 1 - \left(\frac{\phi_{\min, im}}{\phi} \right)^{\theta_m}, \quad \theta_m > \sigma_m - 1. \quad (137)$$

For a cohort born in period b , define the historical minimum entry cutoff through period $t - 1$:

$$\underline{\phi}_{ijm,b,t-1} = \min_{u=b, \dots, t-1} \phi_{ijm}^*(s_{ij,u}). \quad (138)$$

A cohort- b firm remains outside the market at date t only if

$$\phi < \underline{\phi}_{ijm,b,t-1}. \quad (139)$$

Hence old-cohort entry occurs only if

$$\phi_{ijm}^*(s_{ij,t}) < \underline{\phi}_{ijm,b,t-1}. \quad (140)$$

If the initial cohort mass is $B_{im,b}$, entrant mass is

$$N_{ijm,b,t}^E = B_{im,b}(1-\delta)^{t-b} \left[\left(\frac{\phi_{\min,im}}{\phi_{ijm}^*(s_{ij,t})} \right)^{\theta_m} - \left(\frac{\phi_{\min,im}}{\underline{\phi}_{ijm,b,t-1}} \right)^{\theta_m} \right]_+. \quad (141)$$

The corresponding entrant productivity moment is

$$K_{ijm,b,t}^E = B_{im,b}(1-\delta)^{t-b} \frac{\theta_m}{\theta_m - \sigma_m + 1} \phi_{\min,im}^{\theta_m} \left[\left(\phi_{ijm}^*(s_{ij,t}) \right)^{\sigma_m - 1 - \theta_m} - \left(\underline{\phi}_{ijm,b,t-1} \right)^{\sigma_m - 1 - \theta_m} \right]_+. \quad (142)$$

Aggregating across cohorts,

$$K_{ijm,t}^E = \sum_{b \leq t} K_{ijm,b,t}^E. \quad (143)$$

The historical minimum cutoff generates endogenous hysteresis in the extensive margin: a temporary increase in the cutoff can halt entry, while a later decline that remains above the previous historical minimum need not induce renewed entry among old cohorts because the relevant high-productivity firms have already entered.

3 Empirical Findings

The welfare and consumption predictions developed in the previous section and in Appendix A are theoretically rich but difficult to take to the data. Once tariff revenue $R_{j,t+1}^g$ is allowed to respond endogenously, the consumption response to uncertainty depends on three channels—firm profits, the import price index, and tariff revenue—whose relative strengths vary with country-specific parameters such as the tariff level relative to the Laffer peak. Empirically, consumption aggregates and tariff revenue are also unavailable at the monthly bilateral-sector frequency of the trade data.

Propositions 1 and 2, by contrast, deliver bilateral-trade predictions that are sharp in theory and implementable at high frequency. They identify two channels through which uncertainty affects $X_{ij,t}$ —a pessimistic shift, which raises the export cutoff and contracts trade, and a mean-preserving spread (MPS), which lowers the cutoff and expands trade through the convexity of firm profits in $\tau^{1-\sigma}$. Because the two typically arise together, the net effect is theoretically ambiguous, but their relative strength depends systematically on the substitution elasticity σ : the magnitude comparison in Section 2.3.1 implies that the

convexity channel becomes relatively more important as σ rises. The remainder of the paper tests this elasticity-heterogeneity prediction directly in a bilateral panel, and relies on the calibration in Section 4 to quantify the welfare implications that the regressions cannot identify.

This section provides empirical evidence to assess these theoretical predictions. In particular, I test whether the data support the coexistence of these two channels and examine whether the net effect of uncertainty on trade flows is indeed ambiguous across countries and sectors.

3.1 Data Source and Processing

The source of the bilateral trade volume comes from the UN Comtrade dataset, which contains bilateral trade from 2000 to 2024 for 97 two-digit HS sub-categories.

To empirically capture trade policy uncertainty, I use the Economic Policy Uncertainty (EPU) index as the baseline measure (Baker et al., 2016; Cerda et al., 2016; Gil and Silva, 2018; Davis, 2016; Arbatli et al., 2019; Kroese, 2015; Ghirelli et al., 2019; Armelius et al., 2017; Hardouvelis et al., 2018; Zalla, 2017). The choice of EPU is motivated by both theoretical consistency and empirical relevance.

From a theoretical perspective, the model characterizes uncertainty as changes in the distribution of future trade costs, which are driven by policy-related shocks. In practice, such uncertainty arises from fluctuations in policy expectations, including announcements, negotiations, and potential changes in tariffs. The EPU index, constructed from newspaper coverage of policy-related uncertainty, directly captures these fluctuations in expectations and is therefore well aligned with the forward-looking nature of the model.

From an empirical perspective, EPU has been widely used in the literature to measure policy uncertainty and has been shown to affect firm behaviour, investment, and international trade (Baker et al., 2016; Caldara et al., 2020; Handley and Limão, 2017). In particular, trade policy uncertainty is often reflected in broader policy uncertainty measures, as trade policies are frequently embedded in wider economic and political environments.

This interpretation is consistent with how firms form expectations in practice. Firms update their beliefs about future policy environments based on publicly available information, including news coverage, policy announcements, and political developments, and adjust their investment and export decisions accordingly. The EPU index, constructed from newspaper-based measures of policy-related uncertainty, captures these shifts in expectations and aligns closely with the forward-looking structure of the model.

Two natural alternatives to the EPU index are the Trade Policy Uncertainty (TPU) index of Caldara et al. (2020) and the World Trade Uncertainty Index (WTUI) of Ahir et al. (2022). I do not use either as the primary measure for two practical reasons. First, the TPU index is constructed only for the United States from U.S. newspaper coverage of trade-policy news, and is therefore unavailable for the other 20+ countries in the bilateral panel in the dataset. Restricting attention to TPU would force the empirical specification to use a single time series of trade-policy uncertainty that varies only over time, eliminating the country-pair-by-time variation that is essential for identifying the $\text{EPU}_{it} \times \sigma_m$ and $\text{EPU}_{jt} \times \sigma_m$ interactions in equation (144). Second, the WTUI is reported universally at quarterly frequency, which substantially reduces the number of usable observations relative to the monthly EPU panel and reduces the statistical power for identifying the σ - and tariff-modulated interaction effects that are the central objects of the empirical analysis. The country-specific monthly EPU index, by contrast, is available for a broad cross-section of countries at the same frequency as the trade and tariff data, and it captures policy uncertainty more broadly construed—which Caldara et al. (2020) and Handley and Limão (2017) argue is the appropriate object for studying how firms form expectations about future trade-policy environments. I therefore adopt EPU as the primary uncertainty measure throughout, while noting that broader EPU and narrow trade-specific TPU are positively correlated in the U.S. subsample where both are observed.

Taken together, the EPU index provides a suitable proxy for the uncertainty parameter in the model, capturing both the timing and the forward-looking dimension of policy shocks.

My model also reveals an industry heterogeneity differentiated by σ . To address this in the empirical model, I employ the dataset for trade elasticity estimation from Broda and Weinstein (2006) at the 3-digit HS code level and aggregate it into HS2 codes in my dataset.

3.2 The Baseline Model

The closed-form bilateral trade flow in equation (??) carries an explicit dependence on the gross trade cost factor $\tau_{ij,t}^{1-\sigma}$, and the comparative statics in Corollary 1 show that both the convexity and pessimistic channels of uncertainty are scaled by $\bar{\tau}$ and σ . The empirical specification therefore needs to allow the response of trade to uncertainty to depend on both sectoral substitutability and the level of bilateral trade costs from the outset, rather than first estimating an uncertainty-only specification and adding tariffs as a separate robustness exercise. I accordingly take a multi-interaction model as the baseline.

The baseline regression is

$$\begin{aligned}\mathbb{E}[X_{ijmt}] = \exp & \left(\beta_1 \text{EPU}_{it} \sigma_m + \beta_2 \text{EPU}_{jt} \sigma_m + \beta_3 \text{EPU}_{it} \text{Tariff}_{ijmt} + \beta_4 \text{EPU}_{jt} \text{Tariff}_{ijmt} \right. \\ & + \beta_5 \text{EPU}_{it} \sigma_m \text{Tariff}_{ijmt} + \beta_6 \text{EPU}_{jt} \sigma_m \text{Tariff}_{ijmt} \\ & \left. + \gamma_1 \sigma_m \text{Tariff}_{ijmt} + \gamma_2 \text{Tariff}_{ijmt} + \alpha_{it} + \delta_{jt} \right).\end{aligned}\tag{144}$$

where X_{ijmt} denotes exports from country i to country j in HS2 sector m at time t ; EPU_{it} and EPU_{jt} are the Economic Policy Uncertainty indices of the exporter and importer; σ_m is the sectoral elasticity of substitution from Broda and Weinstein (2006); and Tariff_{ijmt} is the bilateral applied tariff from MAcMap-HS6 (Fontagné et al., 2022) aggregated to HS2. The fixed effects α_{it} and δ_{jt} are exporter-time and importer-time, which absorb all country-time-varying shocks (including the EPU level itself, business-cycle conditions, exchange rates, and aggregate trade policy stances). Standard errors are clustered at the country-pair level.

The specification is estimated by Poisson Pseudo-Maximum Likelihood (PPML), which is well suited to gravity-type regressions with heteroskedasticity and zero trade flows (Santos Silva and Tenreyro, 2006, 2011; Head and Mayer, 2014). Sectoral fixed effects η_m and sector-time fixed effects ψ_{mt} are deliberately *not* included in this baseline column, since σ_m is constant within a sector and time-invariant; absorbing the sectoral level would eliminate the σ -heterogeneity that the theory predicts and the empirical specification is designed to identify. I treat η_m and ψ_{mt} as *potential* fixed effects that I add progressively in subsequent columns of Table 1 to assess robustness, recognising the trade-off between absorbing additional unobserved heterogeneity and preserving sufficient variation in σ_m to identify the coefficients of interest.

Differentiating (144) with respect to EPU_{it} and EPU_{jt} delivers the model-implied marginal effect of uncertainty on the log of expected trade:

$$\begin{aligned}\frac{\partial \ln \mathbb{E}[X_{ijmt}]}{\partial \text{EPU}_{it}} &= \beta_1 \sigma_m + \beta_3 \text{Tariff}_{ijmt} + \beta_5 \sigma_m \text{Tariff}_{ijmt}, \\ \frac{\partial \ln \mathbb{E}[X_{ijmt}]}{\partial \text{EPU}_{jt}} &= \beta_2 \sigma_m + \beta_4 \text{Tariff}_{ijmt} + \beta_6 \sigma_m \text{Tariff}_{ijmt}.\end{aligned}\tag{145}$$

Equation (145) maps directly onto Corollary 1. The coefficients β_1 and β_2 capture the elasticity of the trade response to uncertainty in the zero-tariff limit, where Corollary 1 reduces to the MPS-versus-pessimistic-shift balance scaled by $(\theta - \sigma + 1)$ and $\sigma(\theta - \sigma + 1)$; a positive estimate is consistent with the convexity channel dominating over the pessimistic channel for higher- σ sectors. The coefficients β_3 and β_4 capture how trade costs modulate the response of trade to uncertainty independently of sectoral elasticity, isolating the part of the pessimistic channel that is unrelated to the convexity-driven σ -heterogeneity; theoretically they need not have a sharp sign, since Corollary 1 predicts that

Table 1: Uncertainty, Elasticity, and Tariff Heterogeneity

	(1) Baseline	(2) + Pair FE	(3) + HS FE	(4) + HS-time FE	(5) + Pair & HS FE
<i>Panel A. Exporter side</i>					
$\text{EPU} \times \sigma$	0.0000114*** (0.000003)	0.0000096*** (0.000002)	0.0000033 (0.000004)	0.0000029 (0.000004)	-0.0000015 (0.000002)
$\text{EPU} \times \text{Tariff}$	-0.0050417 (0.009713)	-0.0006415 (0.003203)	-0.0050565 (0.009450)	-0.0058283 (0.009819)	-0.0004484 (0.001582)
$\text{EPU} \times \sigma \times \text{Tariff}$	-0.0001591 (0.000278)	-0.0001154 (0.000199)	0.0000135 (0.000194)	0.0000103 (0.000198)	0.0000403 (0.000071)
<i>Panel B. Importer side</i>					
$\text{EPU} \times \sigma$	0.0000336*** (0.000010)	0.0000313*** (0.000007)	0.0000111* (0.000007)	0.0000115* (0.000007)	0.0000110** (0.000005)
$\text{EPU} \times \text{Tariff}$	-0.0100492 (0.008840)	0.0008822 (0.004149)	-0.0124908 (0.009235)	-0.0144498 (0.009922)	-0.0008330 (0.002840)
$\text{EPU} \times \sigma \times \text{Tariff}$	-0.0011664*** (0.000413)	-0.0006165** (0.000251)	-0.0007321** (0.000320)	-0.0007676** (0.000318)	-0.0002758 (0.000173)
<i>Common controls</i>					
$\sigma \times \text{Tariff}$	0.0849035* (0.048965)	0.0517464 (0.038778)	0.0011310 (0.043393)	0.0040429 (0.043630)	-0.0091003 (0.026864)
Tariff	-7.1408261*** (1.974590)	-5.5006056*** (1.278023)	-2.6653607 (2.300364)	-2.3312997 (2.451623)	-0.1398124 (0.754330)
Exporter-time FE	✓	✓	✓	✓	✓
Importer-time FE	✓	✓	✓	✓	✓
Country-pair FE		✓			✓
HS FE			✓		✓
HS-time FE				✓	
Observations	6,166,716	6,166,716	6,166,716	6,166,716	6,166,716
Pseudo R ²	0.390	0.462	0.725	0.729	0.808

Standard errors clustered at the country-pair level.

Robust standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the dominant tariff modulation operates through the triple interaction, but they are a meaningful component of the full marginal effect in equation (145). The triple-interaction coefficients β_5 and β_6 are the central objects: a negative estimate indicates that higher tariffs dampen the convexity channel more strongly than the pessimistic channel, exactly the prediction of the closing paragraph of Section 2.3.1.

Column (1) of Table 1 reports the baseline estimates. The interaction $\text{EPU} \times \sigma$ enters positively and is statistically significant on both the exporter side ($\hat{\beta}_1 = 1.14 \times 10^{-5}$, $p < 0.01$) and, more strongly, on the importer side ($\hat{\beta}_2 = 3.36 \times 10^{-5}$, $p < 0.01$). This is consistent with the relative importance of the convexity channel rising in σ , although the regression cannot rule out alternative interpretations in which the pessimistic channel itself attenuates more sharply in σ .

The triple interaction $\text{EPU} \times \sigma \times \text{Tariff}$ enters with a negative sign on both sides, with the importer-side coefficient strongly significant ($\hat{\beta}_6 = -1.17 \times 10^{-3}$, $p < 0.01$). The exporter-side coefficient ($\hat{\beta}_5 = -1.59 \times 10^{-4}$) is correctly signed but estimated less precisely. This is the model’s sharpest cross-sectional prediction: even though high- σ sectors are on average less negatively affected by uncertainty, the gap is compressed in high-tariff environments. The convexity benefit of dispersion is weaker when the underlying trade-cost level is higher, which makes the pessimistic channel relatively more important.

The pairwise interaction $\text{EPU} \times \text{Tariff}$ is not statistically distinguishable from zero in either panel. This is consistent with the theory: this term corresponds to the $\sigma_m = 0$ counterfactual and is not informative about either the convexity or the pessimistic mechanism on its own. The $\sigma \times \text{Tariff}$ control is positive and marginally significant on the exporter side, and the direct tariff coefficient is large and negative ($\hat{\gamma}_2 \approx -7.14$, $p < 0.01$), as expected from a standard gravity setup.

Taken together, the baseline coefficients show three patterns that the theory predicts jointly. First, the elasticity of substitution σ is a key dimension of cross-sectoral heterogeneity in the response of trade to uncertainty: higher- σ sectors respond more positively, consistent with a stronger convexity channel. Second, this heterogeneity is itself heterogeneous in the level of trade costs: the positive elasticity of trade to uncertainty in high- σ sectors is dampened by higher tariffs, consistent with the closed-form prediction that tariffs attenuate the MPS channel more than the pessimistic channel. Third, the asymmetry between the exporter and importer sides—larger and more precise importer-side coefficients—is consistent with the model, in which the demand-side elasticity in country j governs how firms in country i value the convex payoff of a dispersed future trade-cost distribution.

A natural concern is that the EPU index does not separately measure the dispersion and pessimistic components of an uncertainty shock; in practice it captures their joint occurrence in the policy environment. The empirical strategy in equation (144) therefore identifies the *net* differential impact of EPU across sectors of different σ , rather than separately recovering the intensities w_p and w_d of the pessimistic and dispersion components in Corollary 1. The positive sign of $\hat{\beta}_1, \hat{\beta}_2$ and the negative sign of $\hat{\beta}_5, \hat{\beta}_6$ are consistent with the model in which the relative importance of the dispersion channel rises with σ and is dampened by higher tariffs, but they do not rule out alternative data-generating processes in which the pessimistic channel itself attenuates with σ for unrelated reasons. I therefore frame the empirical results consistent with the convexity-plus-pessimistic-channel mechanism formalised in Section 2, rather than as a sharp identification of w_d relative to w_p .

Columns (2)–(5) of Table 1 progressively add country-pair, HS sector, and HS-time fixed effects to the baseline. The qualitative pattern survives the introduction of HS-level fixed effects in columns (3) and (4): the importer-side $\text{EPU} \times \sigma$ coefficient remains positive and significant at the 10% level, and the importer-side $\text{EPU} \times \sigma \times \text{Tariff}$ coefficient remains negative and significant at the 5% level. The exporter-side coefficients lose statistical significance once HS-level fixed effects are absorbed in columns (3)–(5), although the signs of $\hat{\beta}_1$ remain positive in three of the four columns. I read this conservatively: the exporter-side mechanism is sensitive to the inclusion of HS-level fixed effects, suggesting that the apparent exporter-side σ -heterogeneity in column (1) may partly reflect sectoral-composition variation rather than a pure σ -heterogeneity effect. The importer-side mechanism is the more robust of the two and serves as the primary anchor for the rest of the empirical analysis.

3.3 Robustness Check

This section subjects the main finding—the negative interaction between EPU, sectoral substitution elasticity σ , and bilateral tariffs in the importer-side specifications to a comprehensive battery of identification checks. I organize the exercises around the three econometric concerns most likely to compromise inference in this setting: omitted sector-level trends, business-cycle confounders, and reverse causality. Each concern is addressed by a separate set of specifications. Throughout this section, standard errors are clustered at the country-pair level, and the preferred specification is the one with HS-time fixed effects (Column 3 of Table 1).

3.3.1 Progressive Fixed-Effects Structure

I first verify that the main result is not driven by the choice of fixed-effect structure. Table 1 reports five progressive specifications, from a baseline that includes only exporter-time and importer-time fixed effects (Column 1) to an increasingly demanding control structure that adds country-pair fixed effects

(Column 2), HS sector fixed effects (Column 3), and HS-time fixed effects (Column 4). Column (5) combines country-pair and HS fixed effects.

The triple-interaction coefficient on the importer side, $\text{EPU}_j \times \sigma_m \times \text{Tariff}_{ijmt}$, remains negative and statistically significant at the 5% level or better across Columns (1)–(4). The magnitude attenuates from -1.17×10^{-3} in the baseline to -7.68×10^{-4} under HS-time fixed effects, reflecting the absorption of time-invariant cross-sectoral and country-pair heterogeneity in trade patterns. The stability of the coefficient between Columns (3) and (4) indicates that sector-level time trends contribute little additional explanatory power beyond the HS sector fixed effects, consistent with the small increase in pseudo- R^2 from 0.725 to 0.729 in those columns.

Column (5), which adds country-pair fixed effects on top of HS fixed effects, yields a smaller and statistically insignificant coefficient. This attenuation reflects the limited identifying variation that remains after absorbing both sector and bilateral-pair fixed characteristics: the MAcMap-HS6 tariff data are updated only every three years, so the time-series variation in bilateral tariffs within a country-pair-sector cell is sparse. As noted by Weidner and Zylkin(2021) and Correia et al.(2020), the combination of multiple high-dimensional fixed effects can absorb the signal needed to identify cross-sectional heterogeneity in PPML gravity models. Column (5) should be accordingly treated as a boundary case rather than the main result.

3.3.2 Business-Cycle Controls and Sample Matching

A potential concern with the identification strategy is that high-elasticity sectors are also more cyclically sensitive, in which case the $\text{EPU} \times \sigma$ interaction could conflate uncertainty effects with business-cycle effects on cyclical industries. Table 2 addresses this concern by progressively adding industrial production (IP) controls, while simultaneously verifying that the result is not driven by sample composition.

Columns (1)–(2) report the baseline specification on the full sample without IP controls. Columns (3)–(4) restrict the sample to country-month observations with non-missing IP data—reducing the panel from 6.17M to 5.38M observations—but maintain the same set of regressors. Columns (5)–(6) add the interaction $\text{IP} \times \sigma$ for both exporter and importer countries. Columns (7)–(8) further add the full IP interaction hierarchy, including $\text{IP} \times \text{Tariff}$ and $\text{IP} \times \sigma \times \text{Tariff}$. Within each pair of columns, the first uses only the basic fixed-effect structure while the second additionally includes HS-time fixed effects.

The triple-interaction coefficient is remarkably stable across the key specifications. Comparing Columns (2), (4), and (6)—all with HS-time fixed effects—the coefficient takes values of -7.68×10^{-4} (no IP), -7.67×10^{-4} (sample restriction only), and -8.09×10^{-4} (sample restriction plus IP control). The

near-identical magnitudes across these three columns simultaneously rule out three distinct concerns: sample selection, business-cycle confounding, and sectoral time trends.

A secondary finding emerges from Column (8), which adds the full IP interaction hierarchy. The IP triple-interaction coefficient $IP_j \times \sigma \times \text{Tariff}$ is itself positive and statistically significant at the 10% level ($+2.59 \times 10^{-3}$, $p < 0.10$), in stark contrast to the negative EPU triple coefficient. This sign reversal reveals an underlying economic distinction: business-cycle expansions and policy uncertainty operate through opposing channels in the σ -tariff space. The EPU triple-interaction coefficient attenuates and loses significance in Column (8), consistent with collinearity between the two triple-interaction terms rather than evidence against the underlying mechanism. The preferred specification is therefore Column (6), which controls for the leading IP channel without introducing the over-control problem.

3.3.3 Lagged Uncertainty: Addressing Reverse Causality

A central identification concern is that contemporaneous EPU and trade flows may be jointly determined by unobserved shocks, generating spurious correlation. To address this, Table 3 re-estimates the main specification using EPU lagged by 12 months. Because current trade flows cannot influence uncertainty realized a year prior, any remaining association between $EPU_{j,t-12}$ and X_{ijmt} reflects the forward-looking response of trade to past uncertainty rather than simultaneity bias.

The lagged specification strengthens rather than attenuates the finding. Comparing Column (4) (lagged EPU, basic FE) with Column (1) (contemporaneous, basic FE), the triple-interaction coefficient grows in magnitude from -1.17×10^{-3} to -1.40×10^{-3} , with significance improving from $p = 0.005$ to $p = 0.001$. The same pattern holds in the HS-time fixed-effect specifications: Column (5) yields -8.42×10^{-4} ($p = 0.004$) compared to Column (2)'s -7.68×10^{-4} ($p = 0.016$). Column (6), which combines lagged EPU with both HS-time fixed effects and the $IP \times \sigma$ control, retains a coefficient of -6.04×10^{-4} significant at the 5% level.

These results are inconsistent with simple contemporaneous reverse causality, although persistent confounders correlated with both lagged uncertainty and current trade cannot be ruled out by this exercise alone. Were the contemporaneous estimate driven by simultaneity bias, lagging the right-hand side by twelve months would weaken the association. Instead, lagged uncertainty yields larger and more precise estimates, consistent with the model in which firms condition their export decisions on past realizations of policy uncertainty to form expectations about future trade-cost distributions. The persistence of the effect under the most demanding specification in Column (6)— which simultaneously addresses reverse causality, sectoral time trends, and business-cycle

confounders—provides the strongest single piece of evidence that the σ -tariff heterogeneity in trade response to uncertainty reflects a forward-looking policy mechanism.

Table 2: Robustness: Industrial Production Controls

	(1) Full sample, No IP	(2) + mt	(3) IP sample, No IP	(4) + mt	(5) + IP $\times\sigma$	(6) + mt	(7) + Full IP	(8) + mt
	basic		basic		basic		basic	
<i>Panel A. Exporter side</i>								
EPU $\times\sigma$	0.0000114*** (0.000003)	0.0000029 (0.000004)	0.0000115*** (0.000003)	0.0000030 (0.000004)	0.0000055 (0.000004)	0.0000043 (0.000004)	0.0000050 (0.000004)	0.0000039 (0.000004)
EPU $\times$ Tariff	-0.0050417 (0.009713)	-0.0058283 (0.009819)	-0.0160257 (0.013087)	-0.0115284 (0.011482)	-0.0170642 (0.013116)	-0.0115673 (0.011502)	-0.0154686 (0.012945)	-0.0104253 (0.011242)
EPU $\times\sigma\times$ Tariff	-0.0001591 (0.000278)	0.0000103 (0.000198)	-0.0001822 (0.000339)	-0.0001645 (0.000386)	-0.0001238 (0.000426)	-0.0001406 (0.000369)	-0.0000742 (0.000385)	-0.0000509 (0.000295)
<i>Panel B. Importer side</i>								
EPU $\times\sigma$	0.0000336*** (0.000010)	0.0000115* (0.000007)	0.0000382*** (0.000008)	0.0000069 (0.000007)	0.0000051 (0.000006)	0.0000121* (0.000007)	0.0000047 (0.000006)	0.0000111 (0.000007)
EPU $\times$ Tariff	-0.0100492 (0.008840)	-0.0144498 (0.009922)	-0.0199535 (0.016250)	-0.0261677* (0.014416)	-0.0259546 (0.016114)	-0.0258153* (0.014275)	-0.0255077 (0.016339)	-0.0251075* (0.015126)
EPU $\times\sigma\times$ Tariff	-0.0011664*** (0.000413)	-0.0007676** (0.000318)	-0.0015048*** (0.000440)	-0.0007671** (0.000328)	-0.0008149** (0.000331)	-0.0008088** (0.000339)	-0.0006674** (0.000314)	-0.0005677 (0.000374)
<i>IP controls</i>								
IP $_i \times \sigma$					0.0000756*** (0.000017)	0.0000421* (0.000022)	0.0000448* (0.000026)	0.0000201 (0.000027)
IP $_j \times \sigma$					-0.0000149 (0.000021)	-0.0000499** (0.000024)	0.0000171 (0.000022)	-0.0000262 (0.000026)
IP $_i \times$ Tariff							-0.0868165 (0.093156)	-0.0828195 (0.084750)
IP $_j \times$ Tariff							-0.1230262 (0.091072)	-0.0424844 (0.073784)
IP $_i \times \sigma \times$ Tariff							-0.0028394*** (0.000908)	-0.0015453 (0.001119)
IP $_j \times \sigma \times$ Tariff							0.0026723** (0.001362)	0.0025859* (0.001493)
<i>Common controls</i>								
$\sigma \times$ Tariff	0.0849035* (0.048965)	0.0040429 (0.043630)	0.1324918*** (0.036397)	0.0312424 (0.045533)	0.0306304 (0.056400)	0.0350948 (0.046188)	0.0206881 (0.215052)	-0.1117373 (0.262446)
Tariff	-7.1408261*** (1.974590)	-2.3312997 (2.451623)	-5.6693287** (2.820064)	-0.0916071 (2.488259)	-4.6653958* (2.761175)	-0.1528615 (2.493950)	16.6900132*** (3.416646)	12.3959646*** (3.529910)
Exporter-time FE	✓	✓	✓	✓	✓	✓	✓	✓
Importer-time FE	✓	✓	✓	✓	✓	✓	✓	✓
HS-time FE		✓		✓		✓		✓
IP sample			✓	✓	✓	✓	✓	✓
Observations	6,166,716	6,166,716	5,381,052	5,381,052	5,381,052	5,381,052	5,381,052	5,381,052
Pseudo R ²	0.390	0.729	0.390	0.743	0.391	0.743	0.392	0.744

Notes: Standard errors clustered at the country-pair level. Cols (1)(2) use the full sample. Cols (3)-(8) restrict to country-month observations with non-missing IP data. Cols (5)(6) add IP $\times\sigma$; Cols (7)(8) add the full IP interaction hierarchy. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Robustness: Lagged Uncertainty

	(1)	(2)	(3)	(4)	(5)	(6)
	Contemporaneous EPU			EPU lagged by 12 months		
	basic	+mt	+mt+IP	basic	+mt	+mt+IP
<i>Panel A. Exporter side</i>						
EPU $\times \sigma$	0.0000114*** (0.000003)	0.0000029 (0.000004)	0.0000043 (0.000004)	0.0000111*** (0.000003)	0.0000033 (0.000004)	0.0000044 (0.000004)
EPU $\times$ Tariff	-0.0050417 (0.009713)	-0.0058283 (0.009819)	-0.0115673 (0.011502)	-0.0081387 (0.010730)	-0.0088285 (0.011016)	-0.0091366 (0.011038)
EPU $\times \sigma \times$ Tariff	-0.0001591 (0.000278)	0.0000103 (0.000198)	-0.0001406 (0.000369)	0.0000206 (0.000248)	0.0001429 (0.000185)	-0.0001791 (0.000407)
<i>Panel B. Importer side</i>						
EPU $\times \sigma$	0.0000336*** (0.000010)	0.0000115* (0.000007)	0.0000121* (0.000007)	0.0000360*** (0.000010)	0.0000113* (0.000007)	0.0000098 (0.000007)
EPU $\times$ Tariff	-0.0100492 (0.008840)	-0.0144498 (0.009922)	-0.0258153* (0.014275)	-0.0083742 (0.010169)	-0.0117031 (0.010814)	-0.0278523** (0.013395)
EPU $\times \sigma \times$ Tariff	-0.0011664*** (0.000413)	-0.0007676** (0.000318)	-0.0008088** (0.000339)	-0.0013973*** (0.000423)	-0.0008417*** (0.000295)	-0.0006040** (0.000305)
<i>IP controls</i>						
IP _i $\times \sigma$			0.0000421* (0.000022)			0.0000431* (0.000022)
IP _j $\times \sigma$			-0.0000499** (0.000024)			-0.0000481** (0.000024)
<i>Common controls</i>						
$\sigma \times$ Tariff	0.0849035* (0.048965)	0.0040429 (0.043630)	0.0350948 (0.046188)	0.0784921* (0.047126)	-0.0083028 (0.040705)	0.0162929 (0.050051)
Tariff	-7.1408261*** (1.974590)	-2.3312997 (2.451623)	-0.1528615 (2.493950)	-7.0736905*** (2.004774)	-2.4337896 (2.489945)	-0.4167240 (2.479526)
Exporter-time FE	✓	✓	✓	✓	✓	✓
Importer-time FE	✓	✓	✓	✓	✓	✓
HS-time FE		✓	✓		✓	✓
IP $\times \sigma$ control			✓			✓
Observations	6,166,716	6,166,716	5,381,052	6,166,716	6,166,716	5,381,052
Pseudo R ²	0.390	0.729	0.743	0.390	0.729	0.743

Standard errors clustered at the country-pair level.

Cols (1)-(3) use contemporaneous EPU; Cols (4)-(6) use EPU lagged by 12 months.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4 Calibration

To capture the empirical relevance of the two channels identified in Section 2 and to provide an illustrative quantitative benchmark for the welfare effects of uncertainty, I now bring the closed-form expressions to data. This exercise is best read as a *normalized comparative-static illustration*: it asks whether the model can generate trade-flow and consumption elasticities of plausible sign, σ -heterogeneity, and tariff-modulation, given external estimates of σ and standard parameter values. It does not aim to be a structural calibration in the sense of mapping the EPU index into model probability units, and the welfare magnitudes should be read as illustrative rather than as point estimates. The exercise targets three qualitative facts: (i) pessimistic uncertainty reduces trade; (ii) the mean-preserving-spread channel produces a small positive convexity effect whose relative importance rises with σ ; and (iii) higher baseline tariffs attenuate the σ -gradient of trade responses.

4.1 Closed-Form Expressions for Calibration

The calibration evaluates the closed-form objects derived in Section 2.3.1 (trade flow) and Appendix A (tariff revenue, profits, price index, and welfare). I collect them here for reference, using the same notation throughout. All elasticities are evaluated with current-period objects held fixed and tariff revenue R^g allowed to respond endogenously to η , in the partial-equilibrium spirit of Propositions 1–4.

Probability-weighted aggregates. Two reduced-form quantities summarise the dependence of all building blocks on the probability vector $(\lambda^H, \lambda^M, \lambda^L)$:

$$\bar{\xi}_{ij,t+1}(\lambda) \equiv \mathbb{E}_t[\tau_{ij,t+1}^{1-\sigma}] = \sum_{s \in \{H,M,L\}} \lambda^s (\tau_{ij,t+1}^s)^{1-\sigma}, \quad (146)$$

$$\mathcal{T}_{ij,t+1}(\lambda) \equiv \mathbb{E}_t[(\tau_{ij,t+1} - 1) \tau_{ij,t+1}^{1-\sigma}] = \sum_{s \in \{H,M,L\}} \lambda^s (\tau_{ij,t+1}^s - 1) (\tau_{ij,t+1}^s)^{1-\sigma}. \quad (147)$$

$\bar{\xi}$ is the bilateral demand-cost aggregate that enters the firm's entry decision through Φ_{ijt} (see equation (??)); $\mathcal{T}$ is the per-unit tariff-revenue aggregate that enters $\mathbb{E}_t[R_{ij,t+1}^g]$ through equation (A.4) of Appendix A.

Pareto-implied constants. Under Pareto productivity with shape $\theta > \sigma - 1$,

$$\nu \equiv \frac{\theta - \sigma + 1}{\sigma - 1}, \quad \kappa \equiv \frac{\theta}{\theta - \sigma + 1}, \quad \omega_{ijt} \equiv \frac{\beta A_{ij,t+1} \bar{\xi}_{ij,t+1}}{\Phi_{ijt}} \in (0, 1). \quad (148)$$

ν is the trade-flow elasticity to Φ_{ijt} (equation (??)), κ is the Pareto integral constant, and ω_{ijt} is the share of expected future demand–cost in Φ_{ijt} .

Exposure parameter and income shares. Let $\chi \in (0, 1)$ denote country j ’s aggregate import exposure—imports from all foreign origins relative to GDP—under the interpretation that a common policy-uncertainty shock acts symmetrically on all bilateral pairs (i, j) with $i \neq j$. Under this interpretation the welfare elasticity formulas describe the response of country j ’s consumption to a common-uncertainty shock that propagates through j ’s overall import basket. (This differs from a strictly bilateral interpretation, where χ would be a much smaller per-partner share; I adopt the aggregate interpretation because the empirical EPU index is a country-time variable that affects all of j ’s trading partners simultaneously.) Operating profit and tariff revenue as shares of total household income are

$$s_{\Pi} = \frac{\Pi}{\Pi + wL + R^g} = \frac{1/\sigma}{(\sigma - 1)/\sigma + 1/\sigma + (\bar{\tau} - 1)\chi}, \quad s_R = \frac{(\bar{\tau} - 1)\chi}{(\sigma - 1)/\sigma + 1/\sigma + (\bar{\tau} - 1)\chi}. \quad (149)$$

The first equality in s_{Π} approximates the operating-profit share using the CES markup implication that operating profits are $1/\sigma$ of firm revenue; this should be read as an operating-surplus measure before a full national-accounting treatment of entry costs. Because the calibration is illustrative, I do not separately model the resource cost of sunk export entry f_{ij} in the welfare decomposition. I follow the common practice in calibrated trade models of treating Π_{jt} as profits accruing to a representative household under a lump-sum global (integrated-ownership) rebate; the welfare exercise should therefore be interpreted under integrated-ownership accounting rather than as a strict importer-country national-accounting measure. The alternative closed-national-accounting setup would replace the profit channel by country j ’s exporters’ profits in foreign markets, which the calibration does not analyse. With $1 + r_t = 1/\beta$ in steady state, the future-income share in N_t and the future-component share in D_t both equal

$$\rho_t = q_t = \frac{\beta}{1 + \beta}, \quad \tilde{q}_t = q_t + \frac{1 - q_t}{\gamma}. \quad (150)$$

The Euler-augmented weight $\tilde{q}_t$ governs the price-index loading in $\partial \ln C_{j,t+1}^* / \partial \eta$.

Trade flow. From Propositions 1 and 2,

$$\frac{\partial \ln X_{ijt}}{\partial \eta} = \nu \omega_{ijt} \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta}. \quad (151)$$

Profit, price index, and tariff revenue (bilateral building blocks). From Appendix A.2 (equations (A.12)–(A.14)), the three bilateral elasticities

are

$$\frac{\partial \ln \mathbb{E}_t[\Pi_{ij,t+1}]}{\partial \eta} = (\nu \omega_{ijt} + 1) \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta}, \quad (152)$$

$$\frac{\partial \ln P_{ij,t+1}}{\partial \eta} = \frac{1}{1-\sigma} (\nu \omega_{ijt} + 1) \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta}, \quad (153)$$

$$\frac{\partial \ln \mathbb{E}_t[R_{ij,t+1}^g]}{\partial \eta} = \underbrace{\nu \omega_{ijt} \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta}}_{\text{cutoff/entry}} + \underbrace{\frac{\partial \ln \tau_{ij,t+1}}{\partial \eta}}_{\text{tax-base reallocation}}. \quad (154)$$

Equation (153) is the only bilateral object whose sign reverses relative to $\bar{\xi}$ (because $1 - \sigma < 0$): a pessimistic shift, which lowers $\bar{\xi}$, raises the future price index. The tariff-revenue elasticity (154) combines an entry-channel response (sign of $\partial \ln \bar{\xi} / \partial \eta$) with a reallocation-channel response governed by the curvature of $\mathcal{R}(\tau) \equiv (\tau - 1)\tau^{1-\sigma}$ at $\bar{\tau}$.

The country- j -level elasticities follow from the calibration-based exposure mapping in Appendix A.2 (equation (A.15)):

$$\frac{\partial \ln \mathbb{E}_t[\Pi_{j,t+1}]}{\partial \eta} \approx \chi \frac{\partial \ln \mathbb{E}_t[\Pi_{ij,t+1}]}{\partial \eta}, \quad \frac{\partial \ln P_{j,t+1}}{\partial \eta} \approx \chi \frac{\partial \ln P_{ij,t+1}}{\partial \eta}, \quad \frac{\partial \ln \mathbb{E}_t[R_{j,t+1}^g]}{\partial \eta} \approx \chi \frac{\partial \ln \mathbb{E}_t[R_{ij,t+1}^g]}{\partial \eta}. \quad (155)$$

Consumption. Substituting (152)–(155) into the welfare elasticities derived in Appendix A.2 (equations (A.18)–(A.19)),

$$\frac{\partial \ln C_{jt}^*}{\partial \eta} = \rho_t \left[s_\Pi \frac{\partial \ln \mathbb{E}_t[\Pi_{j,t+1}]}{\partial \eta} + s_R \frac{\partial \ln \mathbb{E}_t[R_{j,t+1}^g]}{\partial \eta} \right] - q_t \left(1 - \frac{1}{\gamma} \right) \frac{\partial \ln P_{j,t+1}}{\partial \eta}, \quad (156)$$

and via the Euler equation,

$$\frac{\partial \ln C_{j,t+1}^*}{\partial \eta} = \rho_t \left[s_\Pi \frac{\partial \ln \mathbb{E}_t[\Pi_{j,t+1}]}{\partial \eta} + s_R \frac{\partial \ln \mathbb{E}_t[R_{j,t+1}^g]}{\partial \eta} \right] - \tilde{q}_t \frac{\partial \ln P_{j,t+1}}{\partial \eta}. \quad (157)$$

Equations (156) and (157) differ only in the price-index loading: $q_t(1 - 1/\gamma)$ vs. $\tilde{q}_t = q_t + (1 - q_t)/\gamma$. The difference, exactly $1/\gamma$, isolates the direct exposure of future consumption to $P_{j,t+1}$ through the Euler equation. These are exactly the welfare expressions of Propositions 3 and 4 in Appendix A. Since current trade costs $\tau_{ij,t}$ are normalized to the medium state and held fixed throughout, the response of C_{jt}^* arises entirely through intertemporal smoothing: the household reoptimizes today in response to revised expectations about period- $t+1$ profits, tariff revenue, and prices.

Discrete shock implementation. The calibration replaces the partial derivatives above with discrete log-changes induced by a probability shift of size $\Delta = 0.10$:

$$\left. \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta} \right|_{\Delta} = \ln \bar{\xi}_{ij,t+1}(\lambda + d\lambda) - \ln \bar{\xi}_{ij,t+1}(\lambda), \quad (158)$$

and analogously for $\mathcal{T}$. The two shocks studied are

$$\text{Pessimistic shift } (\eta_p): \quad d\lambda^H = +\Delta, \quad d\lambda^M = 0, \quad d\lambda^L = -\Delta, \quad (159)$$

$$\text{Mean-preserving spread } (\eta_d): \quad d\lambda^H = +\Delta, \quad d\lambda^M = -2\Delta, \quad d\lambda^L = +\Delta. \quad (160)$$

Because $\bar{\xi}$ and $\mathcal{T}$ are linear in λ , the level changes induced by $d\lambda$ are exactly captured by the analytical directional derivatives. The reported log changes are finite log differences; relative to derivative-times- Δ , they differ only by higher-order $O(\Delta^2)$ terms from the log transform.

4.2 Parameter Choices

The calibration combines (i) external data sources for the elasticity of substitution, (ii) standard literature values for the discount factor and risk aversion, and (iii) model-implied steady-state restrictions for the remaining objects. Table 4 summarises the baseline values.

Substitution elasticity σ . I use the country-sector estimates from Broda and Weinstein (2006), which provide σ at the HS3 level for 73 countries over 1994–2003 (11,293 observations). I trim observations with $\sigma < 1.5$ or $\sigma > 50$ to remove extreme estimation noise (a common practice given known long right-tail issues in the BW data); this retains 10,572 observations (93.6% of the raw sample) with median 3.39, mean 5.28, and an inter-quartile range of [2.47, 5.04].

Pareto shape θ . The benchmark sets $\theta = 5$, consistent with Simonovska and Waugh (2014) who estimate trade elasticities in the 4–6 range. The Pareto regularity condition $\theta > \sigma - 1$ rules out sectors with $\sigma \geq 6$ at this benchmark; the resulting sample contains 8,537 valid observations (80.8% of the trimmed BW sample). I report results for $\theta \in \{5, 10, 20\}$ to assess sensitivity, with the higher values retaining progressively more sectors at the cost of departing from the consensus estimate.

Discount factor, risk aversion, and unit interpretation. $\beta = 0.96$ corresponds to an annual real interest rate of approximately 4%, the textbook calibration. $\gamma = 2$ is the midpoint of the standard range (Mehra and Prescott,

1985). The model is a two-period setup rather than a frequency-indexed dynamic model; β is the discount factor over one period and should be interpreted as the subjective discount factor of the household over the comparative-statics horizon. *The empirical EPU index and the model’s probability shock Δ are in different units*—EPU is an index of newspaper-coverage frequencies, while Δ is a redistribution of probability mass across high/medium/low tariff states. I do not attempt to map one EPU unit into one Δ . I instead use the calibration to assess whether the model can deliver coefficients of *plausible sign and cross-sectoral relative magnitude*; quantitative comparisons with the reduced-form estimates are restricted to unit-free objects such as σ -gradients and cross- $\bar{\tau}$ attenuation ratios.

Tariff structure. The bilateral trade cost takes three values around a baseline $\bar{\tau}$: $\tau^H = \bar{\tau} + \delta$, $\tau^M = \bar{\tau}$, $\tau^L = \bar{\tau} - \delta$. The benchmark $\bar{\tau} = 1.05$ matches the global simple-average applied tariff in the MAcMap–HS6 data used in Section 3.1. Section 4.5 reports results for higher levels $\bar{\tau} \in \{1.10, 1.20, 1.30, 1.50, 1.75\}$ to map out the Laffer-curve nonlinearity in the tariff-revenue channel. The columns with $\bar{\tau} \geq 1.30$ are illustrative of the nonlinear tariff-revenue channel rather than representative of contemporary applied tariffs; for typical OECD trade pairs $\bar{\tau}$ lies in $[1.02, 1.10]$. Throughout the calibration, τ should be interpreted narrowly as the gross tariff factor rather than total trade costs; otherwise the tariff-revenue channel $\mathcal{R}(\tau) = (\tau - 1)\tau^{1-\sigma}$ would not represent revenue accruing to the government.

The half-spread δ . The half-spread δ governs the distance of the high and low future states from the medium state. Combined with the shock magnitude $\Delta = 0.10$, it pins down the units in which the comparative-static elasticities are reported: under the symmetric parameterization, the leading-order pessimistic-shift term in $\partial \ln X / \partial \eta_p$ scales linearly in δ and the mean-preserving-spread term in $\partial \ln X / \partial \eta_d$ scales in δ^2 . I set $\delta = 0.05$ so that the H, M, L states span a tariff range of 10 percentage points, broadly consistent with the magnitude of bilateral tariff revisions observed in the MAcMap–HS6 data over the three-year intervals at which the database is reported. Pinning down an exact empirical distribution of tariff revisions at the bilateral-sector level is beyond the scope of this calibration; δ should therefore be treated as a normalization parameter.

Initial probabilities. The baseline distribution is symmetric: $(\lambda^H, \lambda^M, \lambda^L) = (0.25, 0.50, 0.25)$. Comparative statics are computed for two shocks of magnitude $\Delta = 0.10$ (a 10 pp shift): a pessimistic shift $(d\lambda^H, d\lambda^M, d\lambda^L) = (+\Delta, 0, -\Delta)$, and a mean-preserving spread $(d\lambda^H, d\lambda^M, d\lambda^L) = (+\Delta, -2\Delta, +\Delta)$. The two shocks use the same high-state probability increase Δ , but the MPS shock involves a larger total probability reallocation because probability mass is shifted out of the middle state toward both tails. The two parameterisations admit a

clean decomposition of first-moment vs. second-moment effects on $\bar{\xi}$ and $\mathcal{T}$.

Income shares and exposure. The CES-monopolistic-competition and Melitz structure pins down the operating-profit share of GDP at $\Pi/\text{GDP} = 1/\sigma$, which I impose sector by sector. Wages absorb the rest: $wL/\text{GDP} = (\sigma - 1)/\sigma$. Tariff revenue satisfies the accounting identity

$$\frac{R^g}{\text{GDP}} = (\bar{\tau} - 1) \cdot \chi.$$

The benchmark sets $\chi = 0.30$, interpreted as country j 's aggregate import share in GDP and consistent with the OECD average ratio of imports to GDP for medium-to-large open economies. At $\bar{\tau} = 1.05$ this implies $R^g/\text{GDP} \approx 1.5\%$, consistent with OECD averages for applied tariffs. The welfare magnitudes scale roughly linearly in χ through equation (155); the qualitative patterns (signs, σ -gradients, and cross- $\bar{\tau}$ attenuation rates) are invariant to χ .

The demand-cost share ω_{ijt} . The weight $\omega_{ijt} \equiv \beta A_{ij,t+1} \bar{\xi}_{ij,t+1} / \Phi_{ijt} \in (0, 1)$ measures the share of expected future demand-cost in the firm's entry index Φ_{ijt} and enters every comparative static through the trade-flow elasticity $\nu \omega_{ijt}$. I compute ω_{ijt} sector by sector from its definition, evaluating $\bar{\xi}_{ij,t+1}$ at the initial probability vector $(\lambda^H, \lambda^M, \lambda^L) = (0.25, 0.50, 0.25)$ and the trade-cost states (τ^H, τ^M, τ^L) , and setting the current-period demand-cost component $A_{ijt} \tau_{ij,t}^{1-\sigma}$ to $\tau_M^{1-\sigma}$ under the partial-equilibrium assumption $A_{ij,t+1} = A_{ij,t}$ (consistent with holding wages, expenditure, and the current price index fixed). This delivers

$$\omega_{ijt} = \frac{\beta \bar{\xi}_{ij,t+1}}{\tau_M^{1-\sigma} + \beta \bar{\xi}_{ij,t+1}}. \quad (161)$$

At the benchmark $\tau_M = \bar{\tau} = 1.05$ and the symmetric initial distribution, $\bar{\xi}_{ij,t+1} \approx \tau_M^{1-\sigma} + O(\delta^2)$, so $\omega_{ijt} \approx \beta/(1 + \beta) \approx 0.490$ and is essentially flat across the empirically relevant range of σ_m (varying only within $[0.490, 0.493]$). The sensitivity of the trade-flow and welfare elasticities to σ_m therefore operates primarily through the Pareto-implied factor $\nu(\sigma_m) = (\theta - \sigma_m + 1)/(\sigma_m - 1)$ rather than through ω_{ijt} .

Steady-state cross-period weights. With $1 + r_t = 1/\beta$ in steady state, the future-income share in N_t and the future-component share in D_t both equal $\rho_t = q_t = \beta/(1 + \beta) \approx 0.49$. The Euler-augmented price weight on $C_{j,t+1}^*$ is $\tilde{q}_t = q_t + (1 - q_t)/\gamma \approx 0.74$.

Table 4: Calibrated parameters

Symbol	Description	Value	Source
σ_m	Sector elasticity of substitution	BW (2006), 1.5–50	Broda and Weinstein (2006)
θ	Pareto shape	5 benchmark; {5, 10, 20}	Simonovska and Waugh (2014)
β	Discount factor (subjective)	0.96	4% real rate, annual conv.
γ	Relative risk aversion	2	Mehra and Prescott (1985)
$\bar{\tau}$	Baseline tariff	1.05 benchmark	MAcMap–HS6 average
δ	Tariff state half-spread	0.05	normalization (see text)
λ^s	State probabilities (init.)	(0.25, 0.50, 0.25)	symmetric
Δ	Probability shock magnitude	0.10	comparative-static unit
χ	Aggregate import exposure (imports/GDP)	0.30	OECD open-economy avg.
ω_{ijt}	Demand–cost share in Φ_{ijt}	≈ 0.490	eq. (161)
s_Π	Profit share of total income	$\frac{1/\sigma}{1 + (\bar{\tau} - 1)\chi}$	CES + monopolistic comp.
s_R	Tariff-revenue share of total income	$\frac{(\bar{\tau} - 1)\chi}{1 + (\bar{\tau} - 1)\chi}$	accounting identity
ρ_t, q_t	Future weights	$\beta/(1 + \beta)$	steady state $1 + r_t = 1/\beta$
$\tilde{q}_t$	Euler-augmented weight	$q_t + (1 - q_t)/\gamma$	Euler equation

Notes: Sector-level σ_m comes from Broda and Weinstein (2006), trimmed to [1.5, 50] (retaining 93.6% of obs.). s_Π and s_R are model-implied from σ , $\bar{\tau}$, and χ rather than calibrated independently; their denominator simplifies to $1 + (\bar{\tau} - 1)\chi$ because $(\sigma - 1)/\sigma + 1/\sigma = 1$. ω_{ijt} is computed sector by sector from equation (161) but is essentially flat at ≈ 0.490 across the empirically relevant range of σ_m . χ is interpreted as country j 's aggregate import share in GDP rather than a single bilateral exposure (see Section 4.1). Comparative statics are evaluated at the symmetric initial distribution with shock magnitude $\Delta = 0.10$.

4.3 The σ -Heterogeneity of Trade and Welfare Responses

The empirical specification in Section 3.2 identifies cross-sectoral heterogeneity in the response of trade to uncertainty through interactions with the sectoral substitution elasticity σ_m . The model delivers a sharp prediction for this heterogeneity. From Propositions 1 and 2, $\partial \ln X / \partial \eta_p$ scales with $\omega_{ijt}(\theta - \sigma + 1)$ for the pessimistic shift and $\partial \ln X / \partial \eta_d$ with $\omega_{ijt}\sigma(\theta - \sigma + 1)$ for the mean-preserving spread. The first factor decreases linearly in σ , while the second is hump-shaped in σ with a peak at $\sigma = (\theta + 1)/2$. The MPS effect's absolute magnitude is therefore non-monotonic in σ ; what rises monotonically with σ is its *relative importance* compared with the pessimistic channel.

Figure 1 reports these predictions across the empirically relevant range of σ , holding all other parameters at their benchmarks. Panel (a) plots the bilateral trade-flow elasticity; Panel (b) plots the country- j consumption elasticity, where the profit, price-index, and tariff-revenue building blocks are aggregated through the exposure mapping χ of equation (155). Table 5 reports point values at eight representative levels of σ .



Figure 1: The σ -heterogeneity of trade-flow and consumption elasticities under uncertainty. Panel (a) plots $100 \times \partial \ln X_{ij} / \partial \eta$; Panel (b) plots $100 \times \partial \ln C_{jt}^* / \partial \eta$. Red lines are the pessimistic shift η_p ; blue lines are the mean-preserving spread η_d . Benchmark parameters: $\theta = 5$, $\bar{\tau} = 1.05$, $\delta = 0.05$, $\chi = 0.30$, $\gamma = 2$, $\beta = 0.96$, $\Delta = 0.10$. The grey shaded region at $\sigma = 6$ marks the Pareto regularity boundary $\theta > \sigma - 1$. Profit, price-index, and tariff-revenue building blocks entering C_{jt}^* are evaluated at the country- j level by multiplying the corresponding bilateral elasticities by χ .

The Pareto-implied trade elasticity $\nu(\sigma, \theta)$. Before turning to the figure, it is useful to register the values of $\nu = (\theta - \sigma + 1) / (\sigma - 1)$ that the calibration encounters. At the benchmark $\theta = 5$, ν takes the values 4.000, 2.333, 1.500, 1.000, 0.667, 0.429, 0.250, 0.111 at $\sigma \in \{2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, 5.5\}$, respectively. The collapse of ν from 4.0 to 0.1 over this range is the dominant mechanical force behind the σ -heterogeneity of trade flow. At the BW-median

Table 5: Model-implied elasticities across σ (benchmark $\theta = 5$, $\bar{\tau} = 1.05$, $\delta = 0.05$, $\chi = 0.30$)

σ	ν	ω	s_{Π}	$\Delta \ln X_{ij}$ (%)		$\Delta \ln \mathbb{E}_t[R_{j,t+1}^g]$ (%)		$\Delta \ln C_{jt}^*$ (%)	
				η_p	η_d	η_p	η_d	η_p	η_d
2.0	4.000	0.490	0.493	-1.878	+0.089	+4.790	-0.254	-0.379	+0.018
2.5	2.333	0.490	0.394	-1.648	+0.097	+4.800	-0.392	-0.295	+0.017
3.0	1.500	0.491	0.328	-1.417	+0.100	+4.811	-0.533	-0.249	+0.016
3.5	1.000	0.491	0.281	-1.184	+0.097	+4.823	-0.676	-0.220	+0.016
4.0	0.667	0.491	0.246	-0.951	+0.089	+4.834	-0.821	-0.199	+0.016
4.5	0.429	0.492	0.219	-0.715	+0.075	+4.846	-0.968	-0.183	+0.016
5.0	0.250	0.493	0.197	-0.478	+0.056	+4.859	-1.118	-0.171	+0.016
5.5	0.111	0.493	0.179	-0.240	+0.030	+4.872	-1.269	-0.162	+0.016

Notes: Each row evaluates the closed-form comparative statics from Propositions 1, 2, 3, and 4 at the given σ , holding $\theta = 5$, $\gamma = 2$, $\beta = 0.96$, $\bar{\tau} = 1.05$, $\delta = 0.05$, $\chi = 0.30$ fixed. Pareto regularity $\theta > \sigma - 1$ restricts $\sigma < 6$. Entries are $100 \times \partial \ln(\cdot) / \partial \eta$ for $\Delta = 0.10$. $\Delta \ln \mathbb{E}_t[R_{j,t+1}^g]$ and $\Delta \ln C_{jt}^*$ are country- j -level elasticities, obtained by multiplying the corresponding bilateral elasticities by $\chi = 0.30$. ω is computed from equation (161) and is essentially flat across σ at the benchmark $\bar{\tau} = 1.05$. $s_{\Pi} = (1/\sigma)/[1 + (\bar{\tau} - 1)\chi] = (1/\sigma)/1.015$.

$\sigma = 3.39$, $\nu(\sigma, 5) \approx 1.09$, $\nu(\sigma, 10) \approx 3.18$, and $\nu(\sigma, 20) \approx 7.32$; doubling θ roughly triples ν at this σ . The same factor multiplies the trade-flow comparative static through $\partial \ln \Phi_{ijt} / \partial \eta = \omega_{ijt} \partial \ln \xi_{ij,t+1} / \partial \eta$.

Three features of Figure 1 are worth emphasising. First, the trade-flow response to a pessimistic shift in Panel (a) contracts monotonically in σ , from -1.88% at $\sigma = 2.0$ to -0.24% at $\sigma = 5.5$. This is the direct counterpart of the $(\theta - \sigma + 1)$ factor in Proposition 1: high- σ sectors are mechanically less exposed to the pessimistic channel because the Pareto-implied trade elasticity $\nu(\sigma)$ falls sharply as σ approaches $\theta + 1$.

Second, the trade-flow response to a mean-preserving spread is positive throughout the empirically relevant range and is hump-shaped with a peak near $\sigma \approx 3.0$, where $\sigma(\theta - \sigma + 1)$ is maximised at $\sigma = (\theta + 1)/2 = 3.0$. Although the MPS response is order $O(\delta^2)$ and therefore smaller in absolute magnitude than the pessimistic response, the *ratio* of MPS to pessimistic responses rises with σ : at $\sigma = 2.0$ the convexity gain (+0.09%) is dwarfed by the pessimistic loss (-1.88% , ratio $\approx 5\%$), but at $\sigma = 5.5$ the convexity gain (+0.030%) is over 12% of the pessimistic loss (-0.24%). This rising MPS-to-pessimistic ratio, together with the attenuation of the pessimistic channel itself in high- σ sectors, is consistent with the positive sign of $\hat{\beta}_1$ and $\hat{\beta}_2$ in the baseline regression of Table 1: in the model, high- σ sectors are both less negatively affected by the pessimistic channel and relatively more buffered by the MPS channel, and the data cannot separately attribute the positive cross-sectional slope to one of these forces alone.

Third, the consumption response in Panel (b) inherits the σ -heterogeneity of trade flow but with a much milder slope. Under a pessimistic shift, $\Delta \ln C_{jt}^*$ converges from -0.38% at $\sigma = 2.0$ to -0.16% at $\sigma = 5.5$, roughly a 2.3-fold attenuation as compared with the 7.8-fold attenuation of $\Delta \ln X_{ij}$. The milder slope is the joint outcome of three sub-channels. The profit contribution $\rho s_{\Pi} \partial \ln \mathbb{E}_t[\Pi_{j,t+1}]/\partial \eta_p$ attenuates from -0.21% at $\sigma = 2.0$ to -0.12% at $\sigma = 5.5$, driven by the fall in the profit income share s_{Π} from 0.49 to 0.18 (since profit $\Pi/\text{GDP} = 1/\sigma$ shrinks with σ). The price-index contribution $-q_t(1 - 1/\gamma) \partial \ln P_{j,t+1}/\partial \eta_p$ attenuates from -0.21% to -0.08% , driven by the $1/(1 - \sigma)$ loading on $\partial \ln P_{j,t+1}/\partial \eta_p$ which shrinks in absolute value as σ rises. The tariff-revenue contribution $\rho s_R \partial \ln \mathbb{E}_t[R_{j,t+1}^g]/\partial \eta_p$ is essentially σ -invariant at $+0.035\%$ (since $\Delta \ln \mathbb{E}_t[R_{j,t+1}^g] \approx +4.83\%$ uniformly across σ and s_R is independent of σ). As σ rises, the fixed positive tariff-revenue contribution accounts for an increasing share of the total, partially buffering the attenuation of the other two channels and producing the milder slope on consumption than on trade flow. Under a mean-preserving spread, $\Delta \ln C_{jt}^*$ is small and essentially flat at $+0.016\%$ across the empirically relevant range of σ . As σ rises, the profit contribution to consumption rises (from $+0.010$ pp at $\sigma = 2.0$ to $+0.016$ pp at $\sigma = 5.5$), the tariff-revenue contribution falls (from -0.002 pp to -0.009 pp as the negative reallocation term grows in magnitude), and the price-index contribution is nearly constant. The two opposing trends in profits and tariff revenue approximately cancel, leaving the total near $+0.016\%$. I do not interpret this near-flatness as a sharp theoretical prediction but as a numerical feature of the benchmark calibration.

On the magnitude of the MPS channel. The mean-preserving-spread response of trade is small in absolute terms relative to the pessimistic-shift response: at the benchmark $\delta = 0.05$ and $\bar{\tau} = 1.05$, the peak MPS effect on trade flow is $+0.100\%$ at $\sigma = 3.0$, while the largest pessimistic effect is -1.88% at $\sigma = 2.0$. This is expected: the leading-order term in $\partial \ln X/\partial \eta_d$ is $\omega_{ijt}\sigma(\theta - \sigma + 1)a\delta^2/\bar{\tau}^2$, which is order $O(\delta^2)$ relative to the order- $O(\delta)$ pessimistic term. Two implications follow. First, the MPS channel is empirically detectable not in the *level* of trade-flow responses but in their cross-sectional *ratio* to the pessimistic responses; this ratio rises monotonically with σ as documented above, and the resulting σ -gradient is precisely what the $\text{EPU} \times \sigma$ interaction in the baseline regression of Table 1 is designed to identify. Second, the MPS response scales quadratically in δ , so its quantitative importance grows rapidly with the underlying volatility of trade-cost states; for $\delta = 0.10$ (a 20-percentage-point tariff range), the MPS trade-flow response roughly quadruples. The benchmark $\delta = 0.05$ is therefore a conservative choice; larger δ would only strengthen the role of the convexity channel.

The implied $\text{EPU} \times \sigma$ regression slope—the linear-fit slope of $\partial \ln X/\partial \eta_p$ across $\sigma \in [2.5, 5.0]$ at $\bar{\tau} = 1.05$ —is $+0.00468$. This is in the same direction as

the positive $\hat{\beta}_1$ and $\hat{\beta}_2$ in Table 1, column (1). I do not interpret the level of the model-implied slope as a direct estimate of the reduced-form coefficient, because the EPU index and the model probability shock Δ are in different units; only the sign and the cross- $\bar{\tau}$ attenuation pattern (developed in Section 4.5) admit unit-free comparison with the data.

4.4 Sensitivity to the Pareto Shape Parameter

The benchmark sets $\theta = 5$, the most empirically defensible value but one that constrains $\sigma < 6$. To assess how the calibration depends on θ , Table 6 reports cross-sector means of the model-implied elasticities under three values, $\theta \in \{5, 10, 20\}$. The higher values relax the Pareto regularity constraint at the cost of departing from the consensus 4–6 range estimated by Simonovska and Waugh (2014), but they allow for including progressively more sectors (the coverage rises from 80.8% to 95.4% of the trimmed BW sample) and to identify how the trade-flow elasticity $\nu = (\theta - \sigma + 1)/(\sigma - 1)$ amplifies all responses through Φ_{ijt} .

The qualitative patterns are robust across θ . Pessimistic shifts contract trade flow and profits, raise the future price index $P_{j,t+1}$ (since $\partial \ln \bar{\xi}/\partial \eta_p < 0$ and the $1/(1 - \sigma)$ loading is negative, their product is positive), and depress consumption. Mean-preserving spreads have small positive effects on trade flow that propagate to weakly positive consumption effects through the profit and price-index channels. The trade-flow magnitudes scale roughly proportionally to $\nu(\sigma, \theta)$: a doubling from $\theta = 5$ to $\theta = 10$ multiplies the cross-sector mean of $\Delta \ln X_{ij}$ under a pessimistic shift by 2.6 (from -1.32% to -3.44%), and a further doubling to $\theta = 20$ multiplies it by 2.3. The profit, price-index, and consumption responses scale sub-proportionally to trade flow: the doubling $\theta = 5 \rightarrow 10$ multiplies them by 1.8, 1.6, and 1.7, respectively, and the further doubling by 1.8, 1.8, and 1.9. The sub-proportional scaling reflects an algebraic feature of the model: trade flow scales with $\nu\omega$, while the profit and price-index responses scale with $\nu\omega + 1$ through the unit pass-through of the state-specific $(\tau^s)^{1-\sigma}$ component, so doubling ν less than doubles $\nu\omega + 1$.

Tariff revenue is the notable exception to this pattern. The cross-sector mean $\Delta \ln \mathbb{E}_t[R_{j,t+1}^g]$ under a pessimistic shift *declines* as θ rises, from $+4.82\%$ at $\theta = 5$ to $+4.12\%$ at $\theta = 10$ and $+2.70\%$ at $\theta = 20$. This is because $\Delta \ln R^g$ decomposes into an entry channel ($\chi\nu\omega \cdot \partial \ln \bar{\xi}/\partial \eta_p$, negative and amplified by ν) and a reallocation channel ($\chi\partial \ln \mathcal{T}/\partial \eta_p$, positive and *invariant* to θ). As θ rises, the (negative) entry channel grows in absolute value while the (positive) reallocation channel is unchanged, so the positive net effect contracts. This is not a contradiction with the scaling of trade flow but a property of how the two channels of R^g aggregate; the (small) positive tariff-revenue contribution to $\Delta \ln C_{jt}^*$ shrinks with θ , marginally amplifying the consumption loss but not enough to overturn the dominant role of the profit and price-index channels.

Table 6: Welfare predictions under three Pareto shape parameters

Quantity	Mean elasticity (% per $\Delta = 0.10$ shock)		
	$\theta = 5$	$\theta = 10$	$\theta = 20$
Valid observations	8,537	9,712	10,081
Coverage of trimmed BW	80.8%	91.9%	95.4%
σ range	[1.50, 5.99]	[1.50, 11.00]	[1.50, 20.83]
<i>Pessimistic shift (η_p)</i>			
$\Delta \ln X_{ij}$	-1.318	-3.438	-8.004
$\Delta \ln \mathbb{E}_t[\Pi_{j,t+1}]$	-1.035	-1.827	-3.316
$\Delta \ln \mathbb{E}_t[R_{j,t+1}^g]$	+4.817	+4.119	+2.703
$\Delta \ln P_{j,t+1}$	+0.577	+0.930	+1.664
$\Delta \ln C_{jt}^*$	-0.273	-0.468	-0.869
$\Delta \ln C_{j,t+1}^*$	-0.561	-0.934	-1.701
<i>Mean-preserving spread (η_d)</i>			
$\Delta \ln X_{ij}$	+0.086	+0.265	+0.689
$\Delta \ln \mathbb{E}_t[\Pi_{j,t+1}]$	+0.082	+0.170	+0.341
$\Delta \ln \mathbb{E}_t[R_{j,t+1}^g]$	-0.597	-0.696	-0.684
$\Delta \ln P_{j,t+1}$	-0.038	-0.065	-0.116
$\Delta \ln C_{jt}^*$	+0.017	+0.031	+0.061
$\Delta \ln C_{j,t+1}^*$	+0.036	+0.064	+0.119

Notes: Each column reports the cross-sector mean of $100 \times \partial \ln(\cdot) / \partial \eta$ evaluated at $\Delta = 0.10$, computed sector by sector using σ_m from Broda and Weinstein (2006) (trimmed to [1.5, 50]). The Pareto regularity condition $\theta > \sigma - 1$ excludes high- σ sectors. Other parameters are at their benchmarks: $\beta = 0.96$, $\gamma = 2$, $\bar{\tau} = 1.05$, $\delta = 0.05$, $\chi = 0.30$. All elasticities are reported at the country- j level: each bilateral building-block elasticity is multiplied by χ (the exposure mapping in equation (155)). $\Delta \ln C^*$ predictions are computed from Propositions 3 and 4 in Appendix A and include the endogenous response of tariff revenue $R_{j,t+1}^g$. The samples differ across columns because the regularity condition admits more sectors at higher θ ; a common-sample exercise restricted to $\sigma < 6$ delivers the same qualitative pattern (see footnote below).

The samples in Table 6 differ across columns because the Pareto regularity condition $\theta > \sigma - 1$ admits more sectors at higher θ . To verify that the monotone increase of welfare cost in θ is not driven by sample composition, I re-run the calibration on the common sample $\sigma < 6$ that is admissible for all three values of θ . The pessimistic-shift welfare cost on this common sample is -0.273% at $\theta = 5$, -0.503% at $\theta = 10$, and -0.964% at $\theta = 20$; the trade-flow response is -1.32% , -3.68% , and -8.41% , respectively. The common-sample magnitudes are in fact slightly larger than the full-sample ones at $\theta = 10, 20$, because the full sample includes high- σ sectors with small ν that pull the mean toward zero. The $\theta = 5$ benchmark therefore provides a conservative lower bound on the welfare cost of uncertainty operating through the profit and price-index channels, both within the benchmark sample and within the common-sample comparison. The larger θ values yield welfare costs that are more comparable to recent event-study estimates such as the Brexit-related losses documented by Hassan

et al. (2024) and the U.S.–China trade-war losses estimated by Fajgelbaum and Khandelwal (2022). Importantly, the qualitative ranking of channels does not change: in all three columns, the profit and price-index channels dominate the consumption response, and the tariff-revenue channel is small but non-negligible.

4.5 Sensitivity to the Baseline Tariff

I next vary $\bar{\tau}$ across a wide range while holding all other parameters at their benchmarks ($\theta = 5$, $\delta = 0.05$, $\chi = 0.30$). The exercise serves two purposes. First, it characterises how the welfare cost of uncertainty depends on the underlying trade-cost level, which is a key dimension of cross-sectional heterogeneity in the data. Second, it asks whether the model can rationalise—in sign and in unit-free attenuation pattern—the $\text{EPU} \times \sigma \times \text{Tariff}$ triple-interaction documented in Table 1.

Table 7 reports cross-sector mean elasticities at six baseline-tariff levels, decomposes the tariff-revenue and consumption responses into their components, and reports the implied $\text{EPU} \times \sigma$ regression slope at each value of $\bar{\tau}$. The columns with $\bar{\tau} \geq 1.30$ are illustrative of the nonlinear tariff-revenue channel rather than representative of contemporary applied tariffs.

Three findings stand out.

Trade flow scales with $1/\bar{\tau}$. The trade-flow response to a pessimistic shift attenuates from -1.32% at $\bar{\tau} = 1.05$ to -0.79% at $\bar{\tau} = 1.75$, a factor of $0.60 \approx 1.05/1.75$. This is the exact prediction of the leading-order expression $\partial \ln X / \partial \eta_p \propto 2\delta / \bar{\tau}$ from Proposition 1: holding δ fixed, doubling $\bar{\tau}$ roughly halves the pessimistic-channel response. The MPS response attenuates *faster* than the pessimistic response, because its leading-order term is $\propto \sigma \delta^2 / \bar{\tau}^2$ from Proposition 2: in Table 7 the MPS trade-flow elasticity falls from $+0.086\%$ to $+0.031\%$ over the same range, a factor of $0.36 \approx (1.05/1.75)^2$.

The Laffer-curve crossing in tariff revenue. The reallocation component of $\partial \ln R^g / \partial \eta_p$ collapses from $+5.21\%$ at $\bar{\tau} = 1.05$ to $+0.02\%$ at $\bar{\tau} = 1.75$, and turns negative for high- σ sectors above their respective Laffer peaks $\tau_\sigma^* = (\sigma - 1)/(\sigma - 2)$. Because the reallocation channel is the only positive contributor to $\Delta \ln R^g$ under a pessimistic shift (the entry channel is negative throughout), the net effect on $\Delta \ln R^g$ also collapses, from $+4.82\%$ at $\bar{\tau} = 1.05$ to -0.21% at $\bar{\tau} = 1.75$. The corresponding R^g contribution to $\Delta \ln C_{jt}^*$ shifts from $+0.035$ pp at $\bar{\tau} = 1.05$ to -0.019 pp at $\bar{\tau} = 1.75$: in low-tariff environments the tariff-revenue channel partially *offsets* the profit and price-index losses, while in high-tariff environments it *amplifies* them. Quantitatively, the offset at $\bar{\tau} = 1.05$ is roughly 21% of the gross profit loss ($+0.035$ on -0.166), and the amplification

Table 7: Welfare predictions across baseline tariff levels

Quantity	Baseline tariff $\bar{\tau}$					
	1.05	1.10	1.20	1.30	1.50	1.75
R^g /income implied	1.5%	2.9%	5.7%	8.3%	13.0%	18.4%
<i>Pessimistic shift (η_p)</i>						
$\Delta \ln X_{ij}$	-1.318	-1.257	-1.151	-1.062	-0.919	-0.787
$\Delta \ln \mathbb{E}_t[\Pi_{j,t+1}]$	-1.035	-0.987	-0.904	-0.833	-0.721	-0.617
$\Delta \ln \mathbb{E}_t[R_{j,t+1}^g]$	+4.817	+1.995	+0.601	+0.173	-0.117	-0.214
<i>entry channel</i>	-0.395	-0.377	-0.345	-0.319	-0.276	-0.236
<i>reallocation channel</i>	+5.212	+2.372	+0.947	+0.491	+0.159	+0.022
$\Delta \ln P_{j,t+1}$	+0.577	+0.550	+0.504	+0.465	+0.402	+0.345
$\Delta \ln C_{jt}^*$	-0.273	-0.262	-0.246	-0.232	-0.208	-0.186
Π contribution	-0.166	-0.156	-0.139	-0.125	-0.102	-0.082
R^g contribution	+0.035	+0.029	+0.017	+0.007	-0.008	-0.019
price contribution	-0.141	-0.135	-0.123	-0.114	-0.099	-0.084
$\Delta \ln C_{j,t+1}^*$	-0.561	-0.538	-0.498	-0.464	-0.409	-0.358
<i>Mean-preserving spread (η_d)</i>						
$\Delta \ln X_{ij}$	+0.086	+0.078	+0.066	+0.056	+0.042	+0.031
$\Delta \ln \mathbb{E}_t[R_{j,t+1}^g]$	-0.597	-0.235	-0.077	-0.032	-0.004	+0.004
$\Delta \ln C_{jt}^*$	+0.017	+0.016	+0.014	+0.012	+0.009	+0.007
<i>Implied empirical regression slope (cross-$\bar{\tau}$ pattern only)</i>						
$\partial^2 \ln X / \partial \eta_p \partial \sigma$ (slope)	+0.00468	+0.00447	+0.00409	+0.00378	+0.00327	+0.00280

Notes: All columns are computed at $\theta = 5$, $\beta = 0.96$, $\gamma = 2$, $\delta = 0.05$, $\chi = 0.30$, with σ_m from Broda and Weinstein (2006). Entries are cross-sector means of $100 \times \partial \ln(\cdot) / \partial \eta$ for $\Delta = 0.10$. All elasticities for Π , P , and R^g are reported at the country- j level (each bilateral elasticity multiplied by χ). The *entry channel* of $\Delta \ln \mathbb{E}_t[R^g]$ equals $\chi \cdot \nu \omega \cdot \partial \ln \bar{\xi} / \partial \eta_p$, which is χ times the bilateral trade-flow elasticity $\Delta \ln X_{ij}$. The implied regression slope in the last row is the linear-fit slope of $\partial \ln X / \partial \eta_p$ across $\sigma \in [2.5, 5.0]$. The level of this slope is in arbitrary units (EPU and probability shifts are not commensurable); only its cross- $\bar{\tau}$ attenuation pattern admits unit-free comparison with the empirical $\text{EPU} \times \sigma \times \text{Tariff}$ coefficient in Table 1.

at $\bar{\tau} = 1.75$ is roughly 23% of the gross profit loss (-0.019 on -0.082). The overall consumption-loss magnitude $|\Delta \ln C_{jt}^*|$ nonetheless decreases monotonically with $\bar{\tau}$, from 0.273% to 0.186%, because the attenuation of the profit and price-index channels dominates the amplification of the tariff-revenue channel.

Cross- $\bar{\tau}$ attenuation pattern of the $\text{EPU} \times \sigma$ slope. The model-implied $\text{EPU} \times \sigma$ slope falls from +0.00468 at $\bar{\tau} = 1.05$ to +0.00280 at $\bar{\tau} = 1.75$, attenuating by a factor of about 0.60 over the range. This attenuation pattern is unit-free: it is the proportional decline in the cross-sectoral slope of $\partial \ln X / \partial \eta_p$ as $\bar{\tau}$ rises, and does not depend on how the EPU index is scaled. The reduced-form triple-interaction coefficients $\hat{\beta}_5$ and $\hat{\beta}_6$ in Table 1 have the same sign and identify the same $\bar{\tau}$ -modulation of the cross-sectoral $\text{EPU} \times \sigma$ slope; thus the model reproduces the qualitative direction and the unit-free attenuation pattern of the data. I do not interpret the level of the model-implied slope as a direct

structural estimate of $\hat{\beta}_5$ or $\hat{\beta}_6$, because EPU units and probability shifts are not commensurable; the role of the calibration here is to demonstrate that the proposed mechanism is capable of generating the observed cross-sectoral and cross-tariff patterns, not to map a regression coefficient one-for-one.

The three sensitivity exercises in Sections 4.3–4.5 together deliver a coherent picture. The trade-flow response to uncertainty is mechanically governed by the Pareto-implied factor $\nu(\sigma, \theta)$ and the $1/\bar{\tau}$ scaling of the leading-order term; the consumption response builds on this through three channels whose relative weights depend on the income share s_R of tariff revenue and on the local curvature of $\mathcal{R}(\tau)$. In low-tariff environments, the profit and price-index channels dominate the consumption loss and the tariff-revenue channel acts as a partial buffer through the reallocation gain. In high-tariff environments, the reallocation gain shrinks (the Laffer-curve crossing) and eventually turns negative, so the tariff-revenue channel becomes an amplifier of consumption losses rather than a buffer. The matching between the model-implied and reduced-form attenuation patterns provides a unit-free cross-check of the mechanism against the empirical evidence of Section 3.2, while the absolute welfare magnitudes are best read as illustrative orders-of-magnitude rather than point estimates.

4.6 The Trump Experiment: A Quantitative Reality Check

The closed-form mechanism developed in Section 2 and Appendix A characterizes how trade flows, expected profits, tariff revenue, and welfare respond to first- and second-moment perturbations of the future trade-cost distribution. The 2025 U.S. trade-policy cycle provides a natural setting to test these predictions against realised data. From January to mid-2025, the U.S. average applied tariff rose from approximately 2.5% to 27%, while EPU reached a historical maximum in 2025. I treat that episode as a joint shock to the baseline tariff $\bar{\tau}$ and the future trade-cost distribution, and ask whether the model’s elasticities—when calibrated with empirically grounded parameters—generate predictions that are quantitatively recognizable in the realised data.

The pre-shock state is $(\bar{\tau}, \eta_p) = (1.025, 0)$, calibrated to the U.S. applied MFN tariff in January 2025. The post-shock state is $(\bar{\tau}, \eta_p) = (1.27, 0.10)$, with $\bar{\tau}$ matching the U.S. effective bilateral tariff by mid-2025 and $\eta_p = 0.10$ a moderate pessimistic shift consistent with the 2024–25 tariff rhetoric. The joint shock decomposes into a mechanical channel ($\bar{\tau}$ rising with $\eta_p = 0$), a pure uncertainty channel (η_p rising with $\bar{\tau}$ fixed at the pre-shock level), and an interaction term that captures their non-linear coupling.

The structural parameters are calibrated to U.S. macroeconomic data: substitution elasticity $\sigma = 4.174$ ¹, the cross-sectoral mean of Broda and Weinstein

¹The calibration uses the unweighted cross-sectoral mean of Broda–Weinstein elasticities ($\sigma = 4.174$) rather than a trade-weighted mean. Trade weights are endogenous to σ (high- σ sectors have larger trade flows by the Pareto gravity formula), so trade-weighting risks

(2006) estimates on U.S. HS3 sectors; imports-to-GDP ratio $\chi = 0.14$, from Bureau of Economic Analysis (2025); σ -dependent tariff-revenue and profit income shares $s_R(\sigma, \bar{\tau})$ and $s_\Pi(\sigma)$ as derived in Appendix A; and demand-share weight $\omega = 0.96$.

Two points should be made transparently before this number is used. First, the two-period model is a tractable truncation that is clean for deriving the mechanism, not for calibrating its horizon; I therefore adopt the infinite-horizon limit (Appendix B) for the quantitative exercise, since real export decisions are exposed to many future periods. Second, and importantly, the calibration’s most robust anchor — the tariff-revenue prediction — is essentially invariant to this choice (Table 9), because its level is dominated by the $\bar{\tau}$ -level reallocation term rather than the σ -scaled channel; the trade-flow figure is reported under with the $\omega = 0.49$ counterpart shown alongside it in Table 9 for full transparency. The choice of $\omega = 0.96$ deserves explicit justification. The two-period closed-form derivation of Section 2 implies $\omega \approx \beta/(1 + \beta) = 0.49$, reflecting that the firm’s expected value is split symmetrically between current and one-period-ahead profits. In a multi-period environment, where the firm’s value depends on many future periods of trade-cost realisations, the effective ω approaches β itself. Setting $\omega = 0.96$ in this calibration is a deliberate departure from the two-period derivation, motivated by the multi-period structure of real-world export decisions.² The two-period value would halve the trade and price elasticities and reduce the consumption response in similar proportion; embedding the multi-period structure explicitly through a dynamic Melitz model is the natural extension discussed in Section 5. Other parameters follow Section 4: $\theta = 5$, $\beta = 0.96$, $\gamma = 2$, $\delta = 0.05$.

Pass-through follows the separated specification of Cavallo et al. (2021), who document a sharp distinction between border-level and retail-level pass-through of tariffs. The firm’s entry decision and the producer-side bilateral price index use border-level pass-through $\rho_{\text{border}} = 1.0$, reflecting that U.S. importers face nearly the full statutory tariff at the border. The consumer-side contribution to PCE inflation uses retail-level pass-through $\rho_{\text{retail}} = 0.30$, reflecting that retailers, distributors, and exporters jointly absorb approximately 70% of the border tariff through margin compression before the price reaches the consumer. The tariff-revenue equilibrium uses the statutory rate τ regardless of pass-through, reflecting that the government collects on the statutory rate.³ The joint-shock

circularity in the calibration. The unweighted mean treats each HS3 sector as an independent draw from the underlying elasticity distribution.

²The calibration $\omega = 0.96$ corresponds to the infinite-horizon limit of the demand-share weight, $\omega^\infty = \beta = 0.96$, derived in Appendix B. The two-period model used in the main text is a tractable finite-truncation that yields $\omega^{2p} \approx \beta/(1 + \beta) \approx 0.49$ at the same discount factor; the calibration in this section adopts the infinite-horizon limit to abstract from this truncation. A complete dynamic Melitz extension with explicit firm entry, exit, and uncertainty dynamics is the subject of ongoing work.

³This separated specification follows Cavallo et al.(2021)’s headline finding that border-level pass-through to U.S. import prices was nearly 95% in the 2018 trade war, while retail-

decomposition is reported in Table 8 and visualised in Figure 3.

Table 8: Joint-shock decomposition at the 2025 U.S. baseline

	Mechanical	Uncertainty	Joint	Interaction
$\Delta \ln X_{ij}$	-37.56%	-1.28%	-38.84%	0.00%
$\Delta \ln \mathbb{E}[\Pi]$	-105.60%	-3.60%	-109.20%	0.00%
$\Delta \ln \mathbb{E}[R^g]$	+154.24%	+28.33%	+154.05%	-28.52%
$\Delta \ln P$ (border)	+33.27%	+1.13%	+34.40%	0.00%
$\Delta \ln P$ (retail)	+10.93%	+1.13%	+12.06%	0.00%
$\Delta \ln C^*$	-1.86%	-0.06%	-1.96%	-0.04%

Notes: Calibration: $\sigma = 4.174$ (Broda and Weinstein, 2006, U.S. sectoral mean), $\rho_{\text{border}} = 1.0$ and $\rho_{\text{retail}} = 0.30$ (Cavallo et al., 2021, separated pass-through), $\omega = 0.96$, $\chi = 0.14$, $\theta = 5$, $\delta = 0.05$, $\beta = 0.96$, $\gamma = 2$. Joint = Mechanical + Uncertainty + Interaction (entries closed to two-decimal precision). The interaction term for $\Delta \ln \mathbb{E}[R^g]$ is large in magnitude because the mechanical and uncertainty channels each push $\bar{\tau}$ closer to the cross-sector Laffer peak, attenuating each other's marginal revenue contribution. Consumption decomposition: Π contribution uses ρ_{border} ; R^g contribution uses the statutory rate; price contribution uses ρ_{retail} . Income shares: $s_{\Pi}(\sigma, \bar{\tau}_{\text{avg}}) = 0.235$, $s_R(\sigma, \bar{\tau}_{\text{avg}}) = 0.020$.

Three findings emerge. The mechanical channel dominates each macro aggregate. Raising $\bar{\tau}$ from 1.025 to 1.27 contracts trade flows by 37.6%, raises expected tariff revenue by 154.2%, and raises the border-level bilateral import price by 33.3%. The pure uncertainty channel adds a second-order modification: trade flows contract by an additional 1.3 percentage points, tariff revenue rises by an additional 28.3 percentage points, and the import price rises by an additional 1.1 percentage points.

The interaction term on tariff revenue is negative (-28.5 percentage points). The pure uncertainty channel's contribution to R^g is attenuated when the mechanical channel has already pushed $\bar{\tau}$ closer to the cross-sector Laffer peak. This peak is defined as the maximiser of the cross-sectoral average of the per-sector revenue function $\mathcal{R}_{\sigma}(\bar{\tau}) = \chi(\bar{\tau} - 1)\bar{\tau}^{1-\sigma}$ and lies at $\bar{\tau} \approx 2.2$ in the calibrated panel. This is substantially above the closed-form per-sector peak $\tau_{\sigma_{\text{med}}}^* = (\sigma_{\text{med}} - 1)/(\sigma_{\text{med}} - 2) \approx 1.72$ at the BW median elasticity; the difference reflects that low- σ sectors with τ_{σ}^* far above unity contribute disproportionately to the cross-sector sum (Figure 2). At the post-shock $\bar{\tau} = 1.27$, the same pessimistic probability shift generates a smaller reallocation-channel response than at the pre-shock $\bar{\tau} = 1.025$. This is a direct empirical implication of the Laffer geometry developed in Section 4.5.

I use 2025Q2–Q3 data as the comparison window to exclude three large non-trade shocks that contaminate the macroeconomic series beyond that window:

level pass-through to consumer prices was only 4–7% per unit tariff (approximately 0.30 of the affected-margin pass-through). Applying a single ρ across both channels, as some calibrations do, would conflate the firm-side trade-cost shock with the consumer-side price burden.

the 43-day federal-government shutdown of October–November 2025, the deterioration of the labour market with unemployment rising to 4.4% in September, and PCE-inflation acceleration to +4.5% in 2026Q1. Table 9 summarises the comparison.

Table 9: Joint-shock prediction versus realised 2025 data

Anchor	$\omega = 0.96$ (baseline)		$\omega = 0.49$ (two-period)		Realised
	Model	Gap	Model	Gap	
$\Delta \ln X_{ij}$, 2025Q2 SAAR real goods imports	−38.8%	3.8 pp	−19.8%	−15.2 pp	−35.0%
$\Delta \ln \mathbb{E}[R^g]$, FY24→FY25 customs receipts	+154.1%	4.1 pp	+173.1%	+23.1 pp	+150.0%
$\Delta \ln P \times \chi$, 2025Q3 PCE price index SAAR	+1.7%	−1.1 pp	+1.4%	−1.4 pp	+2.8%
$\Delta \ln C^*$, 2025Q3 SAAR real PCE	−2.0%	−5.5 pp	−1.6%	−5.1 pp	+3.5%

Notes: Model values computed under separated Cavallo pass-through ($\rho_{\text{border}} = 1.0$ for firm-side, $\rho_{\text{retail}} = 0.30$ for consumer-side). The $\omega = 0.96$ column is the infinite-horizon baseline used throughout this section (joint-shock column of Table 8); the $\omega = 0.49$ column reports the same calibration under the two-period demand-share weight $\omega \approx \beta/(1 + \beta)$ of Section 2, shown for comparison. The trade-flow and price anchors scale down by the factor $\omega^{2p}/\omega^\infty \approx 0.51$ derived in Appendix B; the tariff-revenue anchor instead rises, because its level is dominated by the $\bar{\tau}$ -level reallocation term rather than the ω -scaled channel. Realised values are seasonally-adjusted annualised rates (SAAR) from BEA Q2 and Q3 2025 GDP releases (Third Estimate for Q2 2025; Initial Estimate for Q3 2025), NIPA Table 1. The tariff-revenue anchor is the fiscal-year aggregate change (FY 2024 \$78 billion vs FY 2025 \$195 billion; U.S. Customs and Border Protection, 2025; Committee for a Responsible Federal Budget, 2025).

The model matches the realised data in sign on all four anchors. Three of them—trade flow, tariff revenue, and PCE inflation—match in order of magnitude with gaps below 5 percentage points; the consumption gap of 5.5 percentage points is addressed separately below. The trade-flow match at 3.8 percentage points is the closest in absolute terms: the model predicts $\Delta \ln X = -38.8\%$ against a realised -35.0% annualised contraction in real U.S. goods imports in 2025Q2 (Bureau of Economic Analysis, 2025). This match is not the product of parameter tuning— $\sigma = 4.174$ is taken from Broda and Weinstein (2006), $\rho_{\text{border}} = 1.0$ from Cavallo et al. (2021), and $\chi = 0.14$ from Bureau of Economic Analysis (2025) data; the multi-period demand-share weight $\omega = 0.96$ is the only parameter set by modelling convention rather than direct estimation. The match is a central quantitative anchor of this section.

The tariff-revenue prediction closely matches the realised data. The model predicts $\Delta \ln \mathbb{E}[R^g] \approx +154\%$ against a realised $+150\%$ increase between FY 2024 and FY 2025 (U.S. Customs and Border Protection, 2025). The gap of 4.1 percentage points is robust to two assumptions that affect the other anchors: it is invariant to the pass-through coefficient (the government collects on the statutory rate regardless of exporter or retailer absorption), and it is robust to the choice between two-period $\omega = 0.49$ and infinite-horizon $\omega = 0.96$ (because

the reallocation channel, which accounts for most of the level, does not scale with ω).

For PCE inflation, the gap is 1.1 percentage points. The model’s direct prediction for the PCE-inflation contribution is $\chi \cdot \Delta \ln P_{\text{retail}} = 0.14 \times 12.1\% \approx +1.7\%$, using the model-implied retail-CES price index under $\rho_{\text{retail}} = 0.30$.⁴ The realised annualised PCE price index change in 2025Q3 is +2.8% (Bureau of Economic Analysis, Q3 2025 GDP Initial Estimate). Subtracting the Federal Reserve’s 2% annual inflation target as the underlying trend leaves a realised tariff-attributable component of approximately +0.8%, against the model’s +1.7% prediction. The model over-predicts the tariff-attributable component by approximately 0.9 percentage points, consistent with the partial-equilibrium model’s abstraction from supply-chain inventory adjustment and foreign-exporter price absorption documented by Cavallo et al. (2021).

The fourth anchor, the consumption response, has the right sign in its components but exhibits a 5.5 percentage-point sign-reversed gap from the realised +3.5% annualised Q3 real PCE growth. The model’s predicted consumption response of -2.0% under the joint shock is small in absolute terms—an order of magnitude smaller than the predicted responses on trade volume and the bilateral import price. This is informative rather than disappointing; it reflects three structural features of the U.S. economy that together compress the direct transmission channel from tariffs to retail consumption.

The compression arises from three multiplicative income-share weights. U.S. imports account for only $\chi = 0.14$ of GDP; retail pass-through of border tariffs is approximately $\rho_{\text{retail}} = 0.30$ (Cavallo et al., 2021), with the remainder absorbed by foreign exporters and retail markup compression; and tariff revenue accounts for only $s_R \approx 0.02$ of household disposable income at the post-shock baseline. The price contribution to consumption is attenuated by both χ and ρ_{retail} to -0.41 percentage points; the profit contribution at full border pass-through is -1.76 percentage points; and the household-level weight on the +154% expected-tariff-revenue increase evaluates to only +0.21 percentage points via $\rho_t \cdot s_R \cdot \Delta \ln R^g$.

The cumulative effect of these three structural attenuators is to compress the direct trade-uncertainty channel into a consumption response of order -2% . This contrasts with the response sizes typically estimated in quantitative trade-war analyses that assume complete pass-through to consumers or that use cross-country average import shares around 0.30 rather than the U.S.-specific 0.14. The finding is also consistent with the political-economy observation that the 2025 U.S. tariff cycle did not generate immediate retail-side consumption dis-

⁴A simpler linear decomposition gives $\rho_{\text{retail}} \cdot \chi \cdot \Delta \ln P_{\text{border}} = 0.30 \times 0.14 \times 34.4\% \approx +1.4\%$. This is a first-order approximation that omits the non-linearity of the retail-CES aggregator; the two are equivalent in the limit of infinitesimal tariff changes but diverge as $\bar{\tau}$ moves substantially. The model output uses the exact CES index throughout.

ruption: the direct transmission channel from tariff to retail consumption is structurally weak in a relatively closed economy with incomplete retail pass-through and limited fiscal redistribution of tariff revenue.

The full welfare cost of the tariff cycle includes indirect channels not captured in the present closed-form mechanism: firm-profit declines transmitted through investment and employment, real-options effects on capital expenditure (Bloom, 2009), the breakdown of the integrated-ownership rebate convention when tariff revenue is retained as deficit reduction (Committee for a Responsible Federal Budget, 2025), and the long-run reorganisation of global trade flows. These channels lie outside the analytical scope of the closed-form mechanism and represent the natural direction for the general-equilibrium extension discussed in Section 5.

The model’s -2.0% partial-equilibrium prediction also excludes contemporaneous positive shocks to disposable income that lie outside the trade block. Information-processing equipment and software investment grew at an annualised rate of approximately 39% in 2025H1 (Bureau of Economic Analysis, 2025, NIPA Table 5.5.5), making AI-related capital expenditure the dominant positive contributor to U.S. GDP growth in the relevant window. Labour income remained positive through 2025Q3, with monthly payroll gains exceeding 150,000 in the first three quarters of the year (Bureau of Labor Statistics, 2025). U.S. equity-market valuations reached historical peaks during 2025Q2–Q3, generating a wealth effect on consumption that the model does not capture. The model identifies the trade-uncertainty channel’s contribution alone, holding these other determinants of disposable income fixed; the realised $+3.5\%$ annualised Q3 PCE growth reflects the joint action of these positive channels alongside the trade-uncertainty channel that the model formalises.

The Trump experiment yields correctly-signed magnitude predictions: the model is within 5 percentage points of realised data on three of four anchors, with the trade-flow response within 3.8 percentage points of the realised -35.0% contraction. The tariff-revenue match is the model’s strongest quantitative claim. It is invariant to pass-through, only mildly sensitive to firm-horizon assumptions, and reflects the Laffer-geometry mechanism developed in Section 4.5. The exercise also exposes the structural smallness of the direct consumption channel, where the low U.S. import share, incomplete retail pass-through, and modest tariff-revenue income share jointly compress the partial-equilibrium response to order -2% . This identifies the boundary of the closed-form mechanism and motivates specific extensions: explicit fiscal closure for the rebate convention, real-options effects on investment for indirect channels through firm profits, and a multi-period dynamic Melitz block to relax the two-period ω approximation.

The exercise tests the model against the 2025 U.S. tariff episode. A full forecast of the realised series would require a richer general-equilibrium framework, beyond the scope of this paper.

5 Conclusion

This paper has studied how trade-related policy uncertainty affects bilateral trade flows and consumption through a dynamic extension of the Melitz model. Building on the framework of Baley, Veldkamp, and Waugh (2020), I embed uncertainty over future bilateral trade costs into a heterogeneous-firm export-entry model and identify two opposing channels through which uncertainty operates. A pessimistic shift in the expected distribution of future trade costs raises the export cutoff and contracts trade, while a mean-preserving spread in the same distribution lowers the cutoff and expands trade through the convexity of firm payoffs in $\tau^{1-\sigma}$. Under Pareto productivity, the relative strength of the two channels admits closed-form expressions: the pessimistic effect scales with $(\theta - \sigma + 1)$ and the mean-preserving-spread effect with $\sigma(\theta - \sigma + 1)$, so the ratio of the convexity gain to the pessimistic loss rises monotonically with the elasticity of substitution. Higher baseline trade costs attenuate both channels, with the dispersion channel attenuating faster.

The empirical analysis shows patterns consistent with these predictions in a panel of bilateral, sector-level UN Comtrade trade flows. The PPML baseline regression with full multi-interaction structure delivers a positive and statistically significant $\text{EPU} \times \sigma$ coefficient on both the exporter and the importer side, indicating that high-elasticity sectors are systematically less adversely affected by uncertainty. The triple interaction with bilateral tariffs is negative and significant on the importer side, confirming the model’s prediction that trade costs dampen the convexity channel more strongly than the pessimistic channel. The pattern is robust to high-dimensional fixed effects across most specifications, business-cycle controls, and lagged uncertainty measures, with the exception of the saturated pair-plus-HS column where identifying variation is mechanically limited. A complementary calibration on Broda–Weinstein elasticities reproduces the sign and the cross-tariff attenuation pattern of the reduced-form estimates, lending structural support to the proposed mechanism.

Several modelling choices warrant scrutiny. The two-period structure simplifies the firm’s expected-value calculation but understates the exposure to future trade-cost realisations; the calibration in Section 4.6 partially corrects this by setting $\omega = 0.96$ rather than the two-period value $\beta/(1 + \beta)$. The welfare exercise relies on an integrated-ownership rebate convention that the realised 2025 fiscal regime breaks, and the exposure parameter χ aggregates bilateral elasticities to country-level magnitudes without distinguishing import composition.

The natural extension is to embed the closed-form mechanism in a richer general-equilibrium environment. The most consequential modelling addition is fiscal closure for the rebate convention: the present model’s positive consumption response under uncertainty depends on tariff revenue being rebated to households, which the 2025 episode shows is unrealistic. A multi-period dynamic Melitz block would relax the two-period ω approximation and let real-

options effects on investment (Bloom, 2009) enter the firm's entry decision. The σ -heterogeneity prediction also provides a natural test moment for any quantitative extension: if the closed-form scaling $\sigma(\theta - \sigma + 1)$ vs $(\theta - \sigma + 1)$ holds in a fully dynamic model, the cross-sectoral PPML coefficients estimated here should be reproduced as targets.

A Government Tariff Revenue and Welfare under Endogenous R^g

This appendix extends the model of Section 2 by allowing government tariff revenue R^g to respond endogenously to the uncertainty parameter η , and derives the resulting comparative statics for consumption and welfare. The trade-flow results of Propositions 1 and 2 carry over directly through the profit and price-index channels; the new ingredient is a tax-base reallocation channel in tariff revenue whose sign depends on the local curvature of the per-unit revenue function. This is the sole source of indeterminacy in the welfare response.

All comparative statics below are conditional on fixed general-equilibrium objects (wages, current trade costs, the current-period price index, and aggregate expenditure), in the same partial-equilibrium spirit as Propositions 1 and 2. A fully endogenous GE treatment is left for future work.

A.1 Tariff Revenue under Pareto Productivity

Recall from Section 2.2.1 that $\tau_{ijt} = 1 + t_{ijt}$ is the gross bilateral trade cost factor, with t_{ijt} the ad valorem tariff rate. The government in country j collects tariff revenue from imports as

$$R_{jt}^g = \sum_{i \in S} (\tau_{ijt} - 1) \int_{\Omega_{ijt}} p_{ijt}(\omega) q_{ijt}(\omega) d\omega = \sum_{i \in S} (\tau_{ijt} - 1) X_{ijt}. \quad (\text{A.1})$$

At period $t + 1$, expected tariff revenue from origin i is

$$\mathbb{E}_t[R_{ij,t+1}^g] = \sum_{s \in \{H, M, L\}} \lambda^s(\eta) (\tau_{ij,t+1}^s - 1) X_{ij,t+1}(s; \eta). \quad (\text{A.2})$$

Under Pareto productivity with $\theta > \sigma - 1$, the period- $t + 1$ trade flow conditional on state $s \in \{H, M, L\}$ admits

$$X_{ij,t+1}(s) = \Lambda_{ij,t+1} \cdot (\tau_{ij,t+1}^s)^{1-\sigma}, \quad \Lambda_{ij,t+1} \equiv \kappa M_{it} \varphi_{\min,i}^\theta f_{ij}^{-\nu} \tilde{A}_{ij,t+1} \Phi_{ijt}^\nu. \quad (\text{A.3})$$

The scale factor $\Lambda_{ij,t+1}$ contains Φ_{ijt} rather than $\Phi_{ij,t+1}$ because, given the entry timing in Section 2.2.3, the period- t entry decision pins down the mass of firms serving market j in both periods t and $t + 1$; the state-specific trade-cost component $(\tau_{ij,t+1}^s)^{1-\sigma}$ then enters multiplicatively in period $t + 1$.

Substituting (A.3) into (A.2) yields the bilateral Pareto closed form for expected tariff revenue:

$$\boxed{\mathbb{E}_t[R_{ij,t+1}^g] = \Lambda_{ij,t+1} \cdot \mathcal{T}_{ij,t+1}, \quad \mathcal{T}_{ij,t+1} \equiv \sum_{s \in \{H, M, L\}} \lambda^s \mathcal{R}_{ij,t+1}^s,} \quad (\text{A.4})$$

where

$$\mathcal{R}_{ij,t+1}^s \equiv (\tau_{ij,t+1}^s - 1)(\tau_{ij,t+1}^s)^{1-\sigma} \quad (\text{A.5})$$

is the per-unit revenue weight in state s . The object $\mathcal{T}_{ij,t+1}$ summarises how the probability vector shapes per-unit revenue, while $\Lambda_{ij,t+1}$ summarises how period- t entry shapes the tax base. The country-level object is $\mathbb{E}_t[R_{j,t+1}^g] = \sum_{i \in S} \mathbb{E}_t[R_{ij,t+1}^g]$.

Two-channel decomposition. Both $\Lambda_{ij,t+1}$ and $\mathcal{T}_{ij,t+1}$ depend on η through $(\lambda^H, \lambda^M, \lambda^L)$: the former indirectly through Φ_{ijt} , the latter directly through the probability weights on per-unit revenue. Differentiating (A.4) and using $\partial \ln \Lambda_{ij,t+1} / \partial \eta = \nu \omega_{ijt} \partial \ln \bar{\xi}_{ij,t+1} / \partial \eta$,

$$\boxed{\frac{\partial \ln \mathbb{E}_t[R_{ij,t+1}^g]}{\partial \eta} = \underbrace{\nu \omega_{ijt} \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta}}_{\text{(I) cutoff/entry channel}} + \underbrace{\frac{\partial \ln \mathcal{T}_{ij,t+1}}{\partial \eta}}_{\text{(II) tax-base reallocation}},} \quad (\text{A.6})$$

with ω_{ijt} and $\bar{\xi}_{ij,t+1}$ defined as in Propositions 1 and 2.

Channel (I) inherits the sign of $\partial \ln \bar{\xi} / \partial \eta$: a pessimistic shift lowers $\bar{\xi}$ and contracts entry, shrinking the tax base; a mean-preserving spread raises $\bar{\xi}$ through convexity and expands entry. Channel (II) depends on the local curvature of the per-unit revenue function $\mathcal{R}(\tau) \equiv (\tau - 1)\tau^{1-\sigma}$. Direct differentiation gives

$$\mathcal{R}'(\tau) = \tau^{-\sigma} [(2-\sigma)\tau + (\sigma-1)], \quad \mathcal{R}''(\tau) = (\sigma-1)\tau^{-\sigma-1} [(\sigma-2)\tau - \sigma]. \quad (\text{A.7})$$

The shape of $\mathcal{R}$ depends on whether σ exceeds two. For $\sigma > 2$, $\mathcal{R}$ admits a Laffer peak at $\tau_\sigma^* \equiv (\sigma-1)/(\sigma-2)$ and is concave on $(1, \sigma/(\sigma-2))$, switching to convex beyond. For $1 < \sigma \leq 2$, the bracket $(\sigma-2)\tau - \sigma$ is strictly negative for all $\tau > 1$, so $\mathcal{R}$ is strictly concave on $\tau > 1$; the sign of $\mathcal{R}'(\tau)$ is the sign of $(2-\sigma)\tau + (\sigma-1)$, which is positive at $\tau = 1$ and remains positive for all $\tau > 1$ when $\sigma \leq 2$, so $\mathcal{R}$ is strictly increasing without a Laffer peak.

Combining the sign of $\mathcal{R}'$ and $\mathcal{R}''$ across both ranges, in the empirically relevant region $\bar{\tau} < \tau_\sigma^*$ (for $\sigma > 2$) or $\bar{\tau} > 1$ (for $1 < \sigma \leq 2$), $\mathcal{R}$ is increasing, so a pessimistic shift raises $\mathcal{T}$; and $\mathcal{R}$ is concave (over the relevant subset of $\bar{\tau}$), so a mean-preserving spread lowers $\mathcal{T}$ by Jensen's inequality. In both regimes, channels (I) and (II) work in opposite directions, and the overall sign of $\partial \ln \mathbb{E}_t[R_{ij,t+1}^g] / \partial \eta$ is determined by the relative magnitude of ν and the local curvature of $\mathcal{R}(\tau)$ over the trade-cost support.

Taylor expansion under the symmetric parameterization. Under $\tau_{ij,t+1}^H = \bar{\tau} + \delta$, $\tau_{ij,t+1}^M = \bar{\tau}$, $\tau_{ij,t+1}^L = \bar{\tau} - \delta$, second-order Taylor expansion of $\mathcal{R}(\bar{\tau} \pm \delta)$ around $\bar{\tau}$ gives

$$\mathcal{R}(\bar{\tau} + \delta) + \mathcal{R}(\bar{\tau} - \delta) - 2\mathcal{R}(\bar{\tau}) \approx \mathcal{R}''(\bar{\tau}) \delta^2, \quad (\text{A.8})$$

$$\mathcal{R}(\bar{\tau} + \delta) - \mathcal{R}(\bar{\tau} - \delta) \approx 2\mathcal{R}'(\bar{\tau}) \delta. \quad (\text{A.9})$$

For the pessimistic shift ($d\lambda^H/d\eta_p = a$, $d\lambda^L/d\eta_p = -a$),

$$\frac{\partial \ln \mathcal{T}_{ij,t+1}}{\partial \eta_p} \approx \frac{2a \delta \mathcal{R}'(\bar{\tau})}{\mathcal{R}(\bar{\tau})} = \frac{2a \delta}{\bar{\tau} - 1} \cdot \frac{(2 - \sigma)\bar{\tau} + (\sigma - 1)}{\bar{\tau}}. \quad (\text{A.10})$$

For the mean-preserving spread ($d\lambda^H/d\eta_d = d\lambda^L/d\eta_d = a$, $d\lambda^M/d\eta_d = -2a$),

$$\frac{\partial \ln \mathcal{T}_{ij,t+1}}{\partial \eta_d} \approx \frac{a \delta^2 \mathcal{R}''(\bar{\tau})}{\mathcal{R}(\bar{\tau})} = \frac{a(\sigma - 1) \delta^2}{(\bar{\tau} - 1) \bar{\tau}^2} \cdot [(\sigma - 2)\bar{\tau} - \sigma]. \quad (\text{A.11})$$

Expression (A.10) is $O(\delta)$ and positive whenever $\mathcal{R}$ is increasing at $\bar{\tau}$; expression (A.11) is $O(\delta^2)$ and negative whenever $\mathcal{R}$ is concave at $\bar{\tau}$. Both conditions hold in the empirically relevant trade-cost range identified above.

Remark on the limit $\bar{\tau} \rightarrow 1$. The denominators in (A.10)–(A.11) carry $\bar{\tau} - 1$ because $\mathcal{R}(\bar{\tau}) = (\bar{\tau} - 1)\bar{\tau}^{1-\sigma} \rightarrow 0$ as $\bar{\tau} \rightarrow 1$, so the *log* elasticity $\partial \ln \mathcal{T}/\partial \eta$ diverges. This is a mechanical feature of taking logs of an object that approaches zero, not an economic singularity: in the welfare expressions below, $\partial \ln \mathcal{T}/\partial \eta$ is always multiplied by the revenue share s_R , which scales as $\bar{\tau} - 1$ near $\bar{\tau} = 1$, so the product $s_R \cdot \partial \ln \mathcal{T}/\partial \eta$ has a well-defined finite limit corresponding to the level effect on R^g .

A.2 Consumption and Welfare

Bilateral building blocks. Three bilateral elasticities drive the consumption response. First, the bilateral profit elasticity. Aggregating firm-level operating profits across exporters from i serving j and using (??), expected bilateral profits satisfy $\mathbb{E}_t[\Pi_{ij,t+1}] = (1/\sigma) \Lambda_{ij,t+1} \bar{\xi}_{ij,t+1}$, so that

$$\frac{\partial \ln \mathbb{E}_t[\Pi_{ij,t+1}]}{\partial \eta} = (\nu \omega_{ijt} + 1) \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta}, \quad (\text{A.12})$$

where $\nu \omega_{ijt}$ is the trade-flow elasticity through $\Lambda_{ij,t+1}$, and the additional $+1$ reflects the unit pass-through of the state-specific $(\tau^s)^{1-\sigma}$ component through $\bar{\xi}_{ij,t+1}$. Second, the bilateral price-index elasticity, obtained by analogous aggregation of the bilateral ij contribution to $P_{j,t+1}^{1-\sigma}$ and log-differentiating:

$$\frac{\partial \ln P_{j,t+1}}{\partial \eta} = \frac{1}{1 - \sigma} (\nu \omega_{ijt} + 1) \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta}. \quad (\text{A.13})$$

Third, the bilateral tariff-revenue elasticity from (A.6):

$$\frac{\partial \ln \mathbb{E}_t[R_{ij,t+1}^g]}{\partial \eta} = \nu \omega_{ijt} \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta} + \frac{\partial \ln \mathcal{T}_{ij,t+1}}{\partial \eta}. \quad (\text{A.14})$$

Calibration-based exposure mapping. To map the three bilateral elasticities into country- j -level objects required by the welfare formula (A.18)–(A.19) below, I adopt a single calibration-based exposure parameter $\chi \in (0, 1)$, equal to country j 's imports from i relative to GDP. The mapping is

$$\frac{\partial \ln \mathbb{E}_t[\Pi_{j,t+1}]}{\partial \eta} \approx \chi \frac{\partial \ln \mathbb{E}_t[\Pi_{ij,t+1}]}{\partial \eta}, \quad \frac{\partial \ln P_{j,t+1}}{\partial \eta} \approx \chi \frac{\partial \ln P_{ij,t+1}}{\partial \eta}, \quad \frac{\partial \ln \mathbb{E}_t[R_{j,t+1}^g]}{\partial \eta} \approx \chi \frac{\partial \ln \mathbb{E}_t[R_{ij,t+1}^g]}{\partial \eta}. \quad (\text{A.15})$$

Equation (A.15) is *not* a strict aggregation identity: while $R_{j,t+1}^g = \sum_{i \in S} R_{ij,t+1}^g$ holds by accounting, the country-level log elasticity equals the bilateral log elasticity only if the uncertainty shock acts on a single bilateral pair while other origins are held fixed. I adopt (A.15) as the simplest calibration-based exposure mapping consistent with the assumption that the EPU shock is interpreted bilaterally. In the calibration of Section 4, $\chi = 0.30$ corresponds to the OECD-average open-economy import share.

Optimal consumption. Combining the intertemporal budget constraint with the Euler equation,

$$C_{jt}^* = \frac{w_{jt}L_j + \Pi_{jt} + R_{jt}^g + \frac{1}{1+r_t}(w_{j,t+1}L_j + \mathbb{E}_t[\Pi_{j,t+1}] + \mathbb{E}_t[R_{j,t+1}^g])}{P_{jt} + \frac{P_{j,t+1}}{1+r_t}[\beta(1+r_t)P_{jt}/P_{j,t+1}]^{1/\gamma}}, \quad (\text{A.16})$$

and

$$C_{j,t+1}^* = [\beta(1+r_t)P_{jt}/P_{j,t+1}]^{1/\gamma} C_{jt}^*. \quad (\text{A.17})$$

Writing $C_{jt}^* = N_t/D_t$ and differentiating with $R_{j,t+1}^g$ allowed to respond endogenously to η delivers

$$\frac{\partial \ln C_{jt}^*}{\partial \eta} = \rho_t \left[s_\Pi \frac{\partial \ln \mathbb{E}_t[\Pi_{j,t+1}]}{\partial \eta} + s_R \frac{\partial \ln \mathbb{E}_t[R_{j,t+1}^g]}{\partial \eta} \right] - q_t \left(1 - \frac{1}{\gamma} \right) \frac{\partial \ln P_{j,t+1}}{\partial \eta}, \quad (\text{A.18})$$

and via the Euler equation,

$$\frac{\partial \ln C_{j,t+1}^*}{\partial \eta} = \rho_t \left[s_\Pi \frac{\partial \ln \mathbb{E}_t[\Pi_{j,t+1}]}{\partial \eta} + s_R \frac{\partial \ln \mathbb{E}_t[R_{j,t+1}^g]}{\partial \eta} \right] - \tilde{q}_t \frac{\partial \ln P_{j,t+1}}{\partial \eta}, \quad (\text{A.19})$$

where $\rho_t = q_t = \beta/(1+\beta)$ in steady state, $\tilde{q}_t = q_t + (1-q_t)/\gamma$, and the income shares of profits and tariff revenue in total country- j income are

$$s_\Pi = \frac{1/\sigma}{(\sigma-1)/\sigma + 1/\sigma + (\bar{\tau}-1)\chi}, \quad s_R = \frac{(\bar{\tau}-1)\chi}{(\sigma-1)/\sigma + 1/\sigma + (\bar{\tau}-1)\chi}. \quad (\text{A.20})$$

Substituting (A.12)–(A.14) through (A.15) into (A.18)–(A.19),

$$\begin{aligned} \frac{\partial \ln C_{jt}^*}{\partial \eta} = & \chi (\nu \omega_{ijt} + 1) \left[\rho_t s_{\Pi} - \frac{q_t(1 - 1/\gamma)}{1 - \sigma} \right] \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta} \\ & + \rho_t s_R \chi \left[\nu \omega_{ijt} \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta} + \frac{\partial \ln \mathcal{T}_{ij,t+1}}{\partial \eta} \right], \end{aligned} \quad (\text{A.21})$$

$$\begin{aligned} \frac{\partial \ln C_{j,t+1}^*}{\partial \eta} = & \chi (\nu \omega_{ijt} + 1) \left[\rho_t s_{\Pi} - \frac{\tilde{q}_t}{1 - \sigma} \right] \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta} \\ & + \rho_t s_R \chi \left[\nu \omega_{ijt} \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta} + \frac{\partial \ln \mathcal{T}_{ij,t+1}}{\partial \eta} \right]. \end{aligned} \quad (\text{A.22})$$

The first line of each expression collects the profit and price-index channels, whose sign is pinned down by $\partial \ln \bar{\xi} / \partial \eta$ alone. The second line is the tariff-revenue channel, which carries an additional reallocation term $\partial \ln \mathcal{T} / \partial \eta$ whose sign depends on the local curvature of $\mathcal{R}(\tau)$.

Proposition 3 (Pessimistic shift, welfare under endogenous tariff revenue). *Under Pareto productivity with $\theta > \sigma - 1$, endogenous $R_{j,t+1}^g$, exposure $\chi \in (0, 1)$, and the pessimistic-shift parameterization of Proposition 1:*

(i) *The welfare elasticities admit the closed forms (A.21)–(A.22) with $\eta = \eta_p$ and $\partial \ln \bar{\xi} / \partial \eta_p < 0$.*

(ii) *Decomposing the response of current consumption,*

$$\begin{aligned} \frac{\partial \ln C_{jt}^*}{\partial \eta_p} = & \underbrace{\chi (\nu \omega_{ijt} + 1) \left[\rho_t s_{\Pi} - \frac{q_t(1 - 1/\gamma)}{1 - \sigma} \right] \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta_p}}_{(\text{A}) < 0 \text{ for } \gamma \geq 1} + \underbrace{\rho_t s_R \chi \frac{\partial \ln \mathbb{E}_t[R_{ij,t+1}^g]}{\partial \eta_p}}_{(\text{B}) \text{ sign ambiguous}}, \end{aligned} \quad (\text{A.23})$$

and analogously for $C_{j,t+1}^*$ with $q_t(1 - 1/\gamma)$ replaced by $\tilde{q}_t$ (term (A) is then strictly negative for any $\gamma > 0$). Term (B) decomposes via (A.14) into an entry channel (< 0 , inheriting the sign of $\partial \ln \bar{\xi} / \partial \eta_p$) and a reallocation channel (> 0 over the empirically relevant trade-cost range identified in A.1); the relative magnitudes determine the overall sign.

(iii) *Under the symmetric parameterization with $d\lambda^H/d\eta_p = a > 0$ and $d\lambda^L/d\eta_p = -a$, a first-order Taylor expansion in δ yields*

$$\begin{aligned} \frac{\partial \ln C_{jt}^*}{\partial \eta_p} \approx & -\chi (\nu \omega_{ijt} + 1) \left[\rho_t s_{\Pi} - \frac{q_t(1 - 1/\gamma)}{1 - \sigma} \right] \frac{2a\delta(\sigma - 1)}{\bar{\tau}} \\ & + \rho_t s_R \chi \left[-\nu \omega_{ijt} \frac{2a\delta(\sigma - 1)}{\bar{\tau}} + \frac{2a\delta}{\bar{\tau} - 1} \cdot \frac{(2 - \sigma)\bar{\tau} + (\sigma - 1)}{\bar{\tau}} \right]. \end{aligned} \quad (\text{A.24})$$

Both terms are $O(\delta)$. In the low-tariff limit $\bar{\tau} \rightarrow 1$, $s_R \rightarrow 0$ and term (B) vanishes, so the welfare response converges to the exogenous-revenue benchmark and is strictly negative. As $\bar{\tau}$ rises, s_R grows and the reallocation channel becomes quantitatively relevant; the overall sign then depends on the relative magnitudes of these $O(\delta)$ terms.

Proof. The closed forms follow from substituting (A.12)–(A.14) through (A.15) into (A.18)–(A.19). Term (A) is strictly negative because the bracketed coefficient is strictly positive for $\gamma \geq 1$ (since $1 - \sigma < 0$ makes the second contribution positive), $\chi(\nu\omega_{ijt} + 1) > 0$, and $\partial \ln \bar{\xi}/\partial \eta_p < 0$ by Proposition 1. For $C_{j,t+1}^*$ the bracketed coefficient becomes $[\rho_t s_\Pi - \tilde{q}_t/(1 - \sigma)]$, strictly positive for any $\gamma > 0$ since $\tilde{q}_t > 0$. The Taylor expansion uses $\partial \ln \bar{\xi}/\partial \eta_p \approx -2a\delta(\sigma - 1)/\bar{\tau}$ from Proposition 1 (iii) and (A.10). $\square$

Proposition 4 (Mean-preserving spread, welfare under endogenous tariff revenue). *Under the mean-preserving-spread parameterization of Proposition 2 and the same assumptions as Proposition 3:*

- (i) *The welfare elasticities admit the closed forms (A.21)–(A.22) with $\eta = \eta_d$ and $\partial \ln \bar{\xi}/\partial \eta_d > 0$.*
- (ii) *The decomposition reads*

$$\frac{\partial \ln C_{jt}^*}{\partial \eta_d} = \underbrace{\chi(\nu\omega_{ijt} + 1) \left[\rho_t s_\Pi - \frac{q_t(1 - 1/\gamma)}{1 - \sigma} \right] \frac{\partial \ln \bar{\xi}_{ij,t+1}}{\partial \eta_d}}_{(A) > 0 \text{ for } \gamma \geq 1} + \underbrace{\rho_t s_R \chi \frac{\partial \ln \mathbb{E}_t[R_{ij,t+1}^g]}{\partial \eta_d}}_{(B) \text{ sign ambiguous}}. \quad (\text{A.25})$$

Term (B) decomposes via (A.14) into an entry channel (> 0 , inheriting the sign of $\partial \ln \bar{\xi}/\partial \eta_d$) and a reallocation channel (< 0 over the empirically relevant trade-cost range, since $\mathcal{R}$ is concave there for both $\sigma > 2$ and $1 < \sigma \leq 2$); the two work in opposite directions, and the overall sign of term (B) depends on the relative magnitudes.

(iii) *Under the symmetric parameterization with $d\lambda^H/d\eta_d = d\lambda^L/d\eta_d = a > 0$, $d\lambda^M/d\eta_d = -2a$, a second-order Taylor expansion in δ yields*

$$\begin{aligned} \frac{\partial \ln C_{jt}^*}{\partial \eta_d} \approx & \chi(\nu\omega_{ijt} + 1) \left[\rho_t s_\Pi - \frac{q_t(1 - 1/\gamma)}{1 - \sigma} \right] \frac{a\sigma(\sigma - 1)\delta^2}{\bar{\tau}^2} \\ & + \rho_t s_R \chi \left[\nu\omega_{ijt} \frac{a\sigma(\sigma - 1)\delta^2}{\bar{\tau}^2} + \frac{a(\sigma - 1)\delta^2}{(\bar{\tau} - 1)\bar{\tau}^2} [(\sigma - 2)\bar{\tau} - \sigma] \right]. \end{aligned} \quad (\text{A.26})$$

Both terms are $O(\delta^2)$. In the low-tariff limit $\bar{\tau} \rightarrow 1$, $s_R \rightarrow 0$ and the welfare response converges to the exogenous-revenue benchmark, which is strictly positive. As $\bar{\tau}$ rises, the reallocation channel becomes quantitatively relevant; the overall sign depends on the magnitude of s_R and on the local curvature of $\mathcal{R}(\tau)$ at $\bar{\tau}$.

Proof. Symmetric to the proof of Proposition 3, using $\partial \ln \bar{\xi}/\partial \eta_d \approx a\sigma(\sigma - 1)\delta^2/\bar{\tau}^2$ from Proposition 2 (iii) and (A.11). $\square$

Discussion: where the ambiguity comes from. Term (A) in both decompositions inherits the unambiguous sign of $\partial \ln \bar{\xi}/\partial \eta$: pessimistic shifts lower expected profits and raise the future price index, depressing consumption; mean-preserving spreads do the opposite. Its magnitude scales with $(\theta - \sigma + 1)$ for

η_p and $\sigma(\theta - \sigma + 1)$ for η_d , mirroring Propositions 1 and 2. The only new source of indeterminacy is term (B), the tariff-revenue channel. The entry component of term (B) moves with $\partial \ln \bar{\xi} / \partial \eta$, while the reallocation component—reflecting the mechanical effect of probability mass shifting along $\mathcal{R}(\tau)$ —moves in the opposite direction over the relevant trade-cost range. Whether the overall welfare response keeps the benchmark sign depends on the revenue share $s_R = (\bar{\tau} - 1)\chi / [(\sigma - 1)/\sigma + 1/\sigma + (\bar{\tau} - 1)\chi]$. In low-tariff environments ($\bar{\tau}$ close to one, $s_R \rightarrow 0$) and in sectors with high σ (so that the reallocation peak τ_σ^* is closer to one), the tariff-revenue channel is quantitatively small and the benchmark sign predictions are recovered. In high-tariff environments, the reallocation channel can be sizable and may dominate, generating the heterogeneity documented in the calibration of Section 4.

Corollary 1 (Welfare summary). *Combining Propositions 3 and 4, the welfare elasticity of current consumption to an uncertainty shock η that combines a pessimistic shift (intensity w_p) and a mean-preserving spread (intensity w_d) is*

$$\begin{aligned} \frac{\partial \ln C_{jt}^*}{\partial \eta} = & \chi (\nu \omega_{ijt} + 1) \left[\rho_t s_\Pi - \frac{q_t(1 - 1/\gamma)}{1 - \sigma} \right] \left[w_p \frac{\partial \ln \bar{\xi}}{\partial \eta_p} + w_d \frac{\partial \ln \bar{\xi}}{\partial \eta_d} \right] \\ & + \rho_t s_R \chi \left[w_p \frac{\partial \ln \mathbb{E}_t[R_{ij}^g]}{\partial \eta_p} + w_d \frac{\partial \ln \mathbb{E}_t[R_{ij}^g]}{\partial \eta_d} \right]. \end{aligned} \quad (\text{A.27})$$

When $s_R = 0$, the formula collapses to the exogenous-revenue benchmark; when $s_R > 0$, the welfare response acquires an additional term proportional to the endogenous tariff-revenue elasticity, which can either reinforce or offset the profit/price-index channels depending on the local curvature of $\mathcal{R}(\tau)$ at $\bar{\tau}$.

B Infinite-Horizon Derivation of ω

This appendix derives the demand-share weight ω in the infinite-horizon analogue of the two-period model developed in Section 2. The two-period closed-form gives $\omega \approx \beta/(1 + \beta) \approx 0.49$ at $\beta = 0.96$. The infinite-horizon counterpart yields $\omega^\infty = \beta = 0.96$ in the stationary limit, which is the value used in the calibrated application of Section 4.6.

B.1. Setup

Consider the infinite-horizon analogue of the firm's entry decision. A firm with productivity φ active in market j at time t faces the Bellman equation

$$V_{ij}(\varphi, \tau_t) = \pi_{ij}(\varphi, \tau_t) + \beta \mathbb{E}_t[V_{ij}(\varphi, \tau_{t+1})], \quad (\text{B.1})$$

where $\pi_{ij}(\varphi, \tau_t)$ is the per-period profit from exporting to j at productivity φ facing trade cost τ_t , and τ_t follows an ergodic stochastic process with stationary distribution $\bar{\Pi}(\tau)$. The per-period profit takes the closed form

$$\pi_{ij}(\varphi, \tau_t) = B_j \varphi^{\sigma-1} \tau_t^{1-\sigma}, \quad (\text{B.2})$$

where $B_j = \frac{1}{\sigma} \left(\frac{\sigma w_j}{\sigma-1} \right)^{1-\sigma} C_j P_j^{\sigma-1}$ is the country- j demand shifter.

B.2. Stationary value function

In the stationary equilibrium, $\mathbb{E}_t[\pi_{ij,t+s}] = B_j \varphi^{\sigma-1} \mathbb{E}_{\Pi}[\tau^{1-\sigma}] \equiv \pi_{ij}^{\text{ss}}(\varphi)$ for all $s \geq 1$, and the Bellman equation (B.1) can be iterated to give

$$V_{ij}(\varphi, \tau_t) = \pi_{ij}(\varphi, \tau_t) + \sum_{s=1}^{\infty} \beta^s \pi_{ij}^{\text{ss}}(\varphi). \quad (\text{B.3})$$

Summing the geometric series,

$$V_{ij}(\varphi, \tau_t) = \underbrace{\pi_{ij}(\varphi, \tau_t)}_{\text{current-period flow}} + \underbrace{\frac{\beta}{1-\beta} \pi_{ij}^{\text{ss}}(\varphi)}_{\text{discounted future flows}}. \quad (\text{B.4})$$

B.3. Definition of the demand-share weight

In the two-period model, ω is the share of the firm's expected value attributable to next-period (uncertain) flows. Generalising to the infinite-horizon setting, define

$$\omega^\infty(\varphi, \tau_t) \equiv \frac{\text{Discounted future flows}}{\text{Total firm value}} = \frac{\frac{\beta}{1-\beta} \pi_{ij}^{\text{ss}}(\varphi)}{\pi_{ij}(\varphi, \tau_t) + \frac{\beta}{1-\beta} \pi_{ij}^{\text{ss}}(\varphi)}. \quad (\text{B.5})$$

B.4. Stationary limit

In the symmetric stationary case where the current-period profit equals its stationary value, $\pi_{ij}(\varphi, \tau_t) = \pi_{ij}^{\text{ss}}(\varphi)$, the productivity term $\varphi^{\sigma-1}$ and the demand shifter B_j cancel between numerator and denominator in (B.5), yielding

$$\omega^\infty = \frac{\beta/(1-\beta)}{1 + \beta/(1-\beta)} = \frac{\beta/(1-\beta)}{1/(1-\beta)} = \beta. \quad (\text{B.6})$$

At the calibrated $\beta = 0.96$,

$$\boxed{\omega^\infty = \beta = 0.96.} \quad (\text{B.7})$$

B.5. Comparison with the two-period truncation

The two-period closed-form derivation in Section 2 implies

$$\omega^{2p} = \frac{\beta \mathbb{E}[\tau_{ij,t+1}^{1-\sigma}]}{\tau_{ij,t}^{1-\sigma} + \beta \mathbb{E}[\tau_{ij,t+1}^{1-\sigma}]} \approx \frac{\beta}{1 + \beta} \approx 0.49 \quad (\text{B.8})$$

when the stationary expected payoff equals the current payoff. The two-period truncation understates the firm-value share exposed to future trade-cost realisations by approximately a factor of $\omega^{2p}/\omega^\infty \approx 0.51$, because the symmetric two-period weighting effectively caps the firm’s exposure to a single future period. The infinite-horizon analogue (B.7) weights an infinite stream of future realisations and recovers $\omega^\infty = \beta$.

B.6. Implication for the Trump experiment calibration

The Trump experiment in Section 4.6 adopts $\omega = 0.96$, the infinite-horizon limit derived in (B.7). The choice reflects that real-world export decisions are exposed to many future periods of trade-cost realisations, not just one. Using $\omega = 0.49$ from (B.8) instead would attenuate the predicted trade-flow and price elasticities by approximately a factor of 0.51; the qualitative pattern of the predictions would remain unchanged, but the quantitative match with realised 2025 data in Table 9 would deteriorate proportionally

A complete dynamic Pareto-Melitz model with explicit firm entry, exit, and uncertainty dynamics is the subject of ongoing work and is beyond the scope of the present paper. The steady-state derivation (B.1)–(B.7) suffices to justify the $\omega = 0.96$ calibration choice without invoking the full machinery of the dynamic model.

References

- [1] Ahir, H., Bloom, N., & Furceri, D. (2022). The world uncertainty index (No. w29763). National Bureau of Economic Research.
- [2] Al-Thaqeb, S. A., Algharabali, B. G., & Alabdulghafour, K. T. (2022). The pandemic and economic policy uncertainty. *International Journal of Finance & Economics*, 27(3), 2784–2794.
- [3] Altig, D., Baker, S., Barrero, J. M., Bloom, N., Bunn, P., Chen, S., ... & Thwaites, G. (2020). Economic uncertainty before and during the COVID-19 pandemic. *Journal of Public Economics*, 191, 104274.
- [4] Arbatli, E., Davis, S. J., Ito, A., & Miake, N. (2019). Policy uncertainty in Japan. NBER Working Paper No. 23411.
- [5] Arkolakis, C., Costinot, A., & Rodríguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, 102(1), 94–130.
- [6] Armelius, H., Hull, I., & Köhler, H. S. (2017). The timing of uncertainty shocks in a small open economy. *Economics Letters*, 155, 31–34.
- [7] Armington, P. S. (1969). A theory of demand for products distinguished by place of production. *IMF Staff Papers*, 16(1), 159–178.

- [8] Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4), 1593–1636.
- [9] Baley, I., Veldkamp, L., & Waugh, M. (2020). Can global uncertainty promote international trade? *Journal of International Economics*, 126, 103347.
- [10] Benguria, F., Choi, J., Swenson, D. L., & Xu, M. J. (2022). Anxiety or pain? The impact of tariffs and uncertainty on Chinese firms in the trade war. *Journal of International Economics*, 137, 103608.
- [11] Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3), 623–685.
- [12] Boehm, C., Levchenko, A. A., Pandalai-Nayar, N., & Toma, H. (2024). Dynamic models, new gains from trade? NBER Working Paper No. 32565.
- [13] Broda, C., & Weinstein, D. E. (2006). Globalization and the gains from variety. *Quarterly Journal of Economics*, 121(2), 541–585.
- [14] Bureau of Labor Statistics. (2025). The employment situation. Monthly releases, January–September 2025. U.S. Department of Labor.
- [15] Caldara, D., Iacoviello, M., Mollico, P., Prestipino, A., & Raffo, A. (2020). The economic effects of trade policy uncertainty. *Journal of Monetary Economics*, 109, 38–59.
- [16] Cavallo, A., Gopinath, G., Neiman, B., & Tang, J. (2021). Tariff pass-through at the border and at the store: Evidence from US trade policy. *American Economic Review: Insights*, 3(1), 19–34.
- [17] Cerda, R., Silva, A., & Valente, J. T. (2016). Economic policy uncertainty indices for Chile. Working paper.
- [18] Chaney, T. (2008). Distorted gravity: the intensive and extensive margins of international trade. *American Economic Review*, 98(4), 1707–1721.
- [19] Chen, J., Góes, C., Muendler, M. A., & Trottner, F. (2024). Dynamic adjustment to trade shocks. UCSD Mimeo.
- [20] Committee for a Responsible Federal Budget. (2025). Tariff revenue soars in FY 2025 amid legal uncertainty. CRFB analysis, October 2025.
- [21] Correia, S., Guimarães, P., & Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal*, 20(1), 95–115.
- [22] Davis, S. J. (2016). An index of global economic policy uncertainty (No. w22740). National Bureau of Economic Research.
- [23] Eaton, J., & Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5), 1741–1779.

- [24] Fajgelbaum, P. D., & Khandelwal, A. K. (2022). The economic impacts of the US–China trade war. *Annual Review of Economics*, 14(1), 205–228.
- [25] Fontagné, L., Guimbard, H., & Orefice, G. (2022). Tariff-based product-level trade elasticities. *Journal of International Economics*, 137, 103593.
- [26] Ghirelli, C., Perez, J. J., & Urtasun, A. (2019). A new economic policy uncertainty index for Spain. Working paper.
- [27] Gil, M., & Silva, D. (2018). Economic policy uncertainty indices for Colombia. Deutsche Bank Research, 1–9.
- [28] Góes, C., & Bekkers, E. (2022). The impact of geopolitical conflicts on trade, growth, and innovation. *arXiv preprint*.
- [29] Graziano, A., Handley, K., & Limão, N. (2018). Brexit uncertainty and trade disintegration. NBER Working Paper No. 25334.
- [30] Handley, K. (2014). Exporting under trade policy uncertainty: Theory and evidence. *Journal of International Economics*, 94(1), 50–66.
- [31] Handley, K., & Limão, N. (2015). Trade and investment under policy uncertainty. *AEJ: Economic Policy*, 7(4), 189–222.
- [32] Handley, K., & Limão, N. (2017). Policy uncertainty, trade, and welfare. *American Economic Review*, 107(9), 2731–2783.
- [33] Hardouvelis, G. A., Karalas, G., Karanastasis, D., & Samartzis, P. (2018). Economic policy uncertainty, political uncertainty and the Greek economic crisis. Working paper, April 3.
- [34] Hassan, T. A., Hollander, S., van Lent, L., & Tahoun, A. (2024). The global impact of Brexit uncertainty. *Journal of Finance*, 79(1), 413–458.
- [35] Head, K., & Mayer, T. (2014). Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of International Economics* (Vol. 4, pp. 131–195). Elsevier.
- [36] Itakura, K. (2020). Evaluating the impact of the US–China trade war. *Asian Economic Policy Review*, 15(1), 77–93.
- [37] Kroese, L., Kok, S., & Parlevliet, J. (2015). Beleidsonzekerheid in Nederland. *Economisch Statistische Berichten*, 4715, 464–467.
- [38] Krugman, P. (1980). Scale economies, product differentiation, and the pattern of trade. *American Economic Review*, 70(5), 950–959.
- [39] Liu, E., & Yang, D. Y. (2025). International power. NBER Working Paper No. 34006.
- [40] Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2), 145–161.

- [41] Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate productivity. *Econometrica*, 71(6), 1695–1725.
- [42] Novy, D., & Taylor, A. M. (2020). Trade and uncertainty. *Review of Economics and Statistics*, 102(4), 749–765.
- [43] Santos Silva, J. M. C., & Tenreyro, S. (2006). The log of gravity. *Review of Economics and Statistics*, 88(4), 641–658.
- [44] Santos Silva, J. M. C., & Tenreyro, S. (2011). Further simulation evidence on the performance of the Poisson pseudo-maximum likelihood estimator. *Economics Letters*, 112(2), 220–222.
- [45] Simonovska, I., & Waugh, M. E. (2014). The elasticity of trade: Estimates and evidence. *Journal of International Economics*, 92(1), 34–50.
- [46] Steinbach, S. (2023). The Russia–Ukraine war and global trade reallocations. *Economics Letters*, 226, 111075.
- [47] Steinberg, J. B. (2019). Brexit and trade policy uncertainty. *Journal of International Economics*, 117, 175–195.
- [48] U.S. Bureau of Economic Analysis. (2025). “National Income and Product Accounts.” NIPA Tables 1.1.6, 4.2.5; quarterly data through 2025Q3.
- [49] U.S. Customs and Border Protection. (2025). “CBP Announces Record-Breaking \$200 Billion in Tariff Revenue.” National Media Release, December 2025.
- [50] Weidner, M., & Zylkin, T. (2021). Bias and consistency in three-way gravity models. *Journal of International Economics*, 132, 103513.
- [51] Zalla, R. (2017). Economic policy uncertainty in Ireland. *Atlantic Economic Journal*, 45(2), 269–271.




Figure 2: Calibrated tariff-revenue Laffer curve. Panel (a): cross-sector mean of $\chi \cdot \mathbb{E}[(\tau - 1)\tau^{1-\sigma}]$ as a function of baseline tariff $\bar{\tau}$, with peak marked. Panel (b): marginal tariff revenue $\partial R^g / \partial \bar{\tau}$; zero crossing identifies the peak. Panel (c): Laffer curves for five representative values of $\sigma \in \{2.5, 3.0, 3.23, 4.0, 5.0\}$, illustrating how the peak $\tau_\sigma^* = (\sigma - 1)/(\sigma - 2)$ shifts left as σ rises. Calibration: $\theta = 5$, $\delta = 0.05$, $\chi = 0.30$, with σ from Broda and Weinstein (2006).




Figure 3: Joint-shock decomposition at the 2025 U.S. baseline. Panel (a): log-change for five macro aggregates, decomposed into mechanical, uncertainty, and interaction channels; black diamonds mark the joint total. Panel (b): consumption decomposition into Π , R^g , and price contributions under each channel. Calibration uses separated pass-through ($\rho_{\text{border}} = 1.0$ for firm-side decisions, $\rho_{\text{retail}} = 0.30$ for the consumer-side PCE inflation contribution).