E Pluribus Euro:
Minimum Fiscal Capacity for Collective Trade Policy in a Currency Union

Joseph Kopecky*

March 2026

Abstract

As major powers deploy trade policy as coercion, what fiscal capacity does a currency union need to sustain a credible collective response? I embed the multi-sector trade model of Caliendo and Parro (2015) into a monetary union with heterogeneous members, calibrated to the world input-output database (WIOD) for 20 individual Eurozone members. A US tariff escalation of 20% plus EU retaliation requires 0.69% of Eurozone GDP (€97 billion); a Chinese critical minerals restriction requires 0.44% (€62 billion); both simultaneously require 1.12% (€157 billion). A substantial share of the fiscal need arises from the asymmetric costs of collective action itself: the costs that EU counter-tariffs impose on members with concentrated trade exposures. This reframes the fiscal requirement as the price of strategic credibility. Single market deepening generates welfare gains, but barely reduces the fiscal requirement, showing that integration and fiscal capacity are complements. Joint borrowing is needed, as budget-balanced redistribution cannot sustain collective action. However, the headline fiscal requirement is an upper bound. Embedding the model in existing EU institutions (cross-conditionality of EU fiscal flows and qualified majority voting rules) reduces the practical requirement to 0.33% of Eurozone GDP (€46 billion) in the combined scenario, since the EU need only compensate a handful of pivotal large members to prevent a blocking minority.

JEL classification codes: F13, F15, F42, F45, E62, H77

Keywords: Fiscal unions, currency unions, trade policy, economic coercion, Eurozone, strategic autonomy

*Department of Economics, Trinity College Dublin (jkopecky@tcd.ie)

1. Introduction

The Eurozone is one of the world's largest experiments in deep economic integration. Its members share a currency, a common external tariff, and a single monetary policy. However, fiscal integration remains radically incomplete. There is no meaningful cross-country unemployment insurance, limited capacity for counter-cyclical transfers, and joint debt issuance remains politically contentious despite the precedent set by NextGenerationEU. This lack of integration may weaken the ability of the Eurozone to act collectively in a trade environment where global hegemons increasingly act coercively.

The mechanism is straightforward. External trade shocks hit Eurozone members asymmetrically. Germany absorbs automotive and machinery shocks through its export exposure. The Netherlands faces disruption through its intermediate input dependence on China. Ireland is exposed through its multinational pharmaceuticals and technology sector. Without fiscal transfers to pool these asymmetric costs, individual member states face political pressure to seek bilateral accommodation with the coercing power, undermining the EU's collective bargaining position. The common trade policy becomes politically unsustainable precisely when it is most needed.

I formalize this mechanism by defining a *minimum fiscal capacity* ($\bar{F}$): the smallest GDP-weighted transfer budget such that no Eurozone member prefers its bilateral outside option to the collective trade policy. In a repeated game in the spirit of Bagwell and Staiger (1999), cooperation can in principle be sustained without transfers if players are patient enough. I show that under the baseline calibration, the one-shot deviation gain is so large relative to the punishment cost (EU exclusion) that cooperation requires fiscal transfers even at high discount factors, suggesting that the folk theorem alone is insufficient. $\bar{F}$ is then the minimum transfer that relaxes the enforcement constraint.

I make four contributions. First, I develop a simplified two-member analytical model that characterises $\bar{F}$ through four propositions: a defection threshold that depends on sectoral exposure, a complementarity between single market depth and fiscal capacity, an amplification under nominal rigidity, and a repeated-game characterisation that determines when fiscal transfers are needed even with patient players. These results establish that $\bar{F}$ is not merely an accounting object but the enforcement cost of collective trade policy.

Second, I build and calibrate a multi-sector Caliendo-Parro trade model with input-output linkages for 20 individual Eurozone members (plus the US, China, and rest of world) across 10 sectors. Applied to realistic shock scenarios, a US tariff escalation of 20% with symmetric EU retaliation requires $\bar{F} = 0.69\%$ of Eurozone GDP (€97 billion); a Chinese supply restriction requires $\bar{F} = 0.44\%$ (€62 billion); both simultaneously require

$\bar{F} = 1.12\%$ (€157 billion). These numbers are robust to trade elasticity assumptions ($\pm 50\%$), to excluding outlier countries[1] (-14%), and to varying the bilateral deal quality and defection penalty. Without EU retaliation, $\bar{F}$ falls substantially, but remains positive. This is because the quantitative model finds retaliation to be the primary source of fiscal need. Indeed, retaliation makes the EU statically worse off: aggregate welfare falls more under cooperation than under fragmentation. The minimum fiscal capacity therefore includes the cost of the EU's own retaliatory response, reframing $\bar{F}$ as the price of strategic credibility.

Third, I establish two structural results. Single market deepening generates large welfare gains, but barely changes $\bar{F}$. This shows that integration and fiscal capacity are complements. Budget-balanced redistribution as in Farhi and Werning (2017) cannot sustain collective action when all members prefer bilateral deals; only joint borrowing achieves incentive compatibility. A New Keynesian monetary union layer shows that nominal rigidities leave the aggregate $\bar{F}$ unchanged but concentrate fiscal need on small, exposed members. This element of the model suggests that facility *design* is critical in addition to the *size* of fiscal need.

Fourth, I ground the analysis in existing EU institutions. The full compensation estimate ($\bar{F} = 0.69\%$) is an upper bound that assumes every member must be made whole.[2] In practice, existing EU fiscal flows (structural funds, Common Agriculture Policy, etc) already provide cross-conditionality for net recipient members, and the EU's qualified majority voting rules mean that collective policy requires preventing a *blocking minority*, not achieving unanimity. Accounting for both mechanisms reduces the practical fiscal requirement to 0.18% of EZ GDP (€26 billion) for the US tariff scenario, still sizable but roughly a quarter of the headline figure.

A natural objection is that the Pareto-improving response to external tariffs is simply not to retaliate. This is correct in a static, one-shot game, and my cost-benefit analysis of Section 6.6 confirms it. But the static comparison misses two features that dominate EU policymaking. First, retaliatory capacity deters coercion: without it, the hegemon faces no cost from tariff escalation. To illustrate the potential scale, I conduct a bloc bargaining exercise in Section 6.1, where EU retaliation imposes a cost on the US equal to 104% of the tariff's own welfare cost—leverage that no individual member possesses, though whether it translates into negotiated outcomes depends on the hegemon's domestic political economy. Second, the EU has revealed a strong institutional preference for collective trade policy.[3]

[1] Ireland and Luxembourg are small economies with large individual costs that are likely skewed by the presence of multinationals.

[2] Effectively it can be thought of as the cost of ensuring that any pains arising as a result of joint policy are shared evenly across the union.

[3] For example, the Anti-Coercion Instrument (Regulation 2023/2675), expanded trade defence mechanisms, and the common commercial policy under Article 207 TFEU.

This paper does not seek to answer the question of what retaliatory policy is optimal. Instead, I ask what fiscal infrastructure is needed to make such a policy sustainable.

I connect three literatures. The first is the fiscal federalism literature (Farhi and Werning 2017; Evers 2015), which characterises optimal transfers in monetary unions but treats shocks as generic disturbances. The quantitative trade policy literature (Caliendo and Parro 2015; Ossa 2014) evaluates tariff counterfactuals with sectoral granularity, but generally uses flexible-price models with no monetary union structure. Finally I connect these to a geoeconomics literature (Clayton, Maggiori, and Schreger 2024; Farrell and Newman 2019), which formalises coercion through trade networks, but does not model the internal institutional requirements for a bloc to act collectively. I bridge these by embedding trade shocks in a currency union with incentive compatibility constraints, yielding a quantitative answer to a question that each literature alone cannot address.

The remainder of the paper is organised as follows. Section 2 reviews the related literature. Section 3 presents the analytical model. Section 4 sets out the quantitative framework. Section 5 describes the data and calibration. Section 6 presents the quantitative exercises. Section 7 grounds the results in existing EU institutions, computing the practical fiscal requirement under cross-conditionality and qualified majority voting rules. Section 8 concludes.

2. Related Literature

2.1. Fiscal Integration in Monetary Unions

The theoretical case for fiscal integration in currency unions originates with Kenen (1969), who argued that fiscal transfers are necessary to offset asymmetric shocks when exchange rate adjustment is unavailable. Farhi and Werning (2017) formalise this in a New Keynesian open-economy model, showing that even with complete financial markets, privately optimal risk sharing is constrained inefficient in a currency union. Fiscal transfers address the labour wedge that arises when a common monetary policy cannot stabilise heterogeneous members simultaneously. They show that the benefits of a fiscal union are larger when asymmetric shocks are more severe, more persistent, and when member economies exhibit greater home bias.

Asdrubali, Sørensen, and Yosha (1996) decompose interstate risk sharing in the US, finding that capital markets absorb 39% of asymmetric shocks, credit markets 23%, and fiscal transfers only 13%. The Eurozone's private risk-sharing channels, while not absent, are substantially weaker and prone to fragmentation during crises, placing a larger burden on fiscal integration. Evers (2015) takes a quantitative approach using a medium-scale

DSGE model with imperfect risk-sharing and physical capital, finding that stabilising labour income delivers higher welfare than stabilising consumption directly. Galí and Monacelli (2008) analyse optimal fiscal policy for stabilising a currency union but do not consider cross-country transfers.

The key limitation of this literature for my purposes is that shocks are treated as generic productivity or demand disturbances. The source and structure of the shock, whether it is a tariff, a supply chain disruption, or financial coercion, is absent. This matters because the policy response to a trade shock involves not just fiscal stabilisation but also a collective trade policy decision, and the fiscal capacity to sustain that collective response is what I study.

2.2. Quantitative Trade Policy

The trade backbone of the model draws on Caliendo and Parro (2015), who extend the Eaton and Kortum (2002) Ricardian framework to multiple sectors with intermediate goods and input-output linkages. The "exact hat algebra" technique of Dekle, Eaton, and Kortum (2007) permits counterfactual analysis using only observables (trade shares, tariffs, I-O coefficients) and trade elasticities. Arkolakis, Costinot, and Rodríguez-Clare (2012) show that welfare gains from trade across a wide class of models depend on two sufficient statistics: the domestic expenditure share and the trade elasticity. Costinot and Rodríguez-Clare (2014) survey the quantitative trade literature and establish best practices for counterfactual analysis. Ossa (2014) nests terms-of-trade, profit-shifting, and political economy motives in a unified framework, computing Nash and cooperative tariffs for major blocs. Adao, Costinot, and Donaldson (2025) develop an IV-based test for the causal predictions of quantitative trade models, applied to the 2018 US-China tariff war.

The critical gap is that these models use flexible-price, representative-agent economies. When Ossa (2014) computes the EU's optimal tariff, he treats it as a single decision-maker—the internal political economy of sustaining that tariff across heterogeneous members is absent.

The trade agreements literature provides the theoretical underpinning for why countries negotiate common external policies. Bagwell and Staiger (1999) show that trade agreements solve a terms-of-trade externality; members commit to lower tariffs because unilateral defection imposes negative externalities on partners. Bown (2004) documents that the threat of WTO-authorised retaliation effectively enforces trade agreements, suggesting that enforcement mechanisms matter empirically. My framework introduces a new dimension: within a customs union, the "defection" is not a tariff war but a bilateral accommodation that undermines the bloc's collective bargaining position. The minimum fiscal transfer $\bar{F}$ is the price of sustaining this collective commitment. Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2020) and Amiti, Redding, and Weinstein (2019) provide empirical benchmarks for tariff welfare effects against which I validate the model. Recent work by Bonadio, Huo, Levchenko, and Pandalai-Nayar (2024) studies supply chain disruptions and Gopinath, Gourinchas, Presbitero, and Topalova (2025) documents declining inter-bloc trade, but neither addresses the intra-bloc fiscal arrangements needed to sustain collective action.

2.3. Economic Coercion and Geoeconomics

Farrell and Newman (2019) introduced the concept of "weaponised interdependence," identifying how asymmetric global networks create coercive leverage. Clayton et al. (2024) formalise this in a model where a hegemon threatens suspension of trade and financial links. Their key result is that *uncoordinated* defensive policies lead to a "fragmentation doom loop"—excessive fragmentation destroying gains from integration.

Becko and O'Connor (2025) and Becko, Grossman, and Helpman (2025) model hegemon-small country interactions with trade as carrot and stick. Liu and Yang (2025) construct measures of international power based on asymmetric import dependence. Broner, Martin, Meyer, Trebesch, and Zhou (2025) develop a theory of hegemonic globalisation showing that unipolarity promotes alignment while multipolarity fragments. At a broader level, Rodrik (2011) frames the core tension as a trilemma: deep economic integration, national sovereignty, and democratic politics cannot be maintained simultaneously. The EU's struggle to reconcile heterogeneous trade interests within a common policy is a concrete instance of this trilemma. None of these papers model the *internal* institutional requirements for a bloc to act collectively against external coercion.

3. Analytical Model

Before the full quantitative exercise, I derive analytical results in a simplified setting that illustrates the core mechanism.

3.1. Setup

Consider a trade union consisting of two members A and B and one external hegemon H. There are two sectors: a strategic sector s and a residual sector r. Member A has high trade exposure to the hegemon in the strategic sector (π^s_{AH} is large), while B does not (π^s_{BH} is small). The members share a common currency and face Calvo price stickiness with parameter ξ.

The hegemon imposes a tariff $\bar{\tau}$ on both union members. The union's collective response is a common retaliatory tariff at the same rate. Each member's *outside option* is a bilateral deal with the hegemon at a lower tariff $\tau_{\text{bilat}} < \bar{\tau}$,[4] obtained by breaking from the collective stance. Defection from the collective policy incurs a cost: partial loss of single market access (a fraction δ_{pen} of intra-union trade gains).

In this setting, the welfare of member A under the common policy is:

$$W_A^{\text{EU}} = 1 - \underbrace{\left(\pi_{AH}^s \cdot \frac{(\bar{\tau}-1)^2}{2\theta^s} + \pi_{AH}^r \cdot \frac{(\bar{\tau}-1)^2}{2\theta^r}\right)}_{\text{tariff cost}} \cdot \underbrace{\left(1 + \frac{\xi}{1-\xi}\right)}_{\text{NK amplification}} + \underbrace{G_{AB}(d_{AB})}_{\text{SM gains}} \tag{1}$$

where θ^j is the trade elasticity in sector j, d_{AB} is the intra-union trade cost, and $G_{AB}(d_{AB})$ captures gains from single market access that depend on bilateral trade shares and iceberg costs.

Under the bilateral outside option:

$$W_A^{\text{bilat}} = 1 - \left(\pi_{AH}^s \cdot \frac{(\tau_{\text{bilat}}-1)^2}{2\theta^s} + \pi_{AH}^r \cdot \frac{(\tau_{\text{bilat}}-1)^2}{2\theta^r}\right) \cdot \left(1 + \frac{\xi/2}{1-\xi}\right) + (1-\delta_{\text{pen}})G_{AB}(d_{AB}) \tag{2}$$

The bilateral option reduces the tariff cost (from $\bar{\tau}$ to τ_{bilat}) but also reduces single market gains (by the defection penalty δ_{pen}) and partially dampens the NK amplification (the defecting member faces less monetary mismatch as it is partially accommodated).

3.2. Main Results

Proposition 1 (Defection Threshold). *Member A prefers a bilateral deal with H over the common retaliatory tariff whenever its strategic-sector exposure exceeds a threshold:*

$$\pi_{AH}^s > \bar{\pi}(\theta^s, \bar{\tau}, \tau_{\text{bilat}}, d_{AB}, \xi) \tag{3}$$

where $\bar{\pi}$ is the ratio of single market defection costs to the tariff differential gains, scaled by the nominal rigidity amplification. The minimum transfer to prevent defection is:

$$\bar{F}_A = \max\{W_A^{\text{bilat}} - W_A^{\text{EU}},\ 0\} \tag{4}$$

Proof sketch (full derivation in Appendix A). Member A defects if and only if $W_A^{\text{bilat}} > W_A^{\text{EU}}$.

[4] The parameter τ_{bilat} is a reduced-form summary of the bilateral tariff environment. In the quantitative model this decomposes into a hegemon offer and the member's own decision to adjust from collective retaliation; see Section 6.1. This makes the conclusions from this stylized model a lower bound on incentives to defect.

The defection gap $\Delta_A = W_A^{\text{bilat}} - W_A^{\text{EU}}$ is positive when the tariff savings from the bilateral deal exceed the single market losses from defection. From (1)–(2), the tariff savings scale with π_{AH}^s while the defection penalty is independent of exposure. Solving $\Delta_A = 0$ for π_{AH}^s gives the threshold $\bar{\pi}$. □

Proposition 2 (Single Market and Fiscal Capacity are Complements)**.** *The minimum fiscal transfer $\bar{F}_A$ is* ***nearly invariant*** *to the depth of the single market. Formally:*

$$\frac{\partial \bar{F}_A}{\partial d_{AB}} \approx 0 \tag{5}$$

A deeper single market (lower d_{AB}) raises welfare under both the collective policy and the bilateral outside option, because the defecting member retains most single market access. The defection gap, and hence the fiscal transfer needed to prevent it, changes only through the defection penalty channel, which is small.

Proof sketch (full derivation in Appendix A). From (1)–(2):

$$\Delta_A = \underbrace{(\text{tariff savings})}_{\text{independent of } d_{AB}} - \underbrace{\delta_{\text{pen}} \cdot G_{AB}(d_{AB})}_{\text{small if } \delta_{\text{pen}} \text{ small}} + \text{NK terms}$$

The first term depends on $\bar{\tau}, \tau_{\text{bilat}}, \pi_{AH}^s, \theta$ but not on d_{AB}. The defection penalty δ_{pen} operates only through the intra-union trade gains, and G_{AB} enters with coefficient δ_{pen}. For realistic defection penalties ($\delta_{\text{pen}} \approx 0.3$), the sensitivity of Δ_A to d_{AB} is small. □

Proposition 3 (Nominal Rigidity Amplification)**.** *In a currency union with sticky prices ($\xi > 0$), the minimum fiscal capacity $\bar{F}$ is strictly higher than under flexible prices ($\xi = 0$):*

$$\bar{F}(\xi) \geq \bar{F}(0) \quad \forall \xi \in [0, 1) \tag{6}$$

with equality only when members are identical.

Corollary 1 (NK Amplification is Mean-Preserving in a Multi-Member Union)**.** *In a multi-member union where the central bank targets the GDP-weighted average, nominal rigidity amplifies* ***individual*** *defection gaps but leaves the* ***GDP-weighted aggregate*** *$\bar{F}$ unchanged:*

$$\bar{F}^{NK} = \sum_{n \in \mathcal{E}} \omega_n \Delta_n^{NK} = \sum_{n \in \mathcal{E}} \omega_n \Delta_n^{\text{flex}} = \bar{F}^{\text{flex}} \tag{7}$$

where $\omega_n = \text{GDP}_n / \text{GDP}_{\mathcal{E}}$ are the ECB's GDP weights. NK frictions change who needs transfers, not how much the total facility costs.

Proof sketch (full derivation in Appendix A). The NK amplification decomposes each member's welfare deviation into an aggregate component (common to all, absorbed by ECB policy) and an idiosyncratic component (amplified by $1/(1-\xi)$). The idiosyncratic components sum to zero under GDP weights by construction, since they are deviations from the GDP-weighted mean. The GDP-weighted sum of amplified idiosyncratic components therefore also sums to zero. □

3.3. Repeated Game and the Enforcement Constraint

The preceding results characterise $\bar{F}$ as the one-shot cost of preventing defection. A natural objection is that the EU is an infinitely repeated game: can the folk theorem sustain cooperation without transfers?

I now embed the defection decision in a repeated game. In each period, member A chooses whether to cooperate (maintain the common tariff) or defect (pursue a bilateral deal). The punishment for defection is grim-trigger exclusion from the EU: A loses all single market preferences permanently. Let $V_A = W_A^{\text{EU}} - W_A^{\text{excluded}}$ be the per-period cooperation surplus, where W_A^{excluded} is welfare when A faces external trade costs with all union partners.

Proposition 4 (Repeated Game Enforcement)**.** *In the infinitely repeated game with discount factor β and grim-trigger punishment, member A cooperates without transfers if and only if $\beta \geq \beta_A^*$, where:*

$$\beta_A^* = \frac{\Delta_A}{\Delta_A + V_A} \tag{8}$$

Where β_A^* is the minimum patience required to sustain cooperation without transfers. When $\beta < \beta_A^*$, the minimum per-period fiscal transfer in the repeated game is:

$$\bar{F}_A(\beta) = \max\{(1-\beta)\Delta_A - \beta V_A,\ 0\} \tag{9}$$

which interpolates between the one-shot $\bar{F}$ at $\beta = 0$ and zero at $\beta = \beta_A^*$.

Proof sketch (full derivation in Appendix A). Member A cooperates if the one-shot deviation gain Δ_A is weakly less than the discounted loss from punishment: $\Delta_A - F_A \leq \frac{\beta}{1-\beta}(W_A^{\text{EU}} + F_A - W_A^{\text{excluded}})$. With $F_A = 0$, solving for β gives (8). With transfers, rearranging gives the minimum F_A such that the constraint holds, yielding (9). □

The critical discount factor β_A^* depends on the ratio of the defection gain to the cooperation surplus. When external shocks are large (high Δ_A) relative to the value of EU membership (V_A), β_A^* approaches 1. Even very patient players require fiscal transfers. This is the key insight: the folk theorem does not eliminate the need for fiscal capacity when the shocks are large enough.

3.4. Analytical Calibration

Table 1 reports results under the baseline calibration ($\pi^s_{AH} = 0.15$, $\theta^s = 4$, $d_{AB} = 1.20$, $\xi = 0.75$, $\bar{\tau} = 1.20$, $\tau_{\text{bilat}} = 1.05$).

Table 1: *Analytical Model Results*

Result	Value
Defection threshold $\bar{\pi}$ (Prop. 1)	14.39
A's actual exposure π^s_{AH}	0.15
Minimum transfer $\bar{F}_A$ (% of baseline welfare, Prop. 1)	0.64%
Single market depth (Prop. 2)	
$d_{AB} = 1.05$	$\bar{F} = 0.48\%$
$d_{AB} = 1.40$	$\bar{F} = 0.80\%$
Nominal rigidity (Prop. 3)	
$\xi = 0.00$ (flexible)	$\bar{F} = 0.33\%$
$\xi = 0.85$ (very sticky)	$\bar{F} = 0.92\%$

Notes: $\bar{F}$ as percentage of baseline welfare. Comparative statics: retaliation tariff $\bar{\tau}$ has the largest marginal effect on $\bar{F}$ (+0.68pp), followed by nominal rigidity ξ (+0.28pp).

4. Quantitative Model

The quantitative framework has two layers. The first is a workhorse multi-sector trade model, which delivers the welfare effects and defection gaps that determine $\bar{F}$. The second, a New Keynesian monetary union layer, does not change the aggregate $\bar{F}$ (Corollary 1) but reveals how nominal rigidities redistribute fiscal need across members, a key input for facility design. Figure 1 provides a schematic for how each block connects.

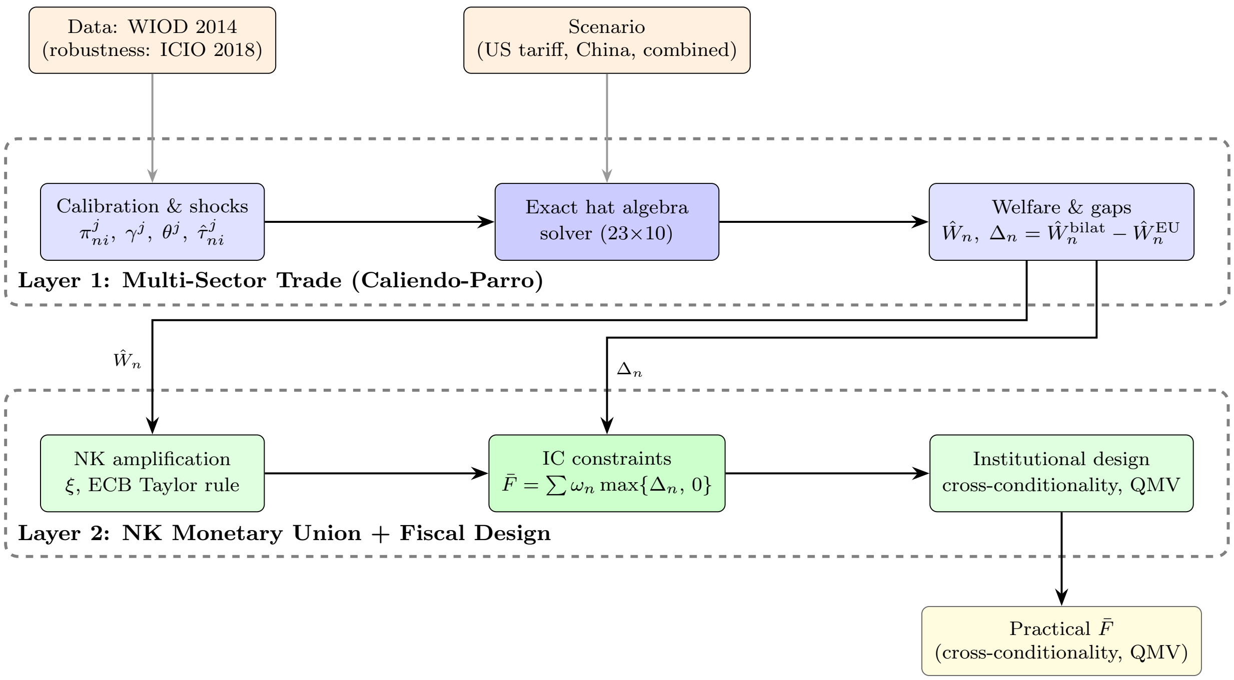

Figure 1: *Model Architecture*

Notes: Layer 1 computes welfare changes $\hat{W}_n$ and defection gaps Δ_n via the Caliendo–Parro trade model. Layer 2 applies NK amplification to obtain headline $\bar{F}$ (See: Table 7), then institutional filters to obtain practical $\bar{F}$ (See: Table 11).

4.1. Layer 1: Multi-Sector Trade with Eurozone Structure

4.1.1 Countries, Sectors, and Trade Structure

There are N countries partitioned into Eurozone members $\mathcal{E} = \{1, \ldots, N_E\}$ and external economies $\mathcal{G} = \{N_E + 1, \ldots, N\}$. In the primary specification, $N_E = 20$ (individual EZ members), with $\mathcal{G} = \{\text{US}, \text{China}, \text{RoW}\}$, giving $N = 23$. There are $J = 10$ sectors, chosen to isolate heterogeneous strategic dependencies across members.

Eurozone members share a *common external tariff* $\bar{\tau}^j$ on imports from outside the EU, while internal tariffs are zero. Intra-EU trade costs $d_{ni}^j > 1$ for $n, i \in \mathcal{E}$ persist due to regulatory divergence, language barriers, and non-tariff barriers.

4.1.2 Production and Trade Shares

Production in each country-sector follows the Caliendo-Parro structure. The unit cost in country n, sector j is:

$$c_n^j = \frac{w_n^{\gamma_n^j}}{z_n^j} \prod_{k=1}^{J} (P_n^k)^{\gamma_n^{kj}} \tag{10}$$

where w_n is the wage, γ_n^j is the value-added share, γ_n^{kj} is the input-output coefficient (share of sector k inputs in sector j production), and z_n^j is Fréchet-distributed productivity with scale T_n^j and shape θ^j.

Bilateral trade shares are:

$$\pi_{ni}^j = \frac{T_i^j(\kappa_{ni}^j c_i^j)^{-\theta^j}}{\sum_h T_h^j(\kappa_{nh}^j c_h^j)^{-\theta^j}} \tag{11}$$

where $\kappa_{ni}^j = d_{ni}^j \cdot \tau_{ni}^j$ combines iceberg trade costs and tariff barriers.

4.1.3 Exact Hat Algebra

Following Dekle et al. (2007) and Caliendo and Parro (2015), I solve for counterfactual equilibria using exact hat algebra. For any change in tariffs $\hat{\tau}_{ni}^j$, the counterfactual equilibrium satisfies:

$$\hat{\pi}_{ni}^j = \frac{(\hat{\kappa}_{ni}^j \hat{c}_i^j)^{-\theta^j}}{\sum_h \pi_{nh}^j(\hat{\kappa}_{nh}^j \hat{c}_h^j)^{-\theta^j}} \tag{12}$$

$$\hat{c}_n^j = \hat{w}_n^{\gamma_n^j} \prod_{k=1}^{J} (\hat{P}_n^k)^{\gamma_n^{kj}}, \qquad \hat{P}_n^j = \left[\sum_i \pi_{ni}^j(\hat{\kappa}_{ni}^j \hat{c}_i^j)^{-\theta^j}\right]^{-1/\theta^j} \tag{13}$$

where hats denote ratios of counterfactual to baseline values. The system is solved by Gauss-Seidel iteration on wages $\hat{w}_n$ until convergence (tolerance 10^{-10}), with an inner loop for unit costs and price indices via the I-O fixed point. Tariff revenue is collected and redistributed lump-sum to domestic households. Let α_{nj} denote country n's expenditure share on sector j. The tariff revenue rate is $r_n = \sum_j \alpha_{nj} \sum_{i \neq n} \frac{\tau_{ni}^j - 1}{\tau_{ni}^j} \pi_{ni}^j$, so total income is $I_n = w_n L_n / (1 - r_n)$ and welfare changes are:

$$\hat{W}_n = \hat{I}_n / \hat{P}_n, \qquad \hat{I}_n = \hat{w}_n \cdot \frac{1 - r_n^{\text{base}}}{1 - r_n'} \tag{14}$$

where r_n^{base} and r_n' are baseline and counterfactual revenue rates, and $\hat{P}_n$ is the aggregate consumer price change.

4.1.4 The Political Economy of Collective Trade Policy

The EU sets the common external tariff to maximise a weighted sum of member welfare:

$$G_{\text{EU}}(\bar{\boldsymbol{\tau}}) = \sum_{n \in \mathcal{E}} \omega_n \left[a_n W_n(\bar{\boldsymbol{\tau}}) + \lambda_n S_n(\bar{\boldsymbol{\tau}})\right] \tag{15}$$

where $\omega_n = \text{GDP}_n/\text{GDP}_\mathcal{E}$ are GDP weights, $a_n = 1$ is the welfare weight and λ_n the security weight for member n, and $S_n = -\sum_{j\in\mathcal{S}} \phi^j \sum_{m\in\mathcal{G}} (\pi^j_{nm})^2$ is a security index penalising import concentration from great powers in strategic sectors $\mathcal{S}$, with ϕ^j the strategic weight assigned to sector j.

4.1.5 Defection and Outside Options

The core collective action problem arises when a member n prefers a bilateral deal with a great power to the collective tariff. The *defection gap* is:

$$\Delta_n = \hat{W}_n^{\text{bilat}} - \hat{W}_n^{\text{EU}} \tag{16}$$

where $\hat{W}_n^{\text{bilat}}$ is welfare when member n is exempted from the external shock (the coercing power offers n a bilateral deal) while all other members maintain the collective stance. Member n prefers the bilateral option if and only if $\Delta_n > 0$.

The bilateral outside option warrants careful discussion. In the baseline specification, the coercing power offers full tariff exemption to the defecting member. This is the most generous possible bilateral deal. This is an appropriate *worst-case* benchmark, as it computes the fiscal capacity needed even when the hegemon plays optimally to break EU solidarity. The hegemon has a strategic incentive to make generous bilateral offers precisely because doing so undermines collective bargaining, creating what Clayton et al. (2024) term a "fragmentation doom loop." I show in sensitivity analysis that $\bar{F}$ declines monotonically as the bilateral offer worsens, but even in the face of no concession from the hegemon the fiscal cost is positive due to individual country gains from adjusting their retaliatory tariff relative to the collective.

4.1.6 Minimum Fiscal Capacity

Definition 1 (Minimum Fiscal Capacity)**.** The minimum fiscal capacity $\bar{F}$ is the smallest GDP-weighted transfer budget such that the common trade policy is incentive-compatible for all members:

$$\bar{F} = \sum_{n\in\mathcal{E}} \omega_n \cdot \max\{\Delta_n, 0\} \tag{17}$$

where $\omega_n = \text{GDP}_n/\text{GDP}_\mathcal{E}$ is member n's GDP share.

This GDP-weighted formulation gives the actual resource cost as a fraction of Eurozone GDP. A 5 percentage point defection gap for Ireland (2% of EZ GDP) costs far less in aggregate terms than a 1 percentage point gap for Germany (28% of EZ GDP). The incentive compatibility constraint requires a transfer vector $F_n \geq \Delta_n$ for each member n with a positive gap; $\bar{F}$ is the minimum aggregate cost of satisfying all such constraints simultaneously.

4.2. Layer 2: New Keynesian Monetary Union

4.2.1 Aggregate-Idiosyncratic Decomposition

Layer 2 does not change the aggregate $\bar{F}$—this is determined entirely by Layer 1—but answers a different question: given a fiscal facility of size $\bar{F}$, how should transfers be allocated across members when the ECB's one-size-fits-all policy amplifies some members' losses and compresses others'?

Following the approach in Farhi and Werning (2017) and Galí and Monacelli (2008), I decompose the NK dynamics into aggregate and idiosyncratic components. This avoids the well-known currency union indeterminacy problem.

The ECB sets the nominal interest rate according to a Taylor rule targeting GDP-weighted union aggregates:

$$i_t = \bar{r} + \phi_\pi \bar{\pi}_t + \phi_y \bar{y}_t \tag{18}$$

where $\bar{\pi}_t = \sum_{n \in \mathcal{E}} \omega_n \pi_{nt}$ and $\bar{y}_t = \sum_{n \in \mathcal{E}} \omega_n y_{nt}$ are GDP-weighted union averages of inflation and the output gap.

4.2.2 Aggregate Dynamics

The aggregate system is a standard closed-economy NK model, well-determined under the Taylor principle ($\phi_\pi > 1$). With AR(1) cost-push shocks $\bar{e}_t = \rho^t \bar{e}_0$ (where $\bar{e}_0$ is the GDP-weighted average welfare deviation from Layer 1):

$$\bar{\pi}_t = \alpha \cdot \bar{e}_t, \qquad \bar{y}_t = \gamma \cdot \bar{e}_t \tag{19}$$

where α and γ are functions of the NK Phillips curve slope κ, the Taylor rule coefficients (ϕ_π, ϕ_y), the shock persistence ρ, and the discount factor β:

$$\alpha = \frac{\sigma(1-\rho) + \phi_y}{(1-\beta\rho)[\sigma(1-\rho) + \phi_y] + \kappa(\phi_\pi - \rho)} \tag{20}$$

The Phillips curve slope is:

$$\kappa = \frac{(1-\xi)(1-\beta\xi)}{\xi}(\sigma + \phi_\ell) \tag{21}$$

where ξ is the Calvo parameter (probability of not adjusting prices).

4.2.3 Idiosyncratic Amplification

Members whose welfare deviations differ from the GDP-weighted aggregate face additional costs because the ECB's one-size-fits-all policy does not match their needs. The idiosyncratic welfare component of member n is:

$$\delta_n^{\text{idio}} = (\hat{W}_n - 1) - \underbrace{\sum_{m\in\mathcal{E}} \omega_m(\hat{W}_m - 1)}_{\delta^{\text{agg}}} \tag{22}$$

Under NK frictions, the idiosyncratic component is amplified:

$$\delta_n^{\text{NK}} = \delta^{\text{agg}} + \delta_n^{\text{idio}} \cdot \underbrace{\left(1 + \frac{\xi}{1-\xi}\right)}_{\text{amplification factor}} \tag{23}$$

By construction, $\sum_n \omega_n \delta_n^{\text{idio}} = 0$, so:

$$\sum_n \omega_n \delta_n^{\text{NK}} = \delta^{\text{agg}} \cdot \underbrace{\sum_n \omega_n}_{=1} + \frac{1+\xi/(1-\xi)}{1} \cdot \underbrace{\sum_n \omega_n \delta_n^{\text{idio}}}_{=0} = \delta^{\text{agg}} \tag{24}$$

This proves Corollary 1: the GDP-weighted aggregate is invariant to NK frictions.

4.2.4 Fiscal Transfer Instruments

I consider four fiscal instruments, nested by generality:

(i) **No transfers** (status quo): Each member absorbs its own shock through national fiscal policy.

(ii) **Contingent transfers** (Farhi-Werning optimal): Budget-balanced state-contingent transfers that equalise welfare across members:

$$F_n^{\text{FW}} = \bar{W} - W_n, \qquad \sum_{n\in\mathcal{E}} F_n^{\text{FW}} = 0 \tag{25}$$

(iii) **Common unemployment insurance**: Transfers proportional to the deviation of national unemployment from the union average:

$$F_n^{\text{UI}} = \beta_{\text{UI}}(u_n - \bar{u}) \cdot L_n \tag{26}$$

(iv) **IC-binding joint borrowing**: The EU borrows B_{EU} on capital markets and allocates to members proportional to their defection gaps:

$$F_n^{\text{IC}} = \Delta_n^+, \qquad B_{\text{EU}} = \sum_n F_n^{\text{IC}} \cdot \omega_n = \bar{F} \tag{27}$$

A key result, established formally in Section 6.4, is that instrument (ii) *cannot* satisfy incentive compatibility when all members have positive defection gaps. Budget balance ($\sum F_n = 0$) is inconsistent with $F_n \geq \Delta_n > 0$ for all n. Only instrument (iv), joint borrowing, achieves IC. Note that the IC-binding transfer $F_n = \Delta_n^+$ is the *minimum-cost* allocation that prevents defection; it is not necessarily welfare-optimal, and alternative instruments (e.g., ex-ante vulnerability investment) could in principle dominate.

5. Calibration

5.1. Data

The model is calibrated using the World Input-Output Database (WIOD) 2016 release, which provides bilateral trade flows and input-output tables for 43 countries and 56 sectors in 2014. I aggregate this to 23 countries and 10 sectors as described below.

5.2. Country Structure

The primary specification uses 20 individual Eurozone members plus three external economies. Table 2 lists the countries.

Table 2: *Country Structure*

Eurozone (20)	AUT, BEL, CYP, DEU, ESP, EST, FIN, FRA, GRC, HRV, IRL, ITA, LTU, LUX, LVA, MLT, NLD, PRT, SVK, SVN
External (3)	US, China, Rest of World

Notes: Individual country disaggregation avoids ad hoc bloc aggregation (e.g., "DE-Core" combining Germany, Austria, and the Netherlands) that masks heterogeneous exposures. This revealed Luxembourg as the second largest outlier after Ireland.

5.3. Sectoral Structure

I use $J = 10$ sectors chosen to isolate heterogeneous strategic dependencies across members: (1) Critical minerals & raw materials; (2) Energy; (3) Semiconductors & electronics; (4)

Pharmaceuticals & chemicals; (5) Agriculture & food; (6) Automotive & transport; (7) Machinery & equipment; (8) Financial & business services; (9) Digital services & ICT; (10) Other manufactures & services. Sectors 1–4, 8, and 9 are designated "strategic" (elevated weight in the security index).

5.4. Parameters

Trade shares π^j_{ni}, input-output coefficients γ^{kj}_n, value-added shares γ^j_n, and country GDP (used as labour endowments L_n) are taken directly from WIOD 2014. Trade elasticities θ^j follow Caliendo and Parro (2015), ranging from 2.5 to 8.0. Intra-EU non-tariff barriers are set uniformly at $d^j_{ni} = 1.15$ for $n, i \in \mathcal{E}$, consistent with Head and Mayer (2021).

NK parameters follow standard calibrations: $\xi = 0.75$ (Calvo), $\sigma = 1.0$ (CRRA), $\phi_\ell = 1.0$ (inverse Frisch), $\beta = 0.99$, $\varepsilon = 6.0$ (within-sector substitution, entering through the markup $\mu = \varepsilon/(\varepsilon - 1)$), $\rho = 0.85$ (shock persistence). The Taylor rule uses $\phi_\pi = 1.5$, $\phi_y = 0.5$.

5.5. Validation: NAFTA Replication

I validate the trade model against two external benchmarks: the NAFTA tariff removal of Caliendo and Parro (2015) and the 2018 US-China tariff war studied by Fajgelbaum et al. (2020) and Amiti et al. (2019).

Table 3 reports the NAFTA replication. The primary validation runs the solver on Caliendo and Parro's own 1993-calibrated inputs: their exact bilateral trade shares, input-output coefficients, tariff rates, and trade elasticities for 31 countries and 40 sectors.[5]

With matched-vintage inputs, the solver reproduces Caliendo and Parro (2015) closely. The US matches at 1.19× and Canada at 1.59×. Mexico undershoots at 0.64×, attributable to the omission of exogenous trade deficits: Caliendo and Parro model Mexico as running a trade deficit of approximately 3% of GDP, which amplifies the gains from tariff removal. The qualitative ordering is correct across all specifications: Mexico gains most, followed by Canada, then the US. The final column extends the same model to WIOD 2014 trade shares, which reflect 20 years of post-NAFTA integration. This results in magnitudes that are substantially larger.

[5] Data from Caliendo and Parro's replication files, as packaged by Aniceto (2023). The only model difference is that the solver does not account for exogenous trade deficits, which Caliendo and Parro include.

Table 3: *Validation: NAFTA Tariff Removal*

	CP2015 inputs	CP2015 target	Ratio	WIOD 2014
US	+0.10	+0.08	1.19×	+0.35
Canada	+0.10	+0.06	1.59×	+0.53
Mexico	+0.84	+1.31	0.64×	+3.23
Calibration details				
Countries	31	31		27
Sectors	40	40		18
Trade share vintage	1993	1993		2014

Notes: Welfare changes (%) from NAFTA tariff elimination. CP2015 inputs: solver run on their exact 1993-calibrated data. CP2015 target: their Table 5. WIOD 2014: post-NAFTA trade shares from the 2016 WIOD release (27 countries aggregated to match CP2015 coverage). Mexico discrepancy from omitted exogenous trade deficits (see text).

As a complementary benchmark, the model applied to the 2018 US-China tariff war produces a US welfare loss of 0.22% and a Chinese loss of 0.47%, qualitatively consistent with the partial-equilibrium estimates of Fajgelbaum et al. (2020) and Amiti et al. (2019) after accounting for the additional GE channels (intermediate inputs, wage adjustments, trade diversion).

Two features of $\bar{F}$ provide further robustness. First, $\bar{F}$ depends on *cross-country heterogeneity* in welfare effects, not on level estimates. Following Adão, Costinot, and Donaldson (2017), cross-country comparative statics are more robust than point predictions because they depend on trade share structure, which comes from data, rather than elasticity levels. Second, recalibrating the model using OECD ICIO tables for 2018 (to address post-2014 shifts including Brexit and US–China decoupling) moves $\bar{F}$ by less than 12% across all scenarios (Spearman $\rho = 0.77$, $p < 0.001$ for the cross-country ranking). The WIOD 2014 calibration is retained as the primary specification for comparability with Caliendo and Parro (2015).

6. Quantitative Results

This section presents the quantitative results. I begin with the core exercises: US tariff escalation, China supply restriction, and single market deepening. I then turn to fiscal instrument design, NK amplification, cost-benefit analysis, repeated game enforcement, and sensitivity analysis.

6.1. Exercise 1: US Tariff Escalation

The US imposes a 20% tariff on all EU exports. The EU retaliates with a symmetric 20% tariff on US imports. Each EZ member's outside option is a bilateral deal where the US exempts that member from its tariff.

Table 4 reports welfare changes for the most affected members under the retaliation scenario (full 20-member results in Appendix Table D4). Ireland loses 3.33% of welfare (driven by MNC pharmaceutical and technology exposure to the US), Luxembourg loses 2.54% (financial services), while Spain loses only 0.25%.

Table 4: *Exercise 1: US 20% Tariff + EU Retaliation*

Member	$\hat{W}_n^{\mathrm{EU}} - 1$ (%)	Gap Δ_n (pp)	GDP wt (%)
IRL	−3.33	4.06	2.0
LUX	−2.54	3.77	0.8
BEL	−0.98	1.16	4.4
DEU	−0.71	0.79	28.2
NLD	−0.71	0.78	6.7
AUT	−0.48	0.62	3.2
FIN	−0.53	0.57	2.1
FRA	−0.50	0.52	20.0
ITA	−0.41	0.45	16.3
ESP	−0.25	0.30	10.3
US	−0.50	—	—

Notes: 19 of 20 EZ members have positive defection gaps (prefer bilateral deals). Gap is: $\Delta_n = \hat{W}_n^{\mathrm{bilat}} - \hat{W}_n^{\mathrm{EU}}$ in percentage points.

The cross-member heterogeneity in Table 4 reflects differences in sectoral exposure. Following Arkolakis et al. (2012), I decompose each country's welfare change into sectoral contributions proportional to $-(\alpha_{nj}/\theta^j)\log\hat{\pi}_{nn}^j$, where $\hat{\pi}_{nn}^j$ is the change in the domestic expenditure share in sector j. Sectors where a country loses domestic market share (as imports become more competitive) contribute negatively; sectors where home production gains contribute positively.

Table 5 reports the top sectoral drivers for the five most exposed members. Ireland's loss is overwhelmingly driven by pharmaceuticals/chemicals and other manufacturing (MNC-intensive sectors with high US export exposure). Germany's loss is spread across autos, chemicals, and machinery. The Netherlands' exposure comes through chemicals/petroleum (intermediate input channels). These sector-specific patterns explain why a uniform tariff produces such heterogeneous welfare effects: the tariff interacts with pre-existing production structure through the input-output network.

Table 5: *Welfare Decomposition: Top Sectoral Drivers*

Member	Top losing sectors	Key channel
IRL	Pharma/chemicals (-0.87pp), other mfg (-0.68pp)	MNC US export exposure
LUX	Financial intermediation, chemicals	Financial services hub
DEU	Pharma/chemicals (-0.16pp), autos (-0.11pp)	Diversified manufacturing
NLD	Chemicals/petroleum (-0.25pp), pharma (-0.14pp)	Intermediate inputs
FRA	Autos (-0.15pp), pharma/chemicals (-0.11pp)	Key export sectors

Notes: ACR-style decomposition of welfare changes from Table 4. Contributions in percentage points of baseline welfare. The total welfare change equals the sum of all sectoral contributions, with I-O linkage amplification distributed proportionally.

Nineteen of twenty Eurozone members prefer bilateral deals (positive defection gaps), meaning nearly every member would have incentive to defect without transfers. The GDP-weighted minimum fiscal capacity is:

$$\bar{F} = \sum_n \omega_n \Delta_n^+ = 0.69\% \text{ of EZ GDP} \approx €97 \text{ billion}$$

Despite Ireland's large individual gap (4.06pp), its GDP-weighted contribution is only $4.06\% \times 2.0\% \approx 0.08$pp of the 0.69% total. The aggregate $\bar{F}$ is dominated by Germany's moderate gap (0.79pp) applied to its large GDP weight (28.2%), contributing 0.22pp. Figure 2 decomposes $\bar{F}$ into individual member contributions, revealing how moderate gaps applied to large GDP weights dominate the aggregate cost.

A natural question is: how much of $\bar{F}$ arises from EU retaliation rather than the US tariff alone? Without EU retaliation (the EU simply absorbs the US tariff), $\bar{F} = 0.47\%$, which is 32% lower than with retaliation. EU retaliation adds 0.22pp to $\bar{F}$ because it amplifies the asymmetries in welfare effects. Members with high US trade exposure suffer disproportionately from the retaliatory tariff as well. This does not imply the EU should forgo retaliation, as retaliation may be optimal for the bloc as a whole. Instead it establishes that the *fiscal cost of collective action includes the cost of the collective response*, not just the external shock. Section 6.6 returns to this observation, asking whether the cooperative strategy is worth pursuing relative to fragmentation.

Figure 3 shows how $\bar{F}$ scales with the US tariff rate. The relationship is sub-linear: at 5%, $\bar{F} = 0.23\%$; at the baseline 20%, $\bar{F} = 0.69\%$; at 50%, it reaches only 1.15%. At low tariff rates Luxembourg has the largest individual defection gap, but Ireland overtakes it above 15% as pharmaceutical and technology exposure dominates.

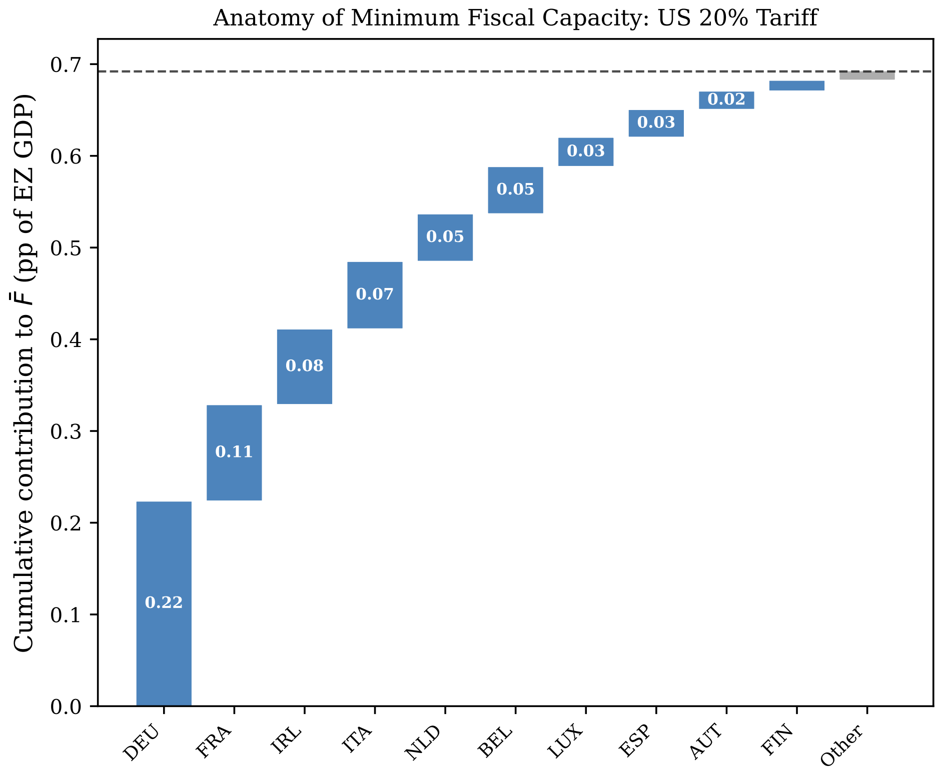

Figure 2: *GDP-Weighted Contributions to $\bar{F}$: US 20% Tariff*

Notes: Each bar shows member n's contribution $\omega_n \Delta_n^+$ to aggregate $\bar{F}$. Germany contributes the most despite having only the fifth largest individual gap, because its 28% GDP weight dominates. Remaining members grouped as "Other".

Endogenous Bilateral Offers. The baseline analysis treats the bilateral outside option as full tariff exemption ($\tau_{\text{bilat}} = 0\%$). This is the most generous offer the hegemon could make. In practice, the US maximises its own welfare subject to each member accepting the bilateral deal. For each member n, the US solves:

$$\max_{\tau_{\text{offer}}} W_{\text{US}}(\tau_{\text{offer}}, n) \quad \text{s.t.} \quad W_n^{\text{bilat}}(\tau_{\text{offer}}) \geq W_n^{\text{EU}}$$

For 19 of 20 members, the hegemon-optimal offer is $\tau_{\text{offer}} = 20\%$, which is to say that the US maintains its full tariff and offers *no concession*. Members still prefer the bilateral deal because they gain from unilaterally dropping the EU's retaliation tariff. The defection incentive comes from avoiding the cost of collective retaliation, not from a US tariff exemption.

This finding is not contradictory with the sectoral decomposition in Table 4, which attributes Ireland's cooperative-regime welfare loss to US tariff exposure in pharmaceuticals and technology. In a GE model with input-output linkages, the EU's retaliatory tariff raises costs for Irish firms through imported intermediates and diverts trade flows away from

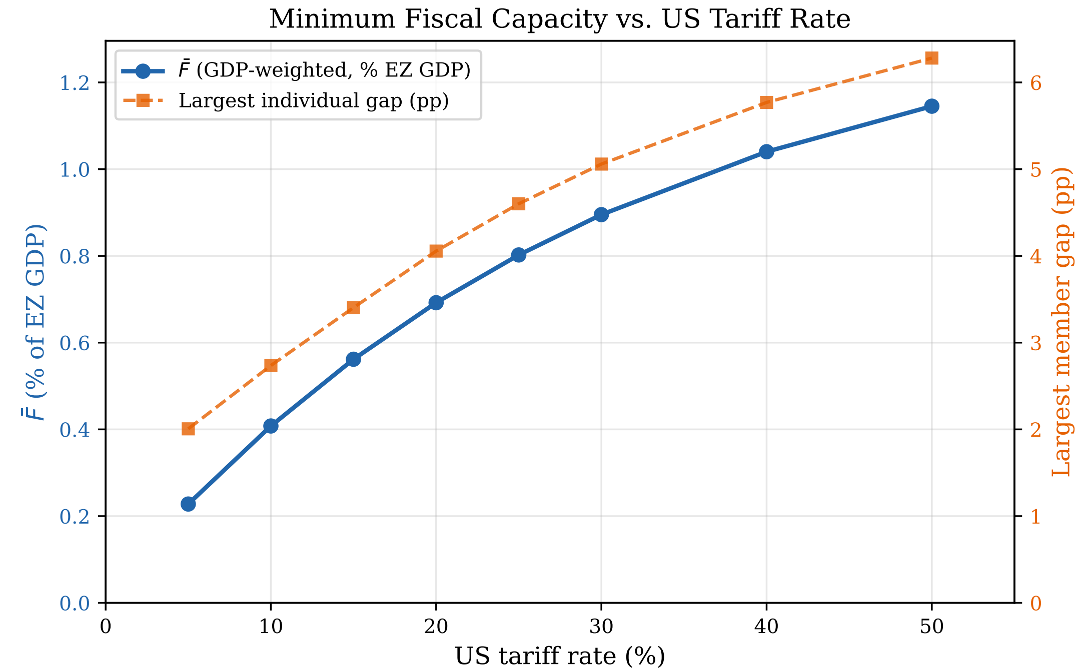

Figure 3: *Minimum Fiscal Capacity and Maximum Gap by US Tariff Rate*
Notes: Left axis shows GDP-weighted $\bar{F}$ (% of EZ GDP). Right axis shows the largest individual defection gap (Luxembourg below 15%, Ireland above). Both exhibit sub-linear scaling: doubling the tariff from 10% to 20% adds only 0.28pp to $\bar{F}$.

efficient supply chains. A member can simultaneously be hurt by tariffs on its *exports* (the source of the cooperative welfare loss) and benefit from dropping retaliation on its *imports* (the source of the defection gain). The two decompositions—"what causes the loss" and "what drives the gap"—operate through different channels in the I-O network.

Under hegemon-optimal offers, $\bar{F} = 0.25\%$ of EZ GDP (€35bn)—36% of the free-trade bilateral estimate. The US's most valuable targets are Germany (+0.10pp US welfare gain), France (+0.09pp), and the Netherlands (+0.06pp), reflecting the large bilateral trade volumes at stake. This endogenises what was previously a free parameter and tightens the range for $\bar{F}$: the *minimum* is 0.25% (hegemon extracts maximum surplus) and the *maximum* is 0.69% (hegemon offers full exemption). The sensitivity analysis to τ_{bilat} in Figure 7 spans this range, validating the sweep approach.

The endogenous bilateral result raises a natural question: if retaliation is individually costly for every member, and the hegemon need not even offer a tariff concession to induce defection, why does the EU retaliate at all? The answer lies in institutional sequencing and collective rationality. Under Article 207 TFEU, the European Commission proposes trade policy measures and the Council approves by qualified majority. Retaliation is decided *collectively* before individual exit options crystallise. At the point of the Council vote, each member may support retaliation for strategic reasons (signalling resolve, protecting specific industries, domestic political economy) even if the *ex post* defection incentive would make unilateral exit attractive.

This is a standard collective action problem: retaliation improves the EU's *aggregate* bargaining position even though it hurts every member individually. No single member's defection changes the aggregate outcome much—but if all defect, the EU loses all bargaining leverage. The fiscal facility solves precisely this problem. Without it, the EU's collective trade stance is not time-consistent: the Council votes to retaliate, but once the hegemon offers bilateral deals, individual members face irresistible pressure to break ranks. $\bar{F}$ is the cost of making the collective decision stick.

Bloc Bargaining: The Value of Collective Leverage. The endogenous bilateral exercise shows that individual members have zero bargaining leverage: the hegemon offers no concession because each member is small relative to the US. But the EU *as a bloc* has substantial leverage through credible retaliation.

EU retaliation costs the US 0.25 percentage points of welfare, roughly matching the welfare cost the US bears from its own tariff. Individual members cannot impose this cost; only the collective can. But this leverage cannot be deployed without fiscal capacity, because the EU's retaliation threat is not credible if members will defect under pressure.

Solving for the precise negotiated tariff would require modelling the hegemon's domestic political economy motivation for the tariff, which is outside the scope of this paper (the tariff is exogenous throughout). Instead, Table 6 reports the EU's welfare and fiscal capacity under a range of *illustrative* negotiated tariffs, each assuming the EU trades the withdrawal of retaliation for a tariff reduction. These should be read as counterfactual scenarios, not equilibrium predictions.

Table 6: *Bloc Bargaining: Outcomes Under Negotiated Tariffs*

US tariff	EU retal.		$\bar{F}$ (% EZ GDP)	EZ welfare (%)	Gain vs confront. (pp)
0%	none	(free trade)	0.02	+0.00	+0.61
5%	none	(negotiated)	0.18	−0.16	+0.45
10%	none	(negotiated)	0.31	−0.28	+0.33
15%	none	(negotiated)	0.41	−0.38	+0.23
20%	none	(no deal)	0.47	−0.45	+0.16
20%	20%	(confrontation)	0.69	−0.61	—

Notes: Rows 1–5: EU withdraws retaliation in exchange for a negotiated US tariff reduction. Last row: confrontation baseline (symmetric 20% tariffs). Bilateral outside option: US exempts individual member. EU retaliation costs the US 0.25pp of welfare.

Individual members extract nothing; the bloc has leverage that individual members lack. At an illustrative negotiated outcome of 10% (halving the tariff), $\bar{F}$ would fall from 0.69% to 0.31% of EZ GDP and aggregate welfare would improve by 0.33pp relative to confrontation. If collective leverage translates into negotiated tariff reductions, the fiscal facility could partially pay for itself by reducing both the external tariff and the internal fiscal requirement. Whether this potential is realised depends on factors outside the model, including the hegemon's domestic political constraints.

6.2. Exercise 2: China Critical Minerals and Semiconductors

China restricts critical mineral exports to the EU by 50% and semiconductor exports by 30%. The EU retaliates with a 20% tariff on Chinese imports. Each member's outside option is a bilateral exemption from the Chinese restriction. Full member-level results are in Appendix Table D5.

The GDP-weighted minimum fiscal capacity is:

$$\bar{F} = 0.44\% \text{ of EZ GDP} \approx €62 \text{ billion}$$

The critical coalition is qualitatively different from Exercise 1. The Netherlands and Germany are the most exposed members (not Ireland), driven by their high China intermediate input shares in machinery and chemicals. This implies that the fiscal facility must be designed to cover different exposure patterns under different shocks.

EU retaliation on China substantially increases the fiscal need (from $\sim$0.27% without retaliation to 0.44% with), echoing the analytical result that the retaliatory tariff level is the strongest driver of $\bar{F}$. In the worst case—a simultaneous US tariff escalation and Chinese supply restriction, with EU retaliation against both—the required fiscal capacity is $\bar{F}^{\text{combined}} = 1.12\%$ of EZ GDP (€157 billion), comparable to NextGenerationEU (€750 billion over 2021–2026, or approximately €125 billion per year). The combined shock is roughly additive ($\bar{F}^{\text{combined}}/(\bar{F}^{\text{US}} + \bar{F}^{\text{China}}) \approx 1.00$), because the two shocks operate through different sectors and bilateral exposures.

Table 7 summarises the headline results across all exercises, and Figure 4 visualises the coalitional shift: the US tariff creates an Ireland/Luxembourg-led critical coalition, while the China shock shifts exposure to the Netherlands and Germany.

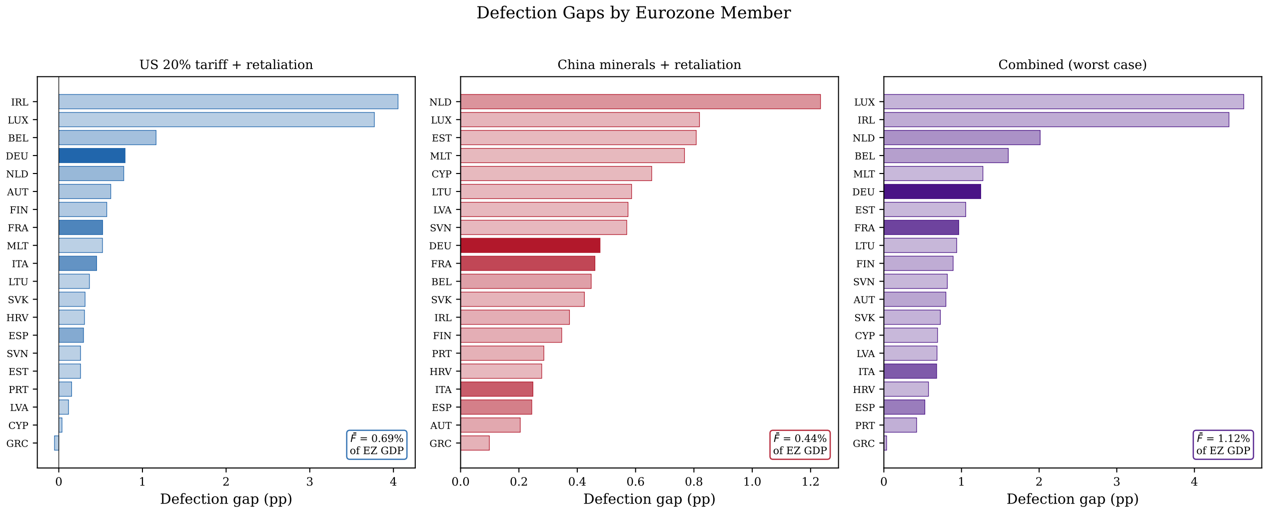

Figure 4: *Defection Gaps by Member and Shock Scenario*
Notes: Horizontal bars show individual defection gaps (pp) for all 20 EZ members, sorted by gap size. Bar opacity encodes GDP weight (darker = larger economy). The three panels show the US tariff, China restriction, and combined shock scenarios respectively.

Table 7: *Cross-Exercise Comparison*

Scenario	$\bar{F}$ (% EZ GDP)	∼€ bn	Most exposed
US 20% tariff + retal.	0.69	97	IRL, LUX
China minerals + retal.	0.44	62	NLD, DEU
Combined (worst case)	1.12	157	IRL, LUX, NLD

Notes: Different shocks create different critical coalitions. All numbers use GDP-weighted $\bar{F}$ with 23-country calibration. Expanded version in Appendix Table D7.

6.3. Exercise 3: Single Market Deepening

To understand how integration across the Eurozone affects these results, I reduce intra-EU non-tariff barriers by 25% (as proposed in the Letta Report). Then I repeat Exercise 1, recomputing $\bar{F}$. Single market deepening generates large welfare gains: Luxembourg gains 3.80%, Slovenia 2.48%, Belgium 2.46%, and small open economies benefit most due to their higher intra-EU trade shares.

However, SM deepening barely changes the minimum fiscal capacity: $\bar{F}$ moves only marginally (from 0.69% to approximately 0.70% of EZ GDP) after a 25% NTB reduction, and falls only to 0.67% even under complete NTB elimination. This confirms Proposition 2: deeper market integration benefits both the "stay in the EU" and "defect to a bilateral deal" options roughly equally. The defection gap, and the fiscal transfer needed to prevent it, is nearly invariant to single market depth. In other words, market integration and fiscal capacity are complements, not substitutes.

6.4. Fiscal Instrument Design and Farhi-Werning Validation

A central result is that no budget-balanced instrument can sustain collective trade policy when all members prefer bilateral deals. The binding gap is Ireland's at 4.06pp. Farhi and Werning (2017)-style contingent transfers redistribute across members but cannot close all gaps simultaneously: budget balance requires $\sum F_n = 0$, which contradicts $F_n \geq \Delta_n > 0$ when every Δ_n is positive. A common unemployment insurance scheme covers only 0.03pp of Ireland's 4.06pp gap. Only deficit-financed joint borrowing, at 0.69% of EZ GDP, achieves incentive compatibility by injecting net resources rather than merely redistributing.

An important subtlety arises: if joint EU debt is repaid by GDP-proportional member contributions, then joint borrowing is present-value neutral. Each member's lifetime transfer nets to zero once repayment is included. Does this mean joint borrowing fails to relax the constraint? No, because the IC constraint binds *contemporaneously*—at the moment of the tariff shock, when governments face pressure to pursue bilateral deals. The per-period repayment cost confirms the asymmetry: for $\bar{F} = 0.69\%$ of EZ GDP, repaid over 30 years at 3% interest, the annual cost is approximately 0.03% of GDP—less than one hundredth of Ireland's acute defection gain (4.06pp).

Joint borrowing therefore solves a *timing* problem: it shifts resources to the present, when the collective action problem is acute, with repayment spread over decades when per-period costs are negligible. The claim is not that joint borrowing creates resources from nothing, but that contemporaneous budget balance is impossible when all members have positive defection gaps.

The allocation rule matters: under GDP-proportional allocation, each member receives the same transfer as a share of its own GDP, so the facility must equal the *maximum* individual gap (Ireland's 4.06pp), yielding a facility of 4.06% of EZ GDP—roughly 6× larger than needs-based allocation. Needs-based allocation, which gives each member exactly its defection gap, is therefore essential to keeping the facility at 0.69%.

6.5. New Keynesian Amplification

Table 8 illustrates a direct implication of Corollary 1: because the ECB targets the GDP-weighted average inflation, NK frictions leave the aggregate minimum fiscal capacity unchanged at 0.69% for all values of the Calvo parameter ξ, while dramatically redistributing individual gaps.

Table 8: *NK Amplification of Defection Gaps (US Tariff Scenario)*

Member	Flex gap (pp)	NK gap (pp, $\xi = 0.75$)	Ratio
IRL	4.06	12.23	3.0×
LUX	3.77	9.57	2.5×
DEU	0.79	1.10	1.4×
FRA	0.52	0.30	0.6×
Aggregate $\bar{F}$ (% EZ GDP):			
Flexible	0.69		
NK ($\xi = 0.75$)	0.69		1.0×

Notes: NK amplification is mean-preserving under GDP weights. Small exposed members (IRL, LUX) see gaps widen dramatically; large members closer to the ECB target (DEU) see modest amplification; France's gap shrinks because its welfare deviation is close to the GDP-weighted mean. Transition dynamics at $\xi = 0.75$: Ireland inflation +10.0%, Germany +2.1%, ECB rate 1.8%, half-life ≈ 4.3 quarters.

The ECB targets the GDP-weighted average. Members close to this average (large economies) are approximately stabilised. Members far from the average (small, exposed economies) face a large monetary mismatch that amplifies their welfare deviation by a factor of $1/(1-\xi) = 4$ at $\xi = 0.75$, but the GDP-weighted sum is unchanged by construction.

Figure 5 visualises this redistribution. Members above the 45° line see their gaps amplified by NK frictions; members below see gaps compressed. Bubble size is proportional to GDP weight, showing that the large members whose gaps shrink dominate the aggregate.

The key takeaway is that while NK frictions do not increase the total cost of the fiscal facility, the €97 billion needed for the US tariff scenario is the same with or without sticky prices, they do dramatically change *who* needs transfers. Members far from the ECB-weighted average see their gaps amplified; those near it see gaps compressed. This is a facility *design* problem, not a facility *sizing* problem. Appendix Figure E1 shows the underlying NK transition dynamics, with a half-life of approximately 4.3 quarters.

6.6. Cost-Benefit Analysis

The preceding analysis establishes *how much* the facility costs. A natural follow-up: is it worth it? I compare EZ-wide GDP-weighted welfare under cooperation (with retaliation and the fiscal facility) against full fragmentation (all members defect, EU retaliation collapses). Modelling the strategic formation of welfare-reducing tariffs is beyond the scope of this paper, but this comparison provides a static cost-benefit benchmark and a basis for assessing how much deterrence value the facility would need to justify its cost.

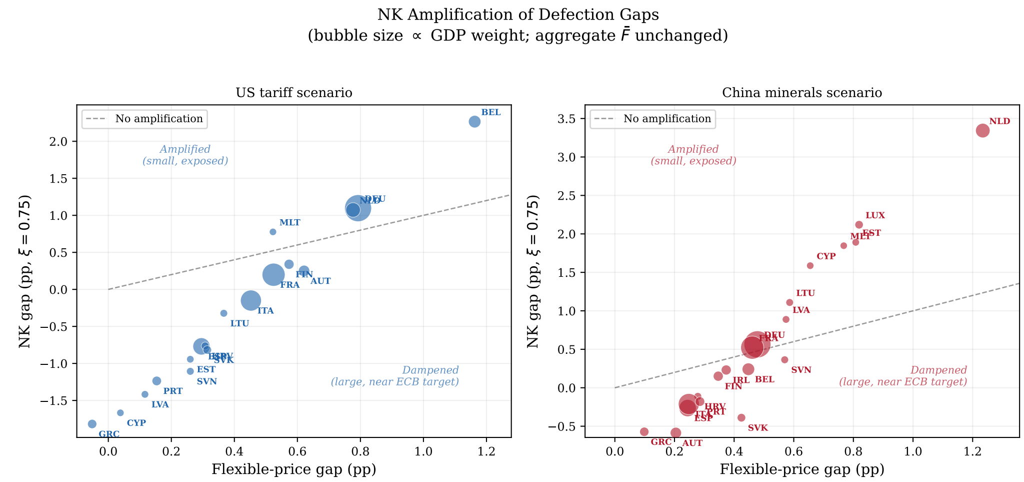

Figure 5: *NK Amplification of Individual Defection Gaps*
Notes: Each bubble is an EZ member. Horizontal axis: flexible-price defection gap. Vertical axis: NK gap ($\xi = 0.75$). The 45° line marks no amplification. Bubble area proportional to GDP weight. Ireland (4.06→12.23pp) and Luxembourg (3.77→9.57pp) omitted from left panel for readability; both are strongly amplified. Large members near the ECB-weighted average (e.g., Italy, Spain) are compressed toward zero.

Under cooperation, the EU retaliates but must finance transfers totalling $\bar{F} = 0.69\%$ of EZ GDP. GDP-weighted EZ welfare falls by 0.61%. Under fragmentation, the US maintains its 20% tariff but no member retaliates. GDP-weighted welfare falls by only 0.45%. Retaliation therefore makes the EU *statically worse off* by 0.16pp. Note that the facility cost $\bar{F} = 0.69\%$ is not an additional welfare loss—it is a transfer within the Eurozone that redistributes resources from net contributors to exposed members. The aggregate welfare cost of cooperation relative to fragmentation is the 0.16pp difference, not $0.16 + 0.69 = 0.85$pp. This is a well-known result in trade theory: retaliatory tariffs reduce welfare for the retaliating bloc when the adversary is a large economy (see, for example: Bagwell and Staiger 1999).

The case for the facility therefore rests on *dynamic* considerations. Without credible retaliation capacity, the EU faces:

1. *Deterrence failure:* the hegemon faces no cost from tariff escalation, inviting further aggression.
2. *Bargaining collapse:* the EU's negotiating position in future trade agreements is undermined.
3. *Precedent effects:* successful coercion signals that the EU's collective stance can be broken at low cost.

Proposition 4 formalised the deterrence channel: in a repeated game, the fiscal facility reduces each member's critical discount factor β_n^*, making cooperative retaliation sustainable at realistic patience levels.

Deterrence value. The logic follows Bagwell and Staiger (1999): trade agreements, and the enforcement infrastructure behind them, generate value by changing the *equilibrium* tariff, not the off-equilibrium payoff. If the EU has fiscal capacity to sustain retaliation, the hegemon's first-period tariff choice incorporates the cost of triggering a credible response. The facility may deter a shock that never materialises *because* the facility exists.

In this context the US tariff costs the EU 0.45% of GDP under fragmentation (no retaliation) and 0.61% under retaliation. If the facility deters the tariff entirely, the EU saves 0.45–0.61% of GDP per period. To illustrate, consider purely hypothetical deterrence probabilities: a facility costing 0.69% that has even a 50% probability of deterring the shock pays for itself in expected value within two years; at a 30% deterrence probability, within four years. A true cost-benefit analysis requires some understanding of this deterrence benefit, but these illustrative examples show that the potential benefits of improved deterrence could reasonably cover the costs estimated here.

Further, the 0.69% should not be understood as annual expenditure. $\bar{F}$ is *contingent capacity*: the EU borrows only if the shock materialises. The standing cost is the administrative infrastructure and the credible political commitment to activate the facility, not €97 billion per year. The NextGenerationEU precedent demonstrates that contingent joint borrowing, while politically difficult to establish, has near-zero carrying cost until activated. The 0.69% is the *size* of the insurance policy, not its *premium*.

The facility cost is therefore the price of strategic credibility: an investment in the EU's capacity to act collectively, whose returns are measured not only in the off-equilibrium payoff but in the equilibrium tariff that credible collective action helps to prevent.

6.7. Repeated Game and Enforcement

Section 3.3 established that in an infinitely repeated game with discount factor $\beta \in (0,1)$, fiscal transfers are needed when $\beta < \beta_n^* = \Delta_n/(\Delta_n + V_n)$. The critical discount factor depends on the punishment V_n—the per-period cooperation surplus that the defecting member forfeits. I consider two punishment regimes that bracket the range of institutional plausibility.

Nuclear punishment: EU exclusion. The upper bound on punishment severity simulates full EU exclusion: a defecting member faces 25% higher NTBs and 5% tariffs on all intra-EU trade, the "nuclear option." Under this severe punishment, cooperation surpluses are large: Ireland's $V_{\text{IRL}} = 3.4$pp, Germany's $V_{\text{DEU}} = 4.7$pp, Belgium's $V_{\text{BEL}} = 10.9$pp. The binding

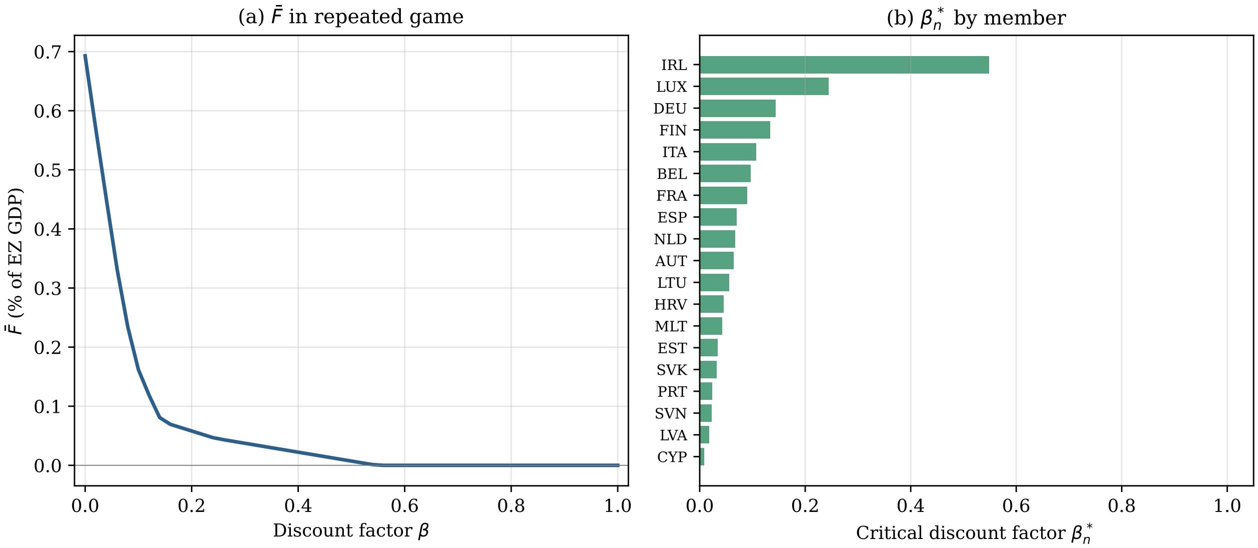

Figure 6: *Repeated Game: $\bar{F}(\beta)$ and Critical Discount Factors*

Notes: Panel (a) shows $\bar{F}$ as a function of the discount factor β under nuclear punishment (EU exclusion). $\bar{F}$ falls from 0.69% (one-shot) to zero at $\beta = 0.55$. Panel (b) shows the critical discount factor β_n^* for each member under nuclear punishment.

constraint is Ireland with $\beta^*_{\text{IRL}} = 0.55$. Cooperation can be sustained without transfers at any discount factor above 0.55 without need for fiscal compensation.

Credible punishment: cross-conditionality. While a policy of EU exclusion could align member incentives, this punishment is politically implausible. The EU has no legal mechanism to deny single market access to a member short of Article 50 proceedings. A more credible punishment is cross-conditionality: defecting members lose their net EU fiscal transfers. Under this regime, the 9 net contributors face zero punishment ($V_n = 0$), giving $\beta_n^* = 1.0$—the folk theorem never sustains cooperation for them regardless of patience. These members account for 94% of GDP-weighted $\bar{F}$. Fiscal transfers are therefore a *substitute for non-credible punishments*: they make cooperation incentive-compatible for the large net contributors whom the EU's existing compliance infrastructure cannot deter. Section 7.1 develops this in detail and estimates the fiscal capacity needed when taking these mechanisms into account.

6.8. Sensitivity and Robustness

Excluding Ireland and Luxembourg. Ireland's and Luxembourg's trade data are potentially distorted by multinational profit shifting. Table 9 shows that excluding these members has a modest impact on aggregate $\bar{F}$, confirming that results are not driven by MNC-distorted data.

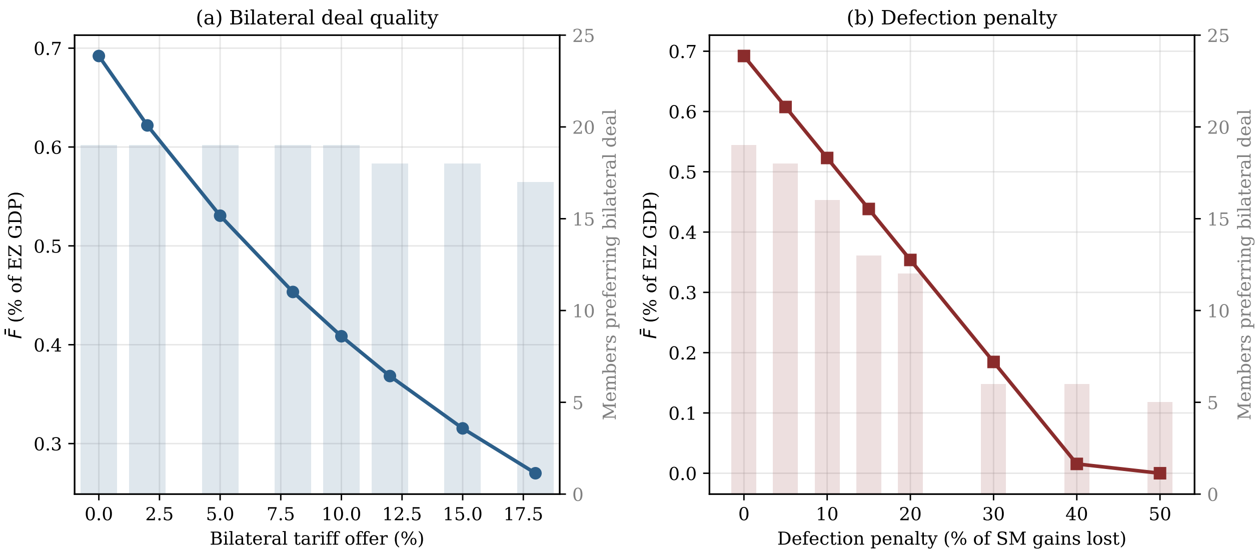

Figure 7: *Sensitivity of $\bar{F}$ to Bilateral Deal Quality and Defection Penalty*
Notes: Panel (a) varies the bilateral tariff the US offers a defecting member (0% = full exemption, 20% = no deal). Panel (b) varies the fraction of single market gains lost upon defection. Bars show the number of members preferring bilateral deals. US 20% tariff + EU retaliation scenario.

Table 9: *Robustness: $\bar{F}$ Excluding Ireland and Luxembourg*

Scenario	Full (20)	Ex-IRL (19)	Ex-IRL/LUX (18)
US tariff	0.69%	0.62% (−10%)	0.59% (−14%)
China shock	0.44%	0.45% (+0%)	0.44% (−0%)
Combined	1.12%	1.05% (−6%)	1.02% (−9%)

Notes: Percentage changes relative to the full 20-member EZ in parentheses. Excluded members' GDP weights are redistributed proportionally to remaining members.

Ireland and Luxembourg have the largest individual defection gaps but their small GDP weights (∼2% and <1% respectively) limit their aggregate impact. The China shock results are entirely unaffected because the critical coalition (NLD, DEU) is different.

Bilateral deal quality and defection penalty. My baseline result assumes the most generous possible outside option: full tariff exemption ($\tau_{\text{bilat}} = 0$) and no penalty for defection ($\delta_{\text{pen}} = 0$), which gives a sensible upper bound for $\bar{F}$. Figure 7 shows how $\bar{F}$ declines as either assumption is relaxed.

Panel (a) shows that $\bar{F}$ declines smoothly from 0.69% under full exemption to 0.27% when the bilateral offer is only a 2pp tariff reduction. Even a modest 5% bilateral tariff reduces $\bar{F}$ by 23%, but 19 members still prefer bilateral deals. The result is robust: $\bar{F}$ remains

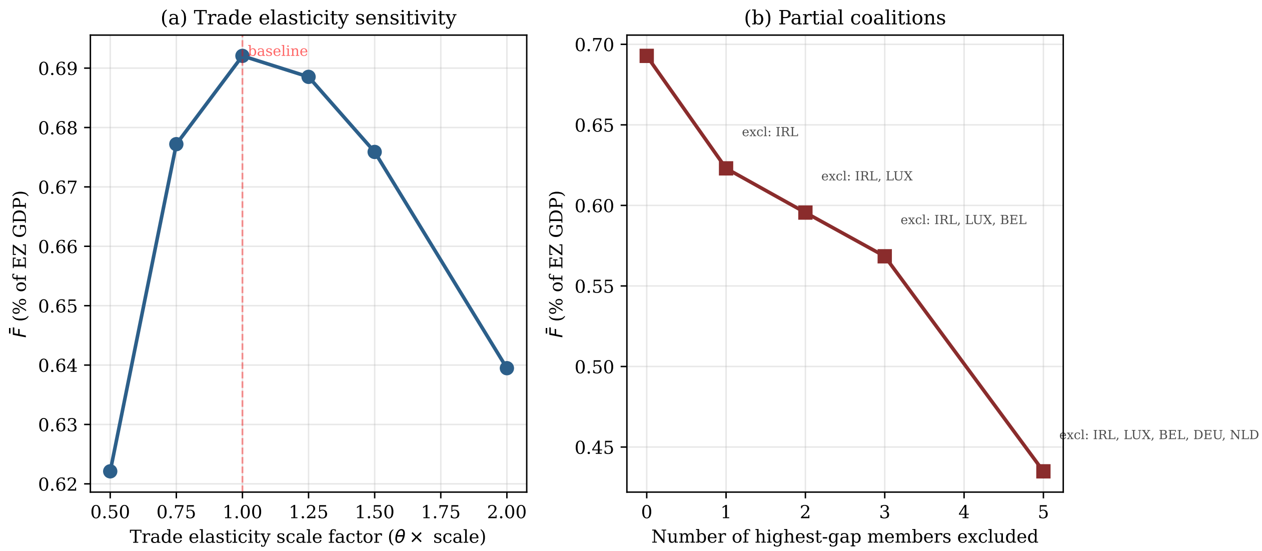

Figure 8: *Robustness: Trade Elasticities and Partial Coalitions*
Notes: Panel (a) scales all trade elasticities by a common factor; baseline is 1.0. Panel (b) sequentially excludes the members with the largest defection gaps. Labels show excluded members.

above 0.25% of EZ GDP unless the bilateral offer is essentially the same as the collective tariff.

Panel (b) shows that $\bar{F}$ is approximately linear in the defection penalty. At a 10% loss of single market gains, four members no longer prefer bilateral deals and $\bar{F}$ falls to 0.52%. Even at 30% SM loss, $\bar{F} = 0.18\%$, and at 50% only five members still prefer bilateral deals. The baseline of $\delta_{\text{pen}} = 0$ is conservative; any positive defection penalty reduces $\bar{F}$ below the headline estimate.

Trade elasticity and partial coalitions. Figure 8 reports two additional robustness exercises. Panel (a) scales all trade elasticities θ^j by a common factor. $\bar{F}$ varies by approximately $\pm 15\%$ around the baseline. The result is robust to substantial uncertainty in trade elasticities.

Panel (b) shows how $\bar{F}$ changes as the highest-gap members are excluded from the coalition. Dropping Ireland reduces $\bar{F}$ by 10%; dropping Ireland and Luxembourg by 14%. The collective action problem is not driven by a few outliers but is embedded in the trade structure of the entire Eurozone.

7. Institutional Mechanisms and Practical Fiscal Requirements

The headline estimate $\bar{F} = 0.69\%$ of Eurozone GDP for a US tariff escalation is an upper bound; it assumes *every* member must be made strictly whole relative to its bilateral outside option. This section grounds the analysis in existing EU institutions to derive a more practical fiscal requirement. Two mechanisms reduce $\bar{F}$ substantially: cross-conditionality with existing EU fiscal flows, and the qualified majority voting rules that govern Council decisions.

Throughout this section, I reframe $\bar{F}$ as the *cost of making collective trade policy Pareto-improving for all members*. In other words, the fiscal transfer budget needed so that no member loses from the collective stance relative to its outside option. This framing applies regardless of whether one interprets the outside option as literal defection from the customs union (which is legally constrained) or as the relevant benchmark for political sustainability: a member that is worse off under the collective policy has a political incentive to block, delay, water down, or circumvent it through the many informal channels available within EU institutions.

7.1. Cross-Conditionality with Existing EU Fiscal Flows

The EU's existing fiscal architecture creates asymmetric compliance incentives. Net recipient member states (those receiving more from the EU budget than they contribute) face an implicit cost to non-compliance: existing transfers can be made conditional on adherence to collective trade policy. This mechanism is not hypothetical; the Rule of Law Conditionality Regulation (2020/2092) already links EU budget disbursements to governance standards, and the Anti-Coercion Instrument (Regulation 2023/2675) establishes a framework for trade policy coordination.

I define the cross-conditionality "stick" for member n as:

$$\mathrm{XC}_n = \max\big(0,\ b_n\big), \tag{28}$$

where b_n is the net EU operating budgetary balance of member n as a share of GNI (positive for net recipients, negative for net contributors). Data are from the European Commission Financial Report 2022. A member's *residual defection gap* after cross-conditionality is then

$$\Delta_n^{\mathrm{XC}} = \max\big(0,\ \Delta_n - \mathrm{XC}_n\big), \tag{29}$$

where $\Delta_n = W_n^{\mathrm{bilat}} - W_n^{\mathrm{EU}}$ is the baseline defection gap from Section 6. Members for whom existing EU transfers exceed their defection gap ($\Delta_n \leq \mathrm{XC}_n$) require no additional fiscal capacity. The threat of losing existing transfers is sufficient to sustain compliance. The

Table 10: *Cross-Conditionality with Existing EU Fiscal Flows*

	US Tariff		China Supply		Combined
	$\bar{F}$	Reduction	$\bar{F}$	Reduction	$\bar{F}$
Full compensation	0.69%	—	0.44%	—	1.12%
After cross-cond.	0.66%	5.3%	0.40%	10.1%	1.06%
Members fully covered	8/20		9/20		8/20

Notes: Full compensation is $\bar{F}$ from Section 6. After cross-conditionality applies Equation (30). "Members fully covered" counts those with $\Delta_n \leq XC_n$. Net EU budgetary balances from European Commission Financial Report 2022.

cross-conditionality-adjusted fiscal capacity is given by:

$$\bar{F}^{XC} = \sum_{n=1}^{N_E} \Delta_n^{XC} \cdot \omega_n, \tag{30}$$

where $\omega_n = \text{GDP}_n/\text{GDP}_\mathcal{E}$ is the GDP weight from Definition 1.

Table 10 reports the results. Of the 20 Eurozone members, 11 are net recipients of EU fiscal flows. In the US tariff scenario, existing EU transfers fully cover the defection gaps of 8 small net recipients (the three Baltic states, Croatia, Slovakia, Slovenia, Portugal, and Cyprus), whose gaps range from 0.04 to 0.37 percentage points and are dwarfed by EU transfers of 0.41 to 2.87% of GNI. However, $\bar{F}^{XC}$ falls only modestly, from 0.69% to 0.66%—a reduction of just 5%. The reason is structural: the members that drive $\bar{F}$ through their GDP weights (Germany (28.2%), France (20.0%), Italy (16.3%), the Netherlands (6.7%)) are all *net contributors* to the EU budget, for whom $\text{XC}_n = 0$. Cross-conditionality has zero bite on the members that matter most.

This asymmetry is a central finding. The existing EU fiscal architecture provides compliance incentives precisely where they are least needed (small, already-dependent members) and none where they are most needed (large net contributors whose defection gaps dominate the GDP-weighted aggregate).

7.2. Qualified Majority Voting (QMV) and Blocking Minorities

EU trade policy under Article 207 TFEU requires qualified majority support in the Council. Collective trade policy does not require *unanimity*. Instead it only requires preventing a *blocking minority*. Under current QMV rules (Article 238(3)(a) TFEU), a blocking minority requires at least 4 member states representing at least 35% of the EU-27 population.[6]

[6]The population threshold is calculated over the full EU-27 (approximately 449 million), not just the Eurozone-20 (349 million). This matters because non-Eurozone EU members (Poland, Czech Republic, Romania, Hungary, Sweden, Bulgaria, Denmark; combined population approximately 100 million) also

This shifts the policy question from "how much must we spend to make every member whole?" to "how much must we spend to ensure that no blocking minority of dissatisfied members can form?" Formally, let $\mathcal{U} \subseteq \{1, \ldots, N_E\}$ denote the set of unsatisfied members (those with positive residual gaps after any transfers). A blocking minority exists if and only if:

$$|\mathcal{U}| \geq 4 \quad \text{and} \quad \sum_{n \in \mathcal{U}} \frac{\text{pop}_n}{\text{pop}_{\text{EU27}}} \geq 0.35. \tag{31}$$

The minimum-cost fiscal capacity to prevent any blocking minority is the solution to the optimisation problem

$$\bar{F}^{\text{QMV}} = \min_{\mathcal{C} \subseteq \mathcal{D}} \sum_{n \in \mathcal{C}} \Delta_n \cdot \omega_n \quad \text{s.t.} \quad |\mathcal{D} \setminus \mathcal{C}| < 4 \;\text{ or }\; \sum_{n \in \mathcal{D} \setminus \mathcal{C}} \frac{\text{pop}_n}{\text{pop}_{\text{EU27}}} < 0.35, \tag{32}$$

where $\mathcal{D} = \{n : \Delta_n > 0\}$ is the set of members that prefer their bilateral outside option (defectors) and $\mathcal{C}$ is the set of compensated members. The problem asks: what is the cheapest subset to compensate such that the remaining defectors cannot form a blocking minority? I solve this by exact enumeration over the defector set ($N_E = 20$ makes this tractable).

The key insight is that the population constraint is binding, not the state count constraint. In the US tariff scenario, 19 of 20 members prefer their bilateral outside option, representing 75.4% of EU-27 population. Because defector count (19) far exceeds the threshold (4), reducing the count is expensive, as one must compensate 16 members. But reducing population below 35% requires compensating only the *largest* defectors. Specifically, compensating just three members (Germany (18.8% of EU-27 pop), Spain (10.7%), and Italy (13.1%)), brings the remaining 16 defectors' population share to 32.8%, below the 35% threshold. The cost is $\bar{F}^{\text{QMV}} = 0.33\%$ of Eurozone GDP (€46 billion)—less than half the full compensation benchmark.

The mechanism is intuitive: Germany, Spain, and Italy are "pivotal" in the QMV sense. Without them, no coalition of the remaining 16 members can muster the population share needed to block. The EU need not make *every* member whole; it need only ensure that no blocking coalition forms, and this is determined by a handful of large members.

7.3. Combined Practical Fiscal Requirement

The most realistic estimate applies both mechanisms sequentially. First, cross-conditionality reduces the gaps of net recipient members. Then, QMV blocking prevention is applied to

participate in Council votes on trade policy.

Table 11: *Minimum Fiscal Capacity Under Alternative Institutional Assumptions*

	US Tariff		China Supply		Combined	
	$\bar{F}$	€bn	$\bar{F}$	€bn	$\bar{F}$	€bn
Full compensation	0.69%	97	0.44%	62	1.12%	157
After cross-conditionality	0.66%	92	0.40%	56	1.06%	149
QMV blocking prevention	0.33%	46	0.25%	35	0.60%	84
Cross-cond. + QMV blocking	0.18%	26	0.13%	19	0.33%	46

Notes: Full compensation is $\bar{F}$ from Section 6. After cross-conditionality applies Equation (30) using net EU operating budgetary balances. QMV blocking prevention solves (32) over baseline gaps. Cross-cond. + QMV applies (33): cross-conditionality first, then QMV blocking on residual gaps. Approximate euro amounts assume EZ GDP of €14 trillion.

the *residual* gaps:

$$\bar{F}^{\text{practical}} = \min_{\mathcal{C} \subseteq \mathcal{D}^{\text{XC}}} \sum_{n \in \mathcal{C}} \Delta_n^{\text{XC}} \cdot \omega_n \quad \text{s.t. blocking minority condition (31) fails for } \mathcal{D}^{\text{XC}} \setminus \mathcal{C}, \tag{33}$$

where $\mathcal{D}^{\text{XC}} = \{n : \Delta_n^{\text{XC}} > 0\}$ is the set of members with positive residual gaps after cross-conditionality. Table 11 reports all four estimates across the three shock scenarios.

Three results stand out. First, cross-conditionality alone has limited impact (5–10% reduction) because the members driving $\bar{F}$ are net contributors. Second, QMV blocking prevention reduces $\bar{F}$ by 43–52% because it targets compensation at the few large-population members whose departure would break the blocking threshold. Third, combining both mechanisms yields the largest reduction—74% for the US tariff scenario, from 0.69% to 0.18% (€26 billion). For the combined US-China shock, the practical requirement is $\bar{F}^{\text{practical}} = 0.33\%$ of EZ GDP (€46 billion), roughly 30% of the headline figure and well within the range of existing EU fiscal instruments.

The practical estimate also identifies which members must be compensated. In the combined US-China scenario under the most conservative (cross-cond. + QMV) approach, the critical members are France, Italy, and Spain—the three largest Eurozone economies after Germany. Germany is compensated under QMV-only but not under the combined approach because, as the largest net contributor, it has no cross-conditionality exposure.

8. Conclusion

I have asked a simple question: what minimum fiscal capacity does the Eurozone need to sustain collective trade policy in a world of great power coercion? The answer, derived from a multi-sector trade model calibrated to 20 individual Eurozone members, yields quantitative benchmarks. An upper bound of around 0.7% of Eurozone GDP is needed for a

US tariff escalation, 0.4% for a Chinese supply restriction, and 1.1% for both simultaneously. When grounded in existing EU institutions including cross-conditionality with EU fiscal flows and qualified majority voting rules, the practical requirement falls to 0.18% for the US tariff, 0.13% for China, and 0.33% for both combined (€46 billion).

These bounds reflect different normative policy objectives. The upper bound is the cost of making collective trade policy Pareto-improving across member states, while the lower bound is the cost of preventing institutional breakdown under existing rules. The appropriate target depends on whether this fiscal capacity is aimed at broad-based political legitimacy or as a practical tool for sustaining needed votes.

Five structural findings emerge from the analysis. Single market deepening and fiscal capacity are complements, not substitutes: reducing intra-EU trade barriers generates large welfare gains but barely changes $\bar{F}$. Budget-balanced redistribution fails when all members prefer bilateral deals, because incentive compatibility requires net resource injection, hence the need for joint borrowing capacity. Nominal rigidities change the *distribution* of fiscal need (Ireland's gap nearly triples under sticky prices) but not its aggregate size, making this a facility design problem rather than a sizing problem. Retaliation makes the EU statically worse off, yet the facility's value lies in deterrence: credible retaliatory capacity changes the hegemon's equilibrium tariff choice, and illustrative bloc bargaining scenarios show $\bar{F}$ falling by more than half. Finally, existing EU institutions (cross-conditionality and QMV blocking prevention) reduce the practical requirement to approximately 30% of the headline figure.

Several limitations should be noted. First, the baseline assumes full tariff exemption and zero defection penalty, making the estimates conservative upper bounds. The endogenous bilateral offer analysis shows that defection is often driven by avoidance of the EU's own retaliation rather than a hegemon concession, and even a modest 10% defection penalty reduces $\bar{F}$ by 25%. A richer bargaining setup with: negotiation frictions, uncertainty, and coalition-wide strategic effects, remains an important extension. Second, defection is modelled as a binary choice; in practice, partial non-compliance is possible, with ambiguous effects on $\bar{F}$. Third, the WIOD trade data for Ireland and Luxembourg are distorted by multinational profit shifting, though $\bar{F}$ is robust to excluding these members (-14%). Additional modelling limitations include uniform NTB calibration, linearised NK dynamics, representative-agent welfare (which abstracts from concentrated sectoral losses that drive trade policy lobbying), and the static nature of the exact hat algebra (see Caliendo, Dvorkin, and Parro, 2019, for dynamic extensions).

The framework could be extended in several directions. Incorporating household heterogeneity would connect fiscal capacity to political sustainability at the voter level. The defection decision would depend on the distribution of losses within each member, not just their aggregate. Endogenising the bilateral deal, by making the hegemon's offer depend on the EU's credible threat and its strategic value of breaking solidarity, would close the model as a full bargaining game. The repeated-game analysis in Section 3.3 takes a first step, but a richer dynamic model with endogenous punishment severity and coalition stability would further strengthen the framework.

Crucially, the facility need not be activated to generate value: its existence changes the hegemon's calculus, making coercive tariffs less attractive when credible retaliation is on the table. The fiscal capacity is the size of the insurance policy, not its annual premium. The EU need only borrow if the shock materialises. A standing facility, pre-committed and automatically activated upon qualifying shocks, is preferable to contingent ad hoc arrangements: it makes the EU's collective stance credible ex ante.

As great power competition intensifies and middle powers are pressed to choose sides, the EU's capacity to sustain a collective trade stance under pressure determines its strategic relevance. Fiscal integration is the price of that credibility.

References

Adão, R., A. Costinot, and D. Donaldson (2017). Nonparametric counterfactual predictions in neoclassical models of international trade. *American Economic Review 107*(3), 633–689.

Adao, R., A. Costinot, and D. Donaldson (2025). Putting quantitative models to the test: An application to Trump's trade war. *Quarterly Journal of Economics 140*(2), 1471–1524.

Amiti, M., S. J. Redding, and D. E. Weinstein (2019). The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives 33*(4), 187–210.

Aniceto, P. F. (2023). cp2015: Caliendo and Parro (2015) quantitative trade model in R. R package, replication data from CP2015 supplementary materials.

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012). New trade models, same old gains? *American Economic Review 102*(1), 94–130.

Asdrubali, P., B. E. Sørensen, and O. Yosha (1996). Channels of interstate risk sharing: United States 1963–1990. *Quarterly Journal of Economics 111*(4), 1081–1110.

Bagwell, K. and R. W. Staiger (1999). An economic theory of GATT. *American Economic Review 89*(1), 215–248.

Becko, J. S., G. M. Grossman, and E. Helpman (2025). Optimal tariffs with geopolitical alignment. Working Paper 34108, National Bureau of Economic Research.

Becko, J. S. and G. O'Connor (2025). Hegemonic trade policies with carrots and sticks. Working Paper.

Bonadio, B., Z. Huo, A. A. Levchenko, and N. Pandalai-Nayar (2024). The role of supply chains in the US-China decoupling. Working Paper.

Bown, C. P. (2004). On the economic success of GATT/WTO dispute settlement. *Review of Economics and Statistics 86*(3), 811–823.

Broner, F., A. Martin, J. Meyer, C. Trebesch, and J. Zhou (2025). Hegemony and international alignment. *AEA Papers and Proceedings 115*, 593–598.

Caliendo, L., M. Dvorkin, and F. Parro (2019). Trade and labor market dynamics: General equilibrium analysis of the China trade shock. *Econometrica 87*(3), 741–835.

Caliendo, L. and F. Parro (2015). Estimates of the trade and welfare effects of NAFTA. *Review of Economic Studies 82*(1), 1–44.

Clayton, C., M. Maggiori, and J. Schreger (2024). A theory of economic coercion and fragmentation. Working Paper 33309, National Bureau of Economic Research. Also BIS Working Paper No. 1224.

Costinot, A. and A. Rodríguez-Clare (2014). Trade theory with numbers: Quantifying the consequences of globalization. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, pp. 197–261. Elsevier.

Dekle, R., J. Eaton, and S. Kortum (2007). Unbalanced trade. *American Economic Review: Papers and Proceedings 97*(2), 351–355.

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779.

Evers, M. P. (2015). Fiscal federalism and monetary unions: A quantitative assessment. *Journal of International Economics 97*(1), 59–75.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2020). The return to protectionism. *Quarterly Journal of Economics 135*(1), 1–55.

Farhi, E. and I. Werning (2017). Fiscal unions. *American Economic Review 107*(12), 3788–3834.

Farrell, H. and A. L. Newman (2019). Weaponized interdependence: How global economic networks shape state coercion. *International Security 44*(1), 42–79.

Galí, J. and T. Monacelli (2008). Optimal monetary and fiscal policy in a currency union. *Journal of International Economics 76*(1), 116–132.

Gopinath, G., P.-O. Gourinchas, A. F. Presbitero, and P. Topalova (2025). Changing global linkages: A new cold war? *Journal of International Economics 153*.

Head, K. and T. Mayer (2021). The united states of europe: A gravity model evaluation of the four freedoms. *Journal of Economic Perspectives 35*(2), 23–48.

Kenen, P. B. (1969). The theory of optimum currency areas: An eclectic view. In R. A. Mundell and A. K. Swoboda (Eds.), *Monetary Problems of the International Economy*, pp. 41–60. University of Chicago Press.

Liu, E. and D. Y. Yang (2025). International power. Working Paper 34006, National Bureau of Economic Research.

Ossa, R. (2014). Trade wars and trade talks with data. *American Economic Review 104*(12), 4104–4146.

Rodrik, D. (2011). *The Globalization Paradox: Democracy and the Future of the World Economy*. W.W. Norton.

Online Appendix

E Pluribus Euro: Minimum Fiscal Capacity for Collective Trade Policy in a Currency Union

Joe Kopecky

A. Proofs

A.1. Proof of Proposition 1 (Defection Threshold)

From equations (1)–(2) in the main text, the defection gap for member A is:

$$\begin{aligned}\Delta_A &= W_A^{\text{bilat}} - W_A^{\text{EU}} \\ &= \left(\pi_{AH}^s \cdot \frac{\bar{\tau}^2 - \tau_{\text{bilat}}^2 - 2(\bar{\tau} - \tau_{\text{bilat}})}{2\theta^s} + \pi_{AH}^r \cdot \frac{\bar{\tau}^2 - \tau_{\text{bilat}}^2 - 2(\bar{\tau} - \tau_{\text{bilat}})}{2\theta^r} \right) \cdot \Phi(\xi) \\ &\quad - \delta_{\text{pen}} \cdot G_{AB}(d_{AB}) \end{aligned} \tag{A.1}$$

where $\Phi(\xi) = 1 + \xi/(1-\xi)$ is the NK amplification factor under the common policy, using the identity $(\tau - 1)^2 = \tau^2 - 2\tau + 1$.

The tariff savings term simplifies to:

$$\frac{\bar{\tau}^2 - \tau_{\text{bilat}}^2 - 2(\bar{\tau} - \tau_{\text{bilat}})}{2\theta^j} = \frac{(\bar{\tau} - \tau_{\text{bilat}})(\bar{\tau} + \tau_{\text{bilat}} - 2)}{2\theta^j}$$

This is strictly positive whenever $\bar{\tau} > \tau_{\text{bilat}} \geq 1$.

Setting $\Delta_A = 0$ defines $\bar{\pi}$ implicitly. To obtain an explicit threshold, assume the two-sector exposure is proportional: $\pi_{AH}^r = \rho \cdot \pi_{AH}^s$ for some fixed ratio $\rho > 0$. Then $\Delta_A = 0$ yields:

$$\bar{\pi} = \frac{\delta_{\text{pen}} \cdot G_{AB}(d_{AB})}{\frac{(\bar{\tau} - \tau_{\text{bilat}})(\bar{\tau} + \tau_{\text{bilat}} - 2)}{2} \left(\frac{1}{\theta^s} + \frac{\rho}{\theta^r} \right) \cdot \Phi(\xi)} \tag{A.2}$$

Member A defects whenever $\pi_{AH}^s > \bar{\pi}$. The minimum transfer $\bar{F}_A = \max\{\Delta_A, 0\}$ equals the gap when positive and zero otherwise. ■

A.2. Proof of Proposition 2 (SM–Fiscal Complement)

Differentiating the gap (A.1) with respect to d_{AB}:

$$\frac{\partial \Delta_A}{\partial d_{AB}} = -\delta_{\text{pen}} \cdot G'_{AB}(d_{AB})$$

The tariff savings term is independent of d_{AB} because it depends only on π^j_{AH}, θ^j, $\bar{\tau}$, and τ_{bilat}.

Single market gains satisfy $G'_{AB}(d_{AB}) < 0$ (lower trade costs increase gains). Therefore $\partial \Delta_A / \partial d_{AB} > 0$: a shallower single market (higher d_{AB}) increases the gap.

However, the magnitude is controlled by δ_{pen}:

$$\left| \frac{\partial \Delta_A}{\partial d_{AB}} \right| = \delta_{\text{pen}} \cdot |G'_{AB}(d_{AB})|$$

For $\delta_{\text{pen}} = 0$ (baseline), $\partial \Delta_A / \partial d_{AB} = 0$ exactly: single market depth has zero effect on the defection gap. Even for $\delta_{\text{pen}} = 0.3$, the effect is small because $|G'_{AB}|$ is modest at realistic trade costs ($d_{AB} \approx 1.20$).

The quantitative exercise confirms this: reducing intra-EU NTBs by 25% barely changes $\bar{F}$ (from 0.69% to ≈0.70%), while a 100% NTB elimination produces only marginal reduction. Single market depth and fiscal capacity are complements in the sense that both are needed: deepening the market does not substitute for fiscal transfers. ■

A.3. Proof of Proposition 3 and Corollary (NK Amplification)

Proposition 3. Under Calvo pricing with parameter ξ in a currency union where the central bank targets the GDP-weighted average, each member's welfare deviation from baseline satisfies the aggregate–idiosyncratic decomposition:

$$\hat{W}_n^{NK} - 1 = \bar{W}^{\text{flex}} + \frac{1}{1-\xi}\,\epsilon_n^{\text{flex}} \tag{A.3}$$

where $\bar{W}^{\text{flex}} = \sum_m \omega_m(\hat{W}_m^{\text{flex}} - 1)$ is the GDP-weighted aggregate welfare change under flexible prices, and $\epsilon_n^{\text{flex}} = (\hat{W}_n^{\text{flex}} - 1) - \bar{W}^{\text{flex}}$ is member n's idiosyncratic deviation from the aggregate.

The NK transformation is *affine*, not multiplicative, in each member's flexible-price welfare change: the aggregate component is unchanged (the ECB stabilises the mean), while the idiosyncratic component is amplified by $1/(1-\xi)$. Members whose shocks deviate most from the GDP-weighted average—small, exposed members like Ireland—experience the largest amplification.[7]

The defection gap under NK pricing is:

$$\Delta_n^{NK} = (\hat{W}_n^{\text{bilat,NK}} - 1) - (\hat{W}_n^{\text{EU,NK}} - 1)$$
$$= \frac{1}{1-\xi}\,\Delta_n^{\text{flex}} - \frac{\xi}{1-\xi}\,\bar{\Delta}^{\text{flex}} \tag{A.4}$$

where $\bar{\Delta}^{\text{flex}} = \bar{W}^{\text{bilat,flex}} - \bar{W}^{\text{EU,flex}}$ is the GDP-weighted aggregate gap under flexible prices. The NK gap is a linear combination of the member's own flexible-price gap (amplified by $1/(1-\xi)$) and the aggregate gap (subtracted, because the ECB partially offsets common shocks). For members with gaps larger than the aggregate, NK amplifies the defection incentive; for members close to the mean, it may reduce it.

For $\xi > 0$ and heterogeneous members, the non-uniform amplification across members implies $\bar{F}^{NK} \geq \bar{F}^{\text{flex}}$ (with equality when members are identical, since then $\epsilon_n = 0$ for all n).

Corollary (Mean-Preservation). The GDP-weighted aggregate satisfies:

$$\sum_n \omega_n \Delta_n^{NK} = \frac{1}{1-\xi}\sum_n \omega_n \Delta_n^{\text{flex}} - \frac{\xi}{1-\xi}\,\bar{\Delta}^{\text{flex}} \underbrace{\sum_n \omega_n}_{=1}$$
$$= \frac{1}{1-\xi}\,\bar{\Delta}^{\text{flex}} - \frac{\xi}{1-\xi}\,\bar{\Delta}^{\text{flex}} = \bar{\Delta}^{\text{flex}} = \sum_n \omega_n \Delta_n^{\text{flex}} \tag{A.5}$$

The GDP-weighted sum of NK gaps equals the GDP-weighted sum of flexible-price gaps. Since all gaps are positive (all members prefer bilateral deals), the positive-part operator $[\cdot]^+$ is redundant: $\bar{F}^{NK} = \bar{F}^{\text{flex}}$. NK frictions change *who* needs transfers, not *how much* the total facility costs. ■

A.4. Proof of Proposition 4 (Repeated Game)

Consider an infinitely repeated game with discount factor $\beta \in (0,1)$. In each period t, member n chooses to *cooperate* (maintain the common tariff, receive transfer F_n) or *defect* (bilateral deal with the hegemon). The punishment is grim-trigger exclusion: permanent loss of EU membership with per-period cost $V_n = W_n^{\text{EU}} - W_n^{\text{excluded}} > 0$.

Cooperation payoff (per period): $W_n^{\text{EU}} + F_n$.

Deviation payoff: W_n^{bilat} in the period of defection, then W_n^{excluded} forever.

IC constraint: Member n cooperates if the continuation value of cooperation exceeds

[7]The simplified analytical model (Equation 1) uses the full amplification factor $1/(1-\xi)$ for member A's tariff cost. This is the limiting case $\omega_A \to 0$, where A's shock is entirely idiosyncratic relative to a large union partner B. The general multi-member decomposition (A.3) nests this as a special case.

the deviation payoff:

$$\begin{aligned}\frac{W_n^{\text{EU}} + F_n}{1-\beta} &\geq W_n^{\text{bilat}} + \frac{\beta}{1-\beta} W_n^{\text{excluded}} \\ W_n^{\text{EU}} + F_n &\geq (1-\beta) W_n^{\text{bilat}} + \beta W_n^{\text{excluded}} \\ F_n &\geq (1-\beta)\Delta_n - \beta V_n \end{aligned} \tag{A.6}$$

where $\Delta_n = W_n^{\text{bilat}} - W_n^{\text{EU}}$ and $V_n = W_n^{\text{EU}} - W_n^{\text{excluded}}$. The right-hand side is $(1 - \beta)\Delta_n - \beta V_n = \Delta_n - \beta(\Delta_n + V_n)$, which is the one-shot deviation gain discounted by the shadow cost of future punishment.

Setting $F_n = 0$ and solving for β:

$$\beta(\Delta_n + V_n) \geq \Delta_n \implies \beta \geq \frac{\Delta_n}{\Delta_n + V_n} = \beta_n^*$$

When $\beta < \beta_n^*$, the minimum per-period transfer is $\bar{F}_n(\beta) = (1-\beta)\Delta_n - \beta V_n > 0$. When $\beta \geq \beta_n^*$, no transfer is needed: $\bar{F}_n(\beta) = 0$.

The aggregate fiscal capacity in the repeated game:

$$\bar{F}(\beta) = \sum_n \omega_n \max\{(1-\beta)\Delta_n - \beta V_n,\ 0\}$$

This interpolates smoothly from $\bar{F}(0) = \sum_n \omega_n \Delta_n^+$ (the one-shot case) to $\bar{F}(\beta^*_{\max}) = 0$, where $\beta^*_{\max} = \max_n \beta_n^*$. ■

B. Data Construction

B.1. Source Data

All trade and production data come from the World Input-Output Database (WIOD), 2016 Release, using the 2014 World Input-Output Table (`WIOT2014_Nov16_ROW.xlsb`). WIOD covers 44 countries (43 individually identified + rest-of-world aggregate) and 56 sectors (ISIC Rev. 4).

B.2. Country Aggregation

Table B1: *Country Structure*

Index	Country	ISO code	GDP weight (%)
0	Austria	AUT	3.2
1	Belgium	BEL	4.4
2	Cyprus	CYP	0.2
3	Germany	DEU	28.2
4	Spain	ESP	10.3
5	Estonia	EST	0.2
6	Finland	FIN	2.1
7	France	FRA	20.0
8	Greece	GRC	1.5
9	Croatia	HRV	0.4
10	Ireland	IRL	2.0
11	Italy	ITA	16.3
12	Lithuania	LTU	0.3
13	Luxembourg	LUX	0.8
14	Latvia	LVA	0.3
15	Malta	MLT	0.1
16	Netherlands	NLD	6.7
17	Portugal	PRT	1.7
18	Slovakia	SVK	0.9
19	Slovenia	SVN	0.4
20	United States	USA	—
21	China	CHN	—
22	Rest of World	RoW	—

The 20 Eurozone members correspond to the 2014 membership. WIOD countries not in the Eurozone are aggregated: the United States and China are kept as individual economies; all others (Australia, Brazil, Canada, etc.) are folded into the rest-of-world aggregate.

GDP weights are constructed from WIOD gross output data (sum of value added across all sectors for each country).

B.3. Sector Aggregation

WIOD's 56 ISIC Rev. 4 sectors are aggregated to 10 model sectors:

Table B2: *Sector Mapping: WIOD to Model*

Index	Model sector	WIOD ISIC codes	Trade elasticity θ^j
0	Critical minerals	B	3.0
1	Energy	C19, D35	5.0
2	Semiconductors/electronics	C26-C27	8.0
3	Pharma/chemicals	C20-C21	4.0
4	Agriculture/food	A, C10-C12	2.5
5	Auto/transport	C29-C30	4.5
6	Machinery	C28	5.5
7	Financial/business services	K65-K66, L68, M71-N	6.0
8	Digital services & ICT	J58-J63, K64	7.0
9	Other manufactures & services	residual (30 WIOD sectors)	4.0

Trade elasticities are taken from Caliendo and Parro (2015) Table 1 (production-side estimates). Where multiple CP sectors are aggregated into one model sector, the elasticity is the trade-weighted average.

B.4. Calibration Parameters

For each country–sector pair (n, j), the following are computed from WIOD:

- **Trade shares** π^j_{ni}: expenditure of country n on goods from country i in sector j, as a fraction of total expenditure on sector j goods. Computed from intermediate use flows plus final demand.
- **Value-added shares** $\gamma^{VA,j}_n$: value added as a fraction of gross output in sector j for country n.
- **Input-output coefficients** $\gamma^{k,j}_n$: intermediate inputs from sector k per unit of gross output in sector j for country n. Cost shares satisfy $\gamma^{VA,j}_n + \sum_k \gamma^{k,j}_n = 1$.
- **Consumption shares** α^j_n: final demand for sector j as a fraction of total final demand for country n.
- **Labour** L_n: proportional to total gross output (normalised to maximum = 1).

C. Solver Properties

C.1. Algorithm

The hat algebra system is solved by fixed-point iteration. Given the shock $\hat{\tau}$, the solver iterates on wages $\hat{w}_n$ and prices $\hat{P}^j_n$:

1. Initialise $\hat{w}_n = 1$ for all n.

2. Compute unit costs: $\hat{c}_n^j = (\hat{w}_n)^{\gamma_n^{VA,j}} \prod_k (\hat{P}_n^k)^{\gamma_n^{k,j}}$.
3. Update trade shares: $\hat{\pi}_{ni}^j = (\hat{\tau}_{ni}^j \hat{c}_i^j)^{-\theta^j} / \sum_{i'} \pi_{ni'}^j (\hat{\tau}_{ni'}^j \hat{c}_{i'}^j)^{-\theta^j}$.
4. Update price indices: $\hat{P}_n^j = \left[\sum_i \pi_{ni}^j (\hat{\tau}_{ni}^j \hat{c}_i^j)^{-\theta^j}\right]^{-1/\theta^j}$.
5. Update wages from the trade balance condition.
6. Dampened update: $\hat{w}_n^{\text{new}} = (1-\lambda)\hat{w}_n^{\text{old}} + \lambda \hat{w}_n^{\text{computed}}$, with $\lambda = 0.2$.
7. Repeat until $\max_n |\hat{w}_n^{\text{new}} - \hat{w}_n^{\text{old}}| < \epsilon$.

C.2. Convergence

The solver uses tolerance $\epsilon = 10^{-10}$. Typical iteration counts:

Table C3: *Solver Convergence*

Exercise	Iterations	Final residual
US 20% tariff (23 countries)	∼800	$< 10^{-10}$
China minerals (23 countries)	∼600	$< 10^{-10}$
NAFTA (27 countries)	∼500	$< 10^{-10}$
Bilateral outside option (per member)	∼400	$< 10^{-10}$

Results are invariant to tolerance: reducing ϵ from 10^{-8} to 10^{-12} changes welfare estimates by less than 10^{-8} percentage points. The RoW numeraire convention does not affect real outcomes.

D. Full Results Tables

D.1. All Members: US Tariff Scenario

Table D4: *Complete Results: US 20% Tariff + EU 20% Retaliation*

Member	$\hat{W}^{EU} - 1$ (%)	$\hat{W}^{bilat} - 1$ (%)	Gap Δ_n (pp)	GDP-wt contr. (pp)
AUT	−0.48	+0.14	0.62	0.020
BEL	−0.98	+0.19	1.16	0.052
CYP	−0.04	−0.00	0.04	0.000
DEU	−0.71	+0.08	0.79	0.224
ESP	−0.25	+0.04	0.30	0.030
EST	−0.21	+0.05	0.26	0.001
FIN	−0.53	+0.04	0.57	0.012
FRA	−0.50	+0.02	0.52	0.105
GRC	−0.02	−0.07	−0.05	0.000
HRV	−0.25	+0.06	0.31	0.001
IRL	−3.33	+0.72	4.06	0.083
ITA	−0.41	+0.05	0.45	0.074
LTU	−0.38	−0.01	0.37	0.001
LUX	−2.54	+1.23	3.77	0.032
LVA	−0.10	+0.02	0.12	0.000
MLT	−0.69	−0.17	0.52	0.001
NLD	−0.71	+0.07	0.78	0.052
PRT	−0.15	+0.01	0.15	0.003
SVK	−0.23	+0.08	0.31	0.003
SVN	−0.15	+0.11	0.26	0.001
Aggregate $\bar{F}$ (% EZ GDP):				0.69

Notes: Aggregate $\bar{F}$ is computed from unrounded values (0.6928%); displayed contributions are rounded to three decimal places and may not sum exactly to the aggregate.

D.2. All Members: China Shock Scenario

Table D5: *Complete Results: China Minerals + Semiconductors + EU Retaliation*

Member	$\hat{W}^{EU} - 1$ (%)	Gap Δ_n (pp)	GDP-wt contr. (pp)
NLD	−0.97	1.23	0.082
LUX	−0.71	0.82	0.007
EST	−0.63	0.81	0.002
MLT	−0.63	0.77	0.001
CYP	−0.58	0.66	0.001
LTU	−0.45	0.59	0.002
DEU	−0.30	0.48	0.135
FRA	−0.19	0.31	0.062
ITA	−0.14	0.24	0.040
Aggregate $\bar{F}$ (% EZ GDP):			0.44

D.3. Hegemon's Optimal Bilateral Offers

Table D6: *Endogenous Bilateral Offers: Hegemon's Optimal Strategy*

Member	τ^* (%)	Endogenous gap (pp)	US welfare gain (pp)
IRL	20.0	2.53	+0.0332
LUX	20.0	3.73	+0.0075
BEL	20.0	0.53	+0.0348
NLD	20.0	0.49	+0.0565
DEU	20.0	0.21	+0.1017
FRA	20.0	0.10	+0.0858
FIN	20.0	0.12	+0.0079
AUT	20.0	0.13	+0.0066
ITA	20.0	0.07	+0.0293
ESP	20.0	0.07	+0.0170
GRC	n/a	−0.05	n/a
Endogenous $\bar{F}$ (% EZ GDP):			0.25
Free-trade $\bar{F}$ (% EZ GDP):			0.69

Notes: τ^* is the welfare-maximising US tariff offer to each member, subject to the member accepting. Endogenous gap is the defection gap under the hegemon-optimal offer. Members with $\tau^* = 20\%$ receive no concession; defection is driven by avoidance of EU retaliation costs. Greece (n/a) has a negative gap and is not targeted. See Section 6.1.

D.4. Cross-Exercise Summary

Table D7: *Cross-Exercise Comparison of Minimum Fiscal Capacity*

Scenario	$\bar{F}$ (% EZ GDP)	≈€bn	Most exposed
US 20% + retaliation (free bilat.)	0.69	97	IRL, LUX
US 20% + retaliation (endo. bilat.)	0.25	35	LUX, IRL
Bloc bargaining (10%, no retal.)	0.31	43	IRL, LUX
China minerals + retaliation	0.44	62	NLD, DEU
Combined shock (worst case)	1.12	157	IRL, LUX, NLD
No-retaliation (US tariff only)	0.47	66	IRL, LUX
Conservative lower bound (headline ÷3; see Section 5.4):			
US 20% + retaliation (adjusted)	0.23	32	IRL, LUX
Combined shock (adjusted)	0.37	52	IRL, LUX, NLD

Notes: "Free bilat." assumes full tariff exemption as the outside option; "endo. bilat." uses the hegemon's welfare-maximising offer (Table D6). Euro amounts use EZ GDP ≈ €14 trillion. Conservative lower bound divides headline estimates by 3 to account for potential data-vintage bias (Section 5.5).

E. Additional Robustness

E.1. Trade Elasticity Sensitivity

Table E8: *Sensitivity to Trade Elasticities*

θ multiplier	Description	$\bar{F}$ (% EZ GDP)
0.5×	Half baseline	0.59
1.0×	Baseline	0.69
1.5×	50% higher	0.79

The range [0.59%, 0.79%] confirms that results are robust to the choice of trade elasticities. Higher elasticities produce modestly larger $\bar{F}$ because they strengthen substitution and trade diversion effects, making the bilateral outside option relatively more attractive and widening defection gaps.

E.2. Bilateral Deal Quality

Table E9: *Sensitivity to Bilateral Tariff Offer*

τ_{bilat}	$\bar{F}$ (% EZ GDP)	Defectors
0%	0.69	19/20
5%	0.53	19/20
10%	0.41	19/20
15%	0.32	18/20
18%	0.27	17/20

E.3. Defection Penalty

Table E10: *Sensitivity to Defection Penalty*

SM gains lost	$\bar{F}$ (% EZ GDP)	Defectors
0%	0.69	19/20
10%	0.52	16/20
20%	0.35	12/20
30%	0.18	6/20
50%	0.00	5/20

E.4. NK Transition Dynamics

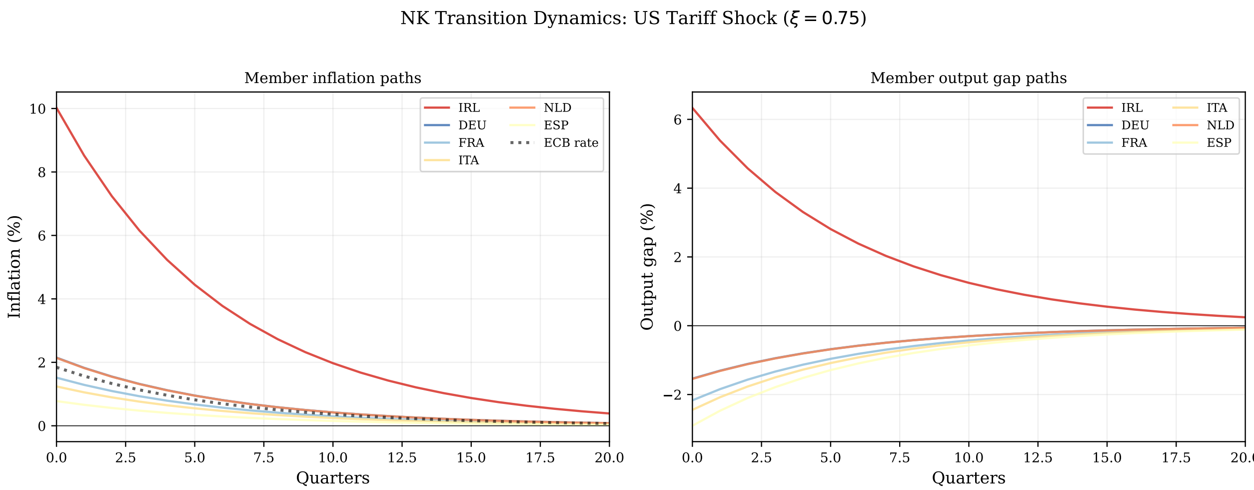

Figure E1: *NK Transition Dynamics (US Tariff Scenario, $\xi = 0.75$)*
Notes: Left panel: inflation paths for six key members and the ECB aggregate rate. Right panel: output gap paths. Ireland is a massive outlier ($\pi_0 = +11.4\%$, $y_0 = +7.1\%$). Half-life approximately 4.3 quarters.