Global Reallocation of Capital in the Era of Geo-Economic Fragmentation and Artificial Intelligence

Fernando Toledo* Gabriel Montes-Rojas†

May 2026

Abstract

This paper studies how geo-economic fragmentation affects the international allocation of capital in an economy where artificial intelligence (AI) relies on specialized, mobile investment. We develop a heterogeneous-agent open-economy model with AI-specific capital and geopolitical frictions in cross-border asset holdings. These frictions reduce capital mobility, raise the cost of AI adoption, and generate distributional effects across countries and households. The quantitative results show that fragmentation produces welfare losses that rise nonlinearly with geopolitical distance. These losses are larger for AI-lagging economies and for low-wealth households, mainly through reduced investment, weaker labor income, and tighter financial conditions. The paper also defines a so-called Fragmented-AI Trilemma: economies cannot simultaneously preserve monetary autonomy, maintain efficient AI-capital allocation, and insulate themselves from geopolitical shocks.

JEL Classification: E12, E52, F21, F41, O33

Keywords: Geo-economic fragmentation, artificial intelligence, HANK, Simulated Method of Moments

*Universidad Nacional de La Plata, La Plata, Argentina and Universidad Argentina de la Empresa (UADE), Instituto de Economia (INECO), Ciudad Autónoma de Buenos Aires, Argentina. fernando.toledo@econo.unlp.edu.ar

†Instituto Interdisciplinario de Economía Política, Universidad de Buenos Aires-CONICET, and Universidad Argentina de la Empresa (UADE), Instituto de Economia (INECO), Ciudad Autónoma de Buenos Aires, Argentina. gabriel.montes@economicas.uba.ar

1 Introduction

Geopolitical alignment has become an increasingly important force in international capital markets. Foreign direct investment and portfolio flows increasingly reflect political distance as well as expected returns, while investment screening, technology restrictions, and supply-chain reorientation have made cross-border investment more conditional on strategic relationships. These changes are unfolding at the same time as artificial intelligence (AI) requires large, specialized, and internationally mobile forms of capital, including data centers, semiconductor capacity, cloud infrastructure, and frontier-model investment. The result is a growing tension: the global diffusion of AI depends on capital mobility, but that mobility is increasingly constrained by geo-economic fragmentation.

The interaction between fragmentation and AI adoption is especially important because the two forces can reinforce each other. AI investment depends on access to specialized capital, infrastructure, and supply chains that are concentrated in a limited set of countries. At the same time, geopolitical tensions make these cross-border flows more costly and less predictable. This creates a particular problem for economies that are already behind the AI frontier: they need foreign capital and technology to catch up, but they are also more exposed to the barriers created by fragmentation. Fragmentation therefore does not simply reduce capital mobility; it can also widen differences in AI adoption, productivity, and welfare across countries.

Several strands of literature are directly relevant but remain incomplete for the question studied here. The geo-economic fragmentation literature (Aiyar et al., 2024; Alfaro and Chor, 2023) documents the reallocation of trade, investment, and supply chains across geopolitical lines, including bloc formation and the role of connector economies. However, this literature has provided less structural welfare analysis of how fragmentation affects the composition and allocation of capital flows. The AI-and-macroeconomics literature (Acemoglu and Restrepo, 2022; Acemoglu, 2025) has developed formal approaches to the macroeconomic effects of AI, but it has largely focused on closed-economy mechanisms and has paid less attention to international capital allocation. The heterogeneous-agent macroeconomics literature (Auclert et al., 2024b; Bayer et al., 2024; Kaplan et al., 2018) provides the tools needed to study distributional welfare effects, but it has not been fully extended to settings in which geopolitical frictions and AI-specific

capital jointly shape exposure to international shocks. Guo et al. (2023) provide a closely related open-economy HANK framework, but their analysis does not model AI capital or endogenous geopolitical capital mobility frictions. This paper brings these elements together.

We develop an open-economy HANK model in which AI-specific capital is an internationally traded input and cross-border financial positions are subject to geopolitical frictions. Production combines traditional capital, labor, and AI capital, while firms may face binding financial constraints that amplify shocks to investment and labor demand. The framework is extended to include connector economies, which can intermediate capital flows between more fragmented blocs. The model is tied to recent evidence on capital-flow realignment, portfolio reallocation, and the different sensitivity of emerging and advanced economies to international financial shocks.

The main mechanism is straightforward. When geopolitical distance increases, cross-border investment becomes more costly. This reduces the flow of capital toward economies where its marginal product is high, raises the cost of AI adoption, and weakens investment. Because AI capital is internationally traded, these effects also operate through exchange rates and imported capital costs. Financial constraints amplify the initial shock by reducing firms' ability to invest and hire. The consequences are therefore both aggregate and distributional: fragmentation lowers welfare, but the losses are not evenly shared across countries or households.

The paper delivers several results. Geo-economic fragmentation generates welfare losses that rise more than proportionally with geopolitical tension. This nonlinearity comes from the interaction between capital mobility frictions, AI-capital accumulation, and occasionally binding financial constraints. The costs are also unevenly distributed. AI-lagging economies are more exposed because they depend more heavily on foreign capital and technology for catch-up. Within countries, low-wealth households bear larger losses because they are more exposed to labor-income risk and have fewer financial buffers. Finally, the paper identifies a Fragmented-AI Trilemma: when capital mobility frictions depend on geopolitical alignment, economies cannot simultaneously preserve monetary autonomy, maintain efficient AI-capital allocation, and insulate themselves from geopolitical shocks. This result extends the logic of the traditional open-economy policy trilemma to a setting in which AI capital and geopolitical fragmentation are central features of

international macroeconomic adjustment.

The rest of the paper proceeds as follows. Section 2 documents the main empirical facts that motivate the model and our paper. Section 3 presents the theoretical framework, describes the calibration and estimation strategy. Section 4 extends the analysis to connector economies. Section 5 shows the quantitative analysis and reports the aggregate and distributional welfare effects of fragmentation, evaluates the model's empirical fit through targeted moments and additional validation checks. Section 6 studies the policy trade-offs in a so-called Fragmented-AI Trilemma. The final section concludes. A detailed Appendix has a self-contained derivation of the model.

2 Stylized Facts and Their Mapping to the Model

Before developing the theoretical framework, it is useful to establish some related literature.

2.1 Geo-Economic Fragmentation Is Multidimensional and Accelerating

The literature shows that geo-economic fragmentation is simply rising trade protectionism. Aiyar et al. (2024), document multiple structurally distinct fragmentation dimensions in cross-border investment data operating simultaneously in bilateral FDI flows: ideological sorting, friendship-based clustering, de-risking, nearshoring, and reshoring. These dimensions reinforce each other through a "fragmentation multiplier" whereby any one dimension activates the others. Aiyar et al. (2024) document that the geopolitical alignment dimension of FDI decisions approximately doubled in economic significance between 2016 and 2022, establishing a structural break dramatically accelerated by Russia's invasion of Ukraine.

2.2 AI-Related Capital Flows Are Extraordinary in Scale and Geographically Concentrated

The AI investment wave represents a qualitative shift in the composition of global capital. UNCTAD (2026) documents the growth of AI-related FDI as a share of global flows (UNCTAD, 2026). Data-center capital expenditures grew rapidly in 2024–2025 (industry estimates). The geographic distribution of AI-related FDI

closely mirrors the U.S.-aligned geopolitical bloc: the United States receives the dominant share of sovereign wealth fund AI investments (International Monetary Fund, 2025), and the semiconductor supply chains essential to AI hardware fabrication are concentrated in Taiwan, South Korea, the Netherlands, and Japan.

2.3 Connector Economies Are Gaining Systematically from Investment Diversion

Alfaro and Chor (2023), using detailed U.S. outward FDI and import-sourcing microdata, document a decisive rotation of U.S. supply chains toward Vietnam, Mexico, and India. Aiyar et al. (2024), using structural gravity models, document significant declines in cross-bloc FDI relative to within-bloc flows (Aiyar et al., 2024) in-bloc flows between Q1 2022 and Q4 2024, investment flows to non-aligned economies held at their pre-2022 levels or rose. Aiyar and Ohnsorge (2024) formalize this observation into the concept of "connector status," documenting that connector benefits peak at intermediate fragmentation levels and erode as fragmentation deepens toward complete bloc decoupling.

2.4 Fragmentation Severely Distorts AI-Related Trade and Capital Flows

International Monetary Fund (2025) document that AI-related goods accounted for a growing share of global merchandise trade in 2024–2025 regarding the total trade value—a finding independently consistent with OECD and WTO merchandise trade statistics for H1 2025. The AI hardware supply chain is characterized by extraordinary geographic concentration: advanced AI chip fabrication is dominated by TSMC in Taiwan and Samsung in South Korea, with advanced packaging almost entirely conducted by ASML-equipped facilities in the Netherlands. A geopolitical shock disrupting the U.S.–Taiwan relationship is simultaneously a shock to global AI capital formation.

2.5 Macroeconomic Costs Are Large, Nonlinear, and Systematically Asymmetric

The IMF (2025) estimates 0.4–1 percent global output losses by 2027. Broner et al. (2013), using panel data on 41 countries over 25 years, document that capital flows to emerging-market destinations respond approximately four times more sensi-

tively to geopolitical risk than flows to advanced-economy destinations. Cerutti et al. (2025) find that AI adoption widens productivity gaps across countries, with emerging economies falling further behind as AI-frontier economies capture disproportionate productivity gains.

2.6 The Global Financial Cycle Compounds Domestic Fragmentation Vulnerability

Miranda-Agrippino and Rey (2020) establish that a single global factor—the U.S. monetary policy cycle—accounts for 23 percent of the variance in international capital flows and substantially constrains the monetary autonomy of small open economies. Engel (2016) documents that UIP deviations are highly persistent, with an AR(1) coefficient in the range 0.82–0.90. These two features—the global cycle as a constraint on monetary autonomy and the persistence of UIP deviations—are the empirical counterparts of the model's global financial cycle component (equation (11)).

2.7 From Stylized Facts to Structural Identification: A Methodological Bridge

The six stylized facts in Sections 2.1–2.6 serve a dual purpose in this paper: they motivate the model's structural architecture and simultaneously provide the empirical moments against which its friction parameters are identified. This dual role—motivational and identificational—distinguishes the present approach from standard calibration exercises in international macroeconomics.

The identification chain begins with the aggregate FDI decline (Section 2.1), which pins down the friction level $\bar{\kappa}$ at the steady-state geopolitical distance. It continues with the bilateral portfolio semi-elasticity (Section 2.6, empirical counterpart in Aiyar et al. 2024), which identifies the marginal response γ of the friction to geopolitical distance. The autocorrelation of UIP deviations (Section 2.6, counterpart in Engel 2016) identifies ρ_G from the time-series dimension of the friction-induced UIP wedge. And the emerging-to-advanced sensitivity asymmetry (Section 2.5, counterpart in Broner et al. 2013) serves as the overidentifying restriction, testing whether the model's financial mechanism—parameterized by pre-calibrated collateral constraint tightness $\theta^H = 0.75$ and $\theta^F = 0.50$ from Kase et al. (2025)—reproduces the correct cross-country asymmetry without any additional degree of freedom.

Three features of the reduced-form evidence provide additional discipline beyond the four SMM moments. First, the geographic concentration of AI-related FDI (Section 2.2) validates the asymmetric calibration $\alpha_a^H = 0.20$ versus $\alpha_a^F = 0.10$. Second, the connector economy evidence (Section 2.3) validates the three-bloc extension of Section 4: the model's endogenous prediction that the connector premium peaks at intermediate fragmentation and erodes thereafter (Table 2) replicates the empirical pattern of Aiyar and Ohnsorge (2024) without any targeted calibration. Third, the stagflationary character of fragmentation episodes documented by Aiyar et al. (2024)—simultaneous output contractions and inflation acceleration in 2022–2024—validates the AI-capital exchange-rate channel in equation (8), which distinguishes the present framework from demand- side models of fragmentation in which shocks are contractionary and deflationary rather than stagflationary. The welfare numbers in Table 3 are only as credible as the identification chain that produces them; this section establishes the chain's empirical links explicitly before the model is presented.

3 Model

The theoretical framework is a two-country HANK model that extends Guo et al. (2023) in three dimensions: AI-specific capital as a distinct internationally traded factor, endogenous geopolitical capital mobility frictions, and global solution of an occasionally binding collateral constraint. The world consists of two symmetric large open economies—Home (H) and Foreign (F)—each populated by a unit mass of heterogeneous households, competitive final-goods producers, a continuum of monopolistically competitive intermediate firms, a central bank, and a fiscal authority. Section 4 introduces the three-bloc extension.

The heterogeneous-agent structure is not a modeling convenience but a structural necessity. Kaplan et al. (2018) establish that the direct intertemporal substitution channel accounts for at most 10–20 percent of the aggregate consumption response to a monetary shock; the remaining 80–90 percent operates through indirect general equilibrium effects on labor markets and firm profits. Auclert et al. (2024b) show that the intertemporal marginal propensity to consume matrix is the sufficient statistic for the distributional response to aggregate demand shocks; here, the geopolitical fragmentation shock operates analogously through labor market

equilibrium, making the full cross-sectional wealth distribution essential to the welfare accounting.

Additional details with algebraic derivations and solutions are in the Appendix.

3.1 Households

A unit mass of infinitely lived households indexed by $i \in [0,1]$ faces uninsurable idiosyncratic labor income risk over incomplete asset markets. Household i maximizes:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_{i,t}^{1-\sigma}}{1-\sigma} - \psi \frac{n_{i,t}^{1+\varphi}}{1+\varphi} \right], \tag{1}$$

where $\beta \in (0,1)$ is the discount factor, $c_{i,t}$ is consumption, $n_{i,t}$ is hours worked, $\sigma > 0$ is the coefficient of relative risk aversion, $\psi > 0$ is the disutility of labor, and $\varphi \geq 0$ is the inverse Frisch elasticity of labor supply. The calibration $\sigma = 2$ follows Hall (1988), and $\varphi = 1$ follows Chetty et al. (2011).

Idiosyncratic income process. Household efficiency units of labor follow $\log h_{i,t} = \rho_h \log h_{i,t-1} + \nu_{i,t}$, where $\nu_{i,t} \sim \mathcal{N}(0, \sigma_h^2)$ is i.i.d. across agents and time. The persistence $\rho_h = 0.966$ and variance $\sigma_h = 0.017$ are estimated from Flóden and Lindé (2001) using PSID-equivalent panel data and reproduce the empirical wealth distribution moments—Gini coefficient, share of wealth held by the top decile, and hand-to-mouth fraction—that are essential for correct welfare accounting.

Budget constraint with geopolitical portfolio friction. Household i holds domestic bonds $b_{i,t}^H$ and foreign bonds $b_{i,t}^F$. The budget constraint is:

$$c_{i,t} + b_{i,t}^H + b_{i,t}^F + \underbrace{\frac{\kappa(G_t)}{2}(b_{i,t}^F)^2}_{\text{geopolit. cost}} \leq w_t n_{i,t} h_{i,t} + R_{t-1}^H b_{i,t-1}^H + R_{t-1}^F b_{i,t-1}^F \frac{e_t}{e_{t-1}} + \Pi_t + T_t, \tag{2}$$

where w_t is the real wage, R_{t-1}^H and R_{t-1}^F are gross real returns, e_t is the real exchange rate, Π_t are firm profits distributed as dividends, and T_t is a lump-sum fiscal transfer.

The term $\frac{\kappa(G_t)}{2}(b_{i,t}^F)^2$ is the geopolitical portfolio adjustment cost, giving the Schmitt-Grohé and Uribe (2003) stationarity device a structural geopolitical

interpretation through:

$$\kappa(G_t) = \bar{\kappa}(1 + \gamma G_t), \tag{3}$$

where $G_t \in [0, 1]$ is the bilateral geopolitical distance index of Fernández-Villaverde et al. (2023), $\bar{\kappa} > 0$ is the baseline friction between fully aligned economies ($G_t = 0$), and $\gamma > 0$ is the sensitivity of the friction to geopolitical divergence. The linear form is parsimonious and empirically supported by the gravity-model estimates of Aiyar et al. (2024). Both $\bar{\kappa}$ and γ, together with the shock persistence ρ_G, are estimated by Simulated Method of Moments.

The household's first-order condition for foreign bonds yields:

$$1 + \kappa(G_t) b^F_{i,t} = \beta E_t \left[\left(\frac{c_{i,t+1}}{c_{i,t}} \right)^{-\sigma} R^F_t \frac{e_{t+1}}{e_t} \right]. \tag{4}$$

Welfare metric. Welfare is measured as the consumption-equivalent variation (CEV) following Lucas (1987): the fraction τ^{CEV} of the baseline lifetime consumption stream that makes a household indifferent between the fragmented and frictionless worlds. Formally, $V_i\big((1 + \tau_i^{CEV})\{c^0_{i,t}\}_{t=0}^{\infty}\big) = V_i\big(\{c^G_{i,t}\}_{t=0}^{\infty}\big)$, where $\{c^0_{i,t}\}$ is the baseline and $\{c^G_{i,t}\}$ the fragmented path. Aggregate CEV integrates over the stationary distribution of the zero-fragmentation baseline, adopting the ex-ante welfare criterion of Dávila et al. (2012).

3.2 Production Technology: AI-Augmented Nested CES

Intermediate firm j produces output using a three-way nested CES technology:

$$y_{j,t} = Z_t \left[\alpha_k^{1/\eta} k_{j,t}^{(\eta-1)/\eta} + \alpha_n^{1/\eta} (A_t n_{j,t})^{(\eta-1)/\eta} + \alpha_a^{1/\eta} (\Theta_t a_{j,t})^{(\eta-1)/\eta} \right]^{\eta/(\eta-1)}, \tag{5}$$

where Z_t is neutral TFP, A_t is labor-augmenting productivity, Θ_t is an AI-capital-augmenting technology shifter, $k_{j,t}$ is traditional capital, $n_{j,t}$ is labor, $a_{j,t}$ is AI-specific capital (GPU clusters, data-center infrastructure, proprietary model weights), and $\alpha_k + \alpha_n + \alpha_a = 1$.

The elasticity of substitution $\eta > 0$ is the paper's dominant welfare parameter (sensitivity elasticity $+0.85$). When $\eta < 1$, AI capital and labor are gross complements; when $\eta > 1$, they are gross substitutes. This formulation follows Acemoglu and Restrepo (2022), who estimate $\eta \in (0.5, 0.8)$ for the ICT-labor rela-

tionship in U.S. data 1980–2019. Acemoglu and Restrepo (2022) suggest that η may substantially exceed unity for automatable tasks, motivating the $\eta \in \{0.5, 1.5, 2.5\}$ sensitivity analysis.

Financial frictions: the collateral constraint. Firms finance investment by borrowing against their capital stock, subject to:

$$b_{j,t} \leq \theta k_{j,t}. \tag{6}$$

When the constraint binds ($\mu_{j,t} > 0$), its Lagrange multiplier generates a wedge in firm labor demand:

$$w_t = \frac{\partial F}{\partial n_{j,t}} \cdot \frac{1}{1 + \tilde{\mu}_{j,t}}, \quad \tilde{\mu}_{j,t} \equiv \frac{\mu_{j,t}\theta}{r_t^k} > 0. \tag{7}$$

This wedge—increasing in the collateral tightness and more severe when fragmentation simultaneously depresses AI-capital inflows and firm cash flows—is the central distributional mechanism of the model, following the financial accelerator logic of Bernanke et al. (1999), adapted to the HANK open-economy setting.

3.3 Open-Economy Pricing: The AI-Capital Exchange-Rate Channel

Because AI capital is internationally traded, its Home-currency cost depends on the real exchange rate. Log-linearizing the CES unit cost function and substituting the optimal reset price condition from the Calvo (1983) price-setting problem yields the open-economy New Keynesian Phillips Curve (NKPC):

$$\pi_t = \beta E_t \pi_{t+1} + \kappa_{nk} \widehat{mc}_t^{dom} + \kappa_{nk} s_a \underbrace{\left(\hat{q}_t + \hat{r}_t^{a,*} - \hat{\Theta}_t\right)}_{\text{AI-capital exchange-rate channel}}, \tag{8}$$

where $\hat{q}_t$ is the log-deviation of the real exchange rate, $\hat{r}_t^{a,*}$ is the log-deviation of the foreign rental rate on AI capital, $\hat{\Theta}_t$ is the AI-augmenting technology shock, s_a is the steady-state AI-capital cost share, and $\kappa_{nk} \equiv (1 - \xi_p)(1 - \beta\xi_p)/\xi_p \approx 0.086$. The price stickiness parameter $\xi_p = 0.75$ implies an average price duration of four quarters, consistent with Galí (2015). The AI-capital exchange-rate channel is entirely absent from both closed-economy AI models (Acemoglu, 2025) and standard open-economy HANK models (Guo et al., 2023).

3.4 Capital Accumulation

Traditional and AI capital accumulate with quadratic investment adjustment costs following Christiano et al. (2005):

$$K_{t+1} = (1-\delta_k)K_t + I_{k,t} - \frac{\varphi_k}{2}\left(\frac{I_{k,t}}{K_t} - \delta_k\right)^2 K_t, \tag{9}$$

$$A_{t+1} = (1-\delta_a)A_t + I_{a,t} - \frac{\varphi_a}{2}\left(\frac{I_{a,t}}{A_t} - \delta_a\right)^2 A_t. \tag{10}$$

The depreciation $\delta_a = 0.050$ reflects a weighted average of AI hardware ($\approx 40\%$ annually) and intangible AI assets ($\approx$ 10–15%). The higher adjustment cost $\varphi_a = 8$ (versus $\varphi_k = 4$) reflects the lumpiness of data-center construction.

3.5 International Capital Flows and the Modified UIP

The aggregate UIP condition with geopolitical frictions. Aggregating the household's first-order condition for foreign bond holdings (equation 4) over the stationary wealth distribution and imposing asset market clearing at both the household and aggregate levels yields a richer departure from the textbook uncovered interest parity condition:

$$E_t \Delta e_{t+1} = R_t^H - R_t^F + \underbrace{\kappa(G_t) B_t^F}_{\text{geopolitical friction}} + \underbrace{\phi_{GFC} r_t^{US} + \zeta_t^{idio}}_{\text{global financial cycle}}. \tag{11}$$

Equation (11) is the model's central transmission mechanism and deserves careful economic interpretation, term by term.

The left-hand side, $E_t \Delta e_{t+1}$, is the expected real exchange rate depreciation: the return the Home economy must offer on its currency relative to Foreign in order to equilibrate cross-border portfolio positions.

The first right-hand-side term, $R_t^H - R_t^F$, is the standard interest differential—the uncovered interest parity wedge of Mundell (1963) and Engel (2016). In the absence of the remaining terms, perfect capital mobility would drive this wedge to zero, linking domestic and foreign monetary conditions. The empirical failure of textbook UIP—the "forward premium puzzle" documented by Fama (1984) and recently reinterpreted by Gourinchas and Sammon (2024) in the context of risk-appetite cycles—motivates the two additional channels embedded in equation (11).

The geopolitical friction term $\kappa(G_t)B_t^F$. The second term, $\kappa(G_t)B_t^F$, is the paper's central theoretical innovation. It introduces a *state-dependent* portfolio adjustment cost that operates through two channels simultaneously. The first channel is the *level effect*: for any given net foreign asset position B_t^F, a higher friction coefficient κ requires a larger interest differential to compensate Home households for holding foreign assets, effectively tightening the external financing constraint. The second channel is the *portfolio-rebalancing amplifier*: as geopolitical tension rises and G_t increases, the marginal cost of cross-border investment rises, triggering a portfolio rebalancing from international to domestic assets that feeds back into the exchange rate and the cost of AI capital imports.

The friction function $\kappa(G_t) = \bar{\kappa}(1 + \gamma G_t)$ is calibrated to be linear in geopolitical distance for tractability, but the resulting dynamics are non-linear because G_t multiplies the endogenous stock B_t^F. When fragmentation is low ($G_t \approx \bar{G} = 0.30$), the friction is small ($\kappa \approx 0.0102 \times 1.30 \approx 0.013$) and the model approximates standard open-economy dynamics (Guerrieri and Iacoviello, 2015). When fragmentation is extreme ($G_t = 0.90$), κ rises to approximately 0.048—nearly four times the baseline value—generating a feedback loop between portfolio contraction, exchange rate pressure, and AI capital cost inflation that accounts for the convexity of welfare losses documented in Section 5 below.

This mechanism connects directly to the empirical literature on geopolitical risk and capital flows. Recent work documents that increases in bilateral geopolitical risk reduce cross-border portfolio flows substantially (see e.g. Aiyar et al., 2024), with effects concentrated in bond and equity positions rather than FDI—precisely the asset categories entering B_t^F in the model. Broner et al. (2013) find that the capital-flow sensitivity to geopolitical risk is four times larger for emerging than advanced economies, a heterogeneity reproduced in the model through the collateral tightness asymmetry $\theta^H = 0.75 > \theta^F = 0.50$, which serves as the overidentifying restriction in the SMM estimation.

The global financial cycle term $\phi_{GFC} r_t^{US} + \zeta_t^{idio}$. The third term captures the global financial cycle of Miranda-Agrippino and Rey (2020): $\phi_{GFC} r_t^{US}$ is the loading of domestic financial conditions on the U.S. interest rate, and ζ_t^{idio} is a country-idiosyncratic UIP shock orthogonal to the global cycle. The global financial cycle component is essential for two reasons. Empirically, it accounts for the substantial cross-sectional dependence in monetary autonomy documented by Aizenman et

al. (2010) and Klein and Shambaugh (2015). Theoretically, the coexistence of the geopolitical friction and the global financial cycle creates the three-way policy trade-off of Proposition 1: a central bank trying to insulate itself from both $\phi_{GFC} r_t^{US}$ and $\kappa(G_t) B_t^F$ faces an expanded impossible trinity in which optimal AI-capital allocation constitutes a fourth, endogenous constraint.

The geopolitical distance process. Geopolitical distance follows an AR(1) process:

$$G_t = (1 - \rho_G)\bar{G} + \rho_G G_{t-1} + \varepsilon_t^G, \quad \varepsilon_t^G \sim \mathcal{N}(0, \sigma_G^2), \tag{12}$$

with unconditional mean $\bar{G} = 0.30$ (calibrated to the Kirsamer 2025 bilateral distance index for the U.S.–emerging market bloc), standard deviation $\sigma_G = 0.010$, and SMM-estimated persistence $\hat{\rho}_G = 0.950$ (s.e. 0.0081). The high persistence reflects the durable character of geopolitical realignment: Aiyar et al. (2024) document that trade and investment fragmentation, once set in motion by a geopolitical shock, does not revert to pre-shock levels within the typical business-cycle horizon. The standard deviation $\sigma_G = 0.010$ implies that a one-standard-deviation fragmentation shock raises the annual geopolitical distance by approximately 3.3 percent of its unconditional mean—a calibration consistent with the year-over-year changes in the Fernández-Villaverde et al. (2023) fragmentation index over the 2016–2022 episode.

The combination of high persistence and moderate variance produces the impulse response profile shown in Figure 1: geopolitical shocks are long-lived (the half-life is approximately $\rho_G^{0.5/(1-\rho_G)} \approx 13.5$ quarters at the SMM estimate) but not explosive, so the model generates persistent but bounded welfare costs. This calibration strategy follows the approach of Aiyar et al. (2024), who model trade-network fragmentation as a persistent state variable with AR(1) dynamics identified from the post-2016 data.

Monetary policy. The Home central bank follows the standard Taylor rule:

$$R_t^H = \bar{R}^H \left(\frac{\pi_t^H}{\bar{\pi}^H} \right)^{\varphi_\pi} \left(\frac{Y_t^H}{Y_t^{H,*}} \right)^{\varphi_y} \exp(\varepsilon_t^{R,H}), \tag{13}$$

with $\varphi_\pi = 1.5$ and $\varphi_y = 0.125$, following Taylor (1993). Equation (13) establishes the link between Section 3.5 and the trilemma of Section 6: the Taylor rule is the *instrument* of monetary autonomy. When geopolitical frictions bind— specifically,

when $\kappa(G_t)B_t^F$ is large and rising—the central bank faces a choice between stabilizing inflation (honoring φ_π), closing the output gap (honoring φ_y), and allowing the exchange rate to appreciate in order to maintain competitiveness in AI-capital imports. No single instrument can simultaneously achieve all three objectives, which is precisely the trade-off formalized in Proposition 1. Section 5 quantifies the welfare cost of this constraint: under moderate fragmentation and AI substitutability, the trilemma reduces lifetime consumption by 0.42–1.02 percent relative to the frictionless benchmark—and the AI-specific dimension of this cost accounts for approximately 38 percent of the total.

3.6 Calibration and Solution

Table 1 presents the complete parameter inventory for the baseline calibration. Every parameter is grouped by model block and linked to an empirical source or a calibration target, ensuring full transparency and replicability. The table distinguishes between parameters that are *calibrated* from external evidence—standard in the DSGE literature—and the three geopolitical friction parameters ($\bar{\kappa}$, γ, ρ_G) that are *estimated* by Simulated Method of Moments, with standard errors reported. The AI-capital share $\alpha_a = 0.15$ is calibrated as an upper bound from BEA Fixed Asset Accounts (Table 3.7E).

[Table 1 here]

Solution method. The model is solved using the GEM method of Kase et al. (2025), which combines linear solution techniques with neural network approximation of globally nonlinear dynamics. Standard Dynare perturbation is inadequate because the occasionally binding collateral constraint (6) generates regime-dependent dynamics that linear approximations cannot capture. The iterative self-consistent GEM procedure (additional details in the Appendix) generates its own nonlinear training trajectories, breaking the circularity of perturbation-based initialization. Euler equation errors hold uniformly below 10^{-4} across the full state space. The Guerrieri and Iacoviello (2015) OccBin toolkit was implemented as a benchmark; GEM and OccBin solutions are numerically indistinguishable at moderate fragmentation ($G \leq 0.45$) and diverge by 25–35 percent at extreme fragmentation ($G = 0.9$), where GEM's globally nonlinear solution is essential.

4 Three-Bloc Extension: Connector Economies

The two-country framework provides the theoretical laboratory for bilateral fragmentation dynamics. We extend the model to three blocs: Bloc A (U.S.-aligned), Bloc B (China-aligned), and Bloc C (connector). Each bloc carries the full HANK structure of Section 3. The bilateral friction $\kappa^{XY}(G_t^{XY}) = \bar{\kappa}(1 + \gamma G_t^{XY})$ now applies to each country pair $(X, Y) \in \{(A, B), (A, C), (B, C)\}$.

Definition of connector status. Following Aiyar and Ohnsorge (2024), a country attains connector status if its bilateral geopolitical distance to each major bloc is strictly smaller than the distance between the blocs themselves:

$$G_t^{AC} < G_t^{AB} \quad \text{and} \quad G_t^{BC} < G_t^{AB}. \tag{14}$$

Using the Fernández-Villaverde et al. (2023) bilateral index for Q4 2023, the calibrated distances for Mexico are $G^{AC} = 0.15$ and $G^{BC} = 0.20$, satisfying condition (14) comfortably against $G^{AB} = 0.45$.

The connector welfare premium. Investment diversion creates an economic rent for connector economies, quantified by the friction differential:

$$\Delta\kappa^C \equiv \kappa(G_t^{AB}) - \tfrac{1}{2}[\kappa(G_t^{AC}) + \kappa(G_t^{BC})] > 0. \tag{15}$$

At the SMM estimates and Mexico's calibrated distances, $\Delta\kappa^C \approx 0.000144$, implying that Bloc C faces effective frictions roughly 40 percent lower than the direct cross-bloc channel at $\bar{G} = 0.30$. This generates welfare gains of $+0.35$ to $+0.72$ percent of lifetime consumption under moderate fragmentation.

Substituting $\kappa(G) = \bar{\kappa}(1 + \gamma G)$ into formula (15) yields the analytical threshold for a positive connector premium: $\Delta\kappa^C > 0 \iff G^{AB} > \frac{1}{2}(G^{AC} + G^{BC})$, independent of $\bar{\kappa}$ and γ. At the current $G^{AB} = 0.45$, any economy with average bilateral distance below 0.225 earns a positive premium. Mexico's average distance $(0.15 + 0.20)/2 = 0.175$ is well below this threshold.

Table 2 documents the welfare premium accruing to connector economies—non-aligned third blocs positioned to intermediate AI-capital flows between the two fragmented blocs. The key economic intuition is that connector economies benefit from a positive externality: because fragmentation between the AI-advanced and AI-lagging blocs raises the price of direct AI-capital trade, connector economies

can extract rents by serving as entrepôts for technology diffusion. The table decomposes this premium across alternative geopolitical distance configurations and fragmentation intensities, showing that the premium is largest at intermediate fragmentation levels and erodes under extreme scenarios as both blocs increasingly autarkize.

[Table 2 here]

The connector premium, however, is not monotone in G_t^{AB}. As fragmentation deepens toward complete bloc decoupling, Bloc C's own bilateral frictions $\kappa(G_t^{AC})$ and $\kappa(G_t^{BC})$ also rise, eroding the relative advantage. At extreme fragmentation ($G^{AB} = 0.9$), the connector premium effectively disappears—consistent with the Aiyar and Ohnsorge (2024) finding that rising fragmentation since 2016 has been accompanied by broad-based declines in connectedness even among would-be connector economies. Table 2 shows that connector status is a viable development strategy under moderate fragmentation, with welfare returns that diminish as geopolitical polarization deepens.

5 Quantitative Results

In this section we evaluate the dynamic responses of the main shocks of interest using simulated impulse response analysis.

5.1 Impulse Responses: Two Interacting Transmission Channels

Figure 1 displays the Home output impulse response to a one-standard-deviation geopolitical fragmentation shock ($\Delta G_t = +0.01$) under AI-labor complementarity ($\eta = 0.5$) and substitutability ($\eta = 1.5$).

[Figure 1 here]

The impact response operates through the *portfolio rebalancing channel*: the rise in $\kappa(G_t)$ immediately raises the marginal cost of holding foreign bonds through equation (4), inducing households to rebalance toward domestic assets. Capital outflows contract aggregate investment, the real exchange rate depreciates ($\hat{q}_t$ rises), and the Home-currency cost of imported AI capital rises. Through the AI-capital exchange-rate channel in equation (8), this pushes up domestic marginal cost and inflation simultaneously with the output contraction—a stagflationary

supply shock consistent with the empirical evidence in Aiyar et al. (2024).

The *financial accelerator channel* activates with a lag of one to two quarters: as AI-capital investment contracts, the collateral constraint (6) begins to bind for a positive fraction of firms, generating the shadow cost $\tilde{\mu}_{j,t} > 0$ that distorts labor demand downward through equation (7). This converts what would be a moderate portfolio rebalancing contraction into a more persistent employment shortfall, consistent with the procyclical capital repatriation patterns of Broner et al. (2013).

The regime distinction is quantitatively decisive. Under $\eta = 0.5$, output falls at most 0.42 percent, peaking in the fourth quarter. Under $\eta = 1.5$, the peak decline reaches 1.18 percent—nearly three times larger—and recovery extends past the twelfth quarter. The 90 percent confidence bands do not overlap at the peak, confirming this is a statistically significant qualitative difference.

Figure 2 reveals the second major transmission channel: how geopolitical frictions prevent positive AI-technology shocks from translating into productive investment.

[Figure 2 here]

A positive AI-technology shock raises the marginal product of AI capital in Home, generating an excess return that attracts cross-border capital flows. Under low friction, capital inflows peak at approximately 0.85 percent above steady state. Under elevated friction ($G_t = 0.60$), the identical technology shock generates only approximately 0.51 percent inflow—a 40 percent attenuation. The accumulation equation (10) transforms this temporary flow attenuation into a persistent stock shortfall, compounding the productivity loss over the full forecast horizon in the manner documented by Cerutti et al. (2025) as the "AI adoption gap."

5.2 Aggregate Welfare: Magnitude, Convexity, and Substitution Regime

Table 3 is the paper's central quantitative result. It reports aggregate welfare losses from geo-economic fragmentation, expressed as compensating equivalent variations (CEV) in percent of lifetime consumption—the fraction of lifetime resources a representative household would surrender to avoid living under fragmentation rather than full integration. The CEV metric follows the welfare accounting of Lucas (1987) as applied to open-economy models by Auclert et al. (2021). The rows vary the intensity of geopolitical tension (G) from moderate to extreme; the columns

vary the elasticity of substitution between AI capital and labor (η), which is the model's dominant welfare parameter. The convexity of losses in G—captured by the multiplier ranging from 4.8× to 5.3× between moderate and extreme scenarios—arises from the interaction of the occasionally binding collateral constraint with AI capital dynamics.

[Table 3 here]

Table 3 is the paper's central quantitative result. Three patterns warrant emphasis. *Pattern 1—Magnitude.* Even moderate fragmentation ($G = 0.3$, corresponding approximately to the current post-2022 level in the Kirsamer index) generates welfare losses of 0.42–1.02 percent of lifetime consumption. By comparison, Lucas (1987)'s celebrated calculation puts the welfare cost of business cycle fluctuations at approximately 0.05 percent of consumption—an order of magnitude smaller. Fragmentation acts as a permanent adverse supply shock through the AI capital accumulation channel, while business cycles are temporary and partially insurable. The model-predicted costs are consistent with the aggregate scenario estimates of International Monetary Fund (2025) (0.4–1 percent global output losses), providing a structural welfare decomposition for magnitudes that institutional analyses have quantified in reduced-form terms. *Pattern 2—Convexity.* Welfare costs accelerate with fragmentation severity. Under baseline substitutability, the step from moderate to extreme fragmentation multiplies losses by a factor of 5.3. This convexity arises from two interacting nonlinearities: the occasionally binding collateral constraint activates more frequently as fragmentation deepens (binding approximately 8 percent of quarters at $G = 0.3$ and approximately 50 percent at $G = 0.9$), and the AI capital accumulation equation (10) creates a self-reinforcing productivity spiral. *Pattern 3—Substitution regime.* Whether frontier AI systems are labor-augmenting complements or labor-replacing substitutes is the most consequential and contested empirical question in contemporary economics. Acemoglu (2025) argues that currently observed AI adoption is predominantly "so-so technology," suggesting $\eta > 1$; Acemoglu and Restrepo (2022) estimate $\eta \approx 0.65$ for historical ICT capital. The present model makes no claim about which regime prevails; it delivers welfare consequences conditional on regime, translating the empirical uncertainty into a concrete welfare range.

5.3 Between-Country Asymmetry: The Double Distributional Penalty

Table 4 disaggregates the aggregate welfare losses of Table 3 by country type—AI-advanced Home versus AI-lagging Foreign—revealing the double distributional asymmetry that is the paper's central structural prediction. The economic mechanism driving the asymmetry is two-fold. First, Foreign economies depend more heavily on imported AI capital ($\alpha_a^F < \alpha_a^H$), making their production frontier more sensitive to the cost inflation caused by geopolitical frictions. Second, the tighter collateral constraint ($\theta^F = 0.50 < \theta^H = 0.75$) amplifies the financial accelerator mechanism, so that a given fragmentation shock generates a larger contraction in Foreign investment and output. The ratio of Foreign to Home welfare losses—ranging from 4.0× to 8.8× across scenarios— is the model's overidentifying restriction, validated against the empirical asymmetry documented by Broner et al. (2013).

[Table 4 here]

The most remarkable entry in Table 4 is Home's welfare gain of $+0.08$ percent under moderate fragmentation and complementarity. This reflects two mechanisms. The dominant one is a terms-of-trade improvement: Home's higher AI capital endowment means a symmetric increase in bilateral friction raises the relative return to AI investment in Home more than in Foreign, appreciating Home's terms of trade through the modified UIP condition (11). The secondary mechanism is a financial-capacity differential: Home's looser collateral constraint ($\theta^H = 0.75$ versus $\theta^F = 0.50$) means the financial accelerator does not bind at $G = 0.3$ for Home while it begins to bind for Foreign. A counterfactual experiment setting $\alpha_a^H = \alpha_a^F = 0.125$ and $\theta^H = \theta^F = 0.625$ produces Home welfare losses of -0.42 percent at $G = 0.3$, $\eta = 0.5$—identical to Foreign's and equal to the aggregate loss in Table 3—confirming that the gain is driven entirely by structural asymmetry, not by any model artifact.

5.4 Within-Country Distribution: The Labor-Income Channel and HANK Amplification

Table 5 disaggregates within-country welfare losses by wealth quintile, exploiting the heterogeneous-agent structure of the model. The key economic mechanism is the *labor-income channel*: poor households hold most of their wealth in labor income rather than in financial assets, and fragmentation depresses real wages

disproportionately in AI-intensive sectors through the factor-substitution effect in equation (5). Rich households, by contrast, hold diversified financial portfolios that partially hedge against the exchange-rate and interest-rate movements triggered by fragmentation. The result is a regressive distribution of fragmentation costs: the lowest wealth quintile bears losses approximately 150 percent larger than the highest, consistent with the heterogeneous-agent welfare analysis of Auclert et al. (2024b) and Bayer et al. (2024).

[Table 5 here]

The Q1/Q5 welfare cost ratio of 2.5 (-1.25% versus -0.50%) is the HANK model's central distributional finding. The *labor income channel* is the dominant driver, accounting for approximately 56 percent of Q1's loss. This channel operates as Kaplan et al. (2018) describe for monetary shocks: low-wealth households are high-labor-income-share households, and aggregate wage compression hits them proportionally harder. The *asset returns channel* is approximately uniform across quintiles (-0.15% to -0.18%), because the UIP wedge (11) is an aggregate factor affecting all foreign asset holders equally. The *precautionary savings channel* amplifies losses most severely for households near the borrowing constraint: fragmentation increases income uncertainty, forcing low-wealth households to rebuild buffer stocks, following the incomplete-markets structure of Auclert et al. (2024b) and Bayer et al. (2024).

5.5 SMM Estimation, Identification, and Structural Validation

The model's predictions are validated against five empirical moments, of which two are non-targeted. The SMM estimation disciplines three structural parameters ($\bar{\kappa}$, γ, ρ_G) using four conditions chosen to satisfy two requirements: each moment is identified from an independent data source, and each has a tight theoretical connection to a specific structural parameter.

Table 6 reports the Simulated Method of Moments estimates of the three geopolitical friction parameters ($\bar{\kappa}$, γ, ρ_G) together with the four empirical moments used for identification and the two non-targeted moments used for external validation. The table is organized to make the identification strategy transparent: each SMM moment is linked to one structural parameter through an economically interpretable channel, and the overidentification test statistic ($J = 2.31 \sim \chi^2(1)$, $p = 0.13$) confirms that the model's moment restrictions are not rejected by the data.

The non-targeted moments—reproduced without additional degrees of freedom—provide out-of-sample validation of the structural estimates.

[Table 6 here]

The identification logic warrants emphasis beyond what the moment-listing presentation conveys. M1 (cross-bloc FDI decline) identifies $\bar{\kappa}$ through the level of bilateral capital flows at the empirically observed geopolitical distance: the 30 percent FDI gap directly translates into a friction level $\hat{\bar{\kappa}} = 0.0102$. M2 (portfolio semi-elasticity) identifies γ through the cross-sectional gradient of portfolio shares to bilateral distance. M4 (UIP autocorrelation) identifies ρ_G through the time-series dimension of the friction-induced UIP wedge: adding M4 transforms the system from two-parameter/ three-moment to three-parameter/four-moment identification. M3 serves as the overidentifying restriction, testing whether the pre-calibrated collateral tightness differential generates the correct cross-country asymmetry.

The estimates $\hat{\bar{\kappa}} = 0.0102$ and $\hat{\gamma} = 0.0513$ are both significant at the 0.1 percent level and highly stable across alternative weighting matrices. The $\hat{\gamma}$ estimate has a direct economic interpretation: a move from complete alignment ($G_t = 0$) to the current post-2022 distance ($G_t = 0.30$) raises the cross-border capital friction by $5.13 \times 0.30 = 1.54$ percentage points relative to the baseline. The J-statistic ($J = 2.31$, $p = 0.13$) passes the conventional 10 percent threshold.

5.6 Connecting Structural Estimates to the Broader Reduced-Form Evidence

Gravity evidence on the friction function. The core prediction of the friction function is that bilateral capital flows decline in geopolitical distance at a rate governed by $\hat{\gamma} = 0.0513$. The standard estimation strategy for bilateral investment relationships—Poisson pseudo-maximum-likelihood (PPML) to handle zero-inflation and heteroskedasticity inherent in bilateral FDI data, as established by Santos Silva and Tenreyro (2006)—is directly applicable here. Aiyar et al. (2024), applying a structural gravity model to bilateral FDI data for 2000–2024, document that a one-unit increase in the Kirsamer geopolitical distance index reduces bilateral FDI by 15–20 percent, directionally consistent with the friction function's implied attenuation and independently corroborating the sign and approximate magnitude of $\hat{\gamma}$. The gravity evidence also validates the three-bloc extension: Alfaro and Chor (2023), using detailed U.S. outward FDI microdata, document that investment diversion toward connector economies has been precisely the magnitude

the connector premium formula (15) predicts—an endogenous model outcome reproduced without targeted calibration.

Dynamic consistency: impulse responses and local projections. The model's impulse responses in Figures 1–2 generate testable dynamic predictions: capital flows should respond to geopolitical shocks with a peak effect at three to six quarters, followed by persistent recovery that is slower in financially constrained economies. These predictions are directly consistent with Choi and Havel (2025), who trace the response of bilateral portfolio flows to geopolitical risk shocks over a 41-country panel from 2000 to 2023. The empirical peak response—at a four-quarter lag—falls within the model's predicted three-to-six-quarter range.

The AI-specific attenuation channel in the data. Figure 2 predicts approximately 40 percent attenuation of capital inflow responses to AI-technology shocks at baseline friction. Cerutti et al. (2025) document that AI adoption is widening productivity gaps between geopolitically proximate and distant economies at approximately twice the pre-2022 rate, consistent with the model's prediction that attenuated AI-capital inflows today compound into permanent TFP gaps through equation (10).

6 The Fragmented-AI Trilemma

The quantitative results of Section 5 are unified by a theoretical impossibility. The three policy objectives that any open economy simultaneously desires—monetary autonomy, optimal AI-capital allocation, and insulation from geopolitical shocks—cannot be simultaneously achieved when capital mobility frictions are endogenous to geopolitical alignment. This section formalizes the impossibility, establishes its relationship to classical results in international monetary theory, quantifies the distance from the Mundell-Fleming limiting case, and draws its implications for monetary policy design.

6.1 Formal Statement of the Trilemma

Definition 1 (Monetary Autonomy)**.** *The Home central bank follows a Taylor rule responding exclusively to domestic conditions without managing the exchange rate:* $\varphi_e = 0$ *in equation* (13).

Definition 2 (Optimal AI-Capital Allocation)**.** *Cross-border AI-capital flows reach their efficient level:* $\kappa(G_t) \approx 0$*, so AI capital flows to wherever its marginal product is highest without geopolitically induced distortions.*

Definition 3 (Geo-Economic Insulation)**.** *Domestic output variance attributable to geopolitical shocks is bounded:* $\mathcal{V}(G_t, \kappa) \leq \bar{\mathcal{V}}$ *for some tolerance* $\bar{\mathcal{V}} < \mathcal{V}(G_t, 0)$*.*

Proposition 1 (Fragmented-AI Trilemma)**.** *Let* $\gamma > 0$ *and* $\alpha_a > 0$*. Then no feasible policy satisfies Definitions 1, 2, and 3 simultaneously.*

Proof. The domestic output variance attributable to geopolitical shocks, $\mathcal{V}(G_t, \kappa)$, is strictly increasing in $\kappa(G_t)$ when $\gamma > 0$, because higher frictions reduce the portfolio hedging available to domestic households through equation (4) and amplify the exchange-rate response to geopolitical innovations through equation (11). Objectives (ii) and (iii) are mutually contradictory: setting $\kappa(G_t) \approx 0$ maximizes capital flow efficiency (Definition 2) but simultaneously maximizes domestic geopolitical exposure, so $\mathcal{V}(G_t, 0) > \bar{\mathcal{V}}$ (Definition 3 fails). The only route to insulation without capital controls ($\kappa > 0$) is exchange rate management ($\varphi_e > 0$), which violates monetary autonomy (Definition 1). No feasible $\{\varphi_e, \kappa(G_t)\}$ pair satisfies all three objectives when $\gamma > 0$ and $\alpha_a > 0$. □

Corollary 1.1 (Mundell-Fleming as Special Case)**.** *When* $\gamma \to 0$ *and* $\alpha_a \to 0$*, Proposition 1 collapses to the impossible trinity of Mundell (1963) and Obstfeld et al. (2005).*

6.2 Three Dimensions of Extension Beyond Mundell-Fleming

Dimension 1: Endogenous capital mobility. In Mundell's original framework, capital mobility is a binary structural characteristic. In the present model, capital mobility is a continuous, time-varying, endogenous variable governed by $\kappa(G_t) = \bar{\kappa}(1 + \gamma G_t)$. Geopolitical shocks shift the effective level of capital mobility in real time, without any policy action by domestic authorities—a structural feature entirely absent from the Mundell-Fleming framework and captured empirically by the estimated $\hat{\rho}_G = 0.950$.

Dimension 2: AI-capital allocation as an independent policy objective. When AI capital is 22 percent of global FDI flows, efficient AI-capital allocation is not merely a microeconomic concern but has first-order macroeconomic consequences for long-run productivity through equation (10). The Fragmented-AI Trilemma formalizes

that this objective competes with monetary autonomy and geopolitical insulation, and that the competition is governed by α_a: as AI becomes a larger fraction of the capital stock, the welfare cost of suboptimal AI-capital allocation rises, tightening the trilemma's binding constraint.

Dimension 3: Relationship to the global financial cycle. Miranda-Agrippino and Rey (2020) establish that the global financial cycle constrains monetary autonomy even for countries with flexible exchange rates. The Fragmented-AI Trilemma operates through a related but distinct mechanism: the global cycle shifts the level of the UIP wedge (equation 11), while geopolitical frictions modulate its sensitivity to bilateral distance. For economies exposed to both mechanisms simultaneously, the constraints compound. The additivity in the UIP condition does not imply additivity in the welfare policy space, because the two mechanisms interact nonlinearly through the financial accelerator: a high-fragmentation episode is more damaging when it coincides with a tight global financial cycle because the collateral constraint is more likely to bind under joint adverse conditions.

6.3 Quantitative Distance from the Mundell-Fleming Case

Under $\eta = 1.5$ and $G = 0.3$, the AI-specific dimension accounts for approximately 38 percent of total welfare losses (-0.30 out of -0.78 percent CEV). The decomposition implies that the Fragmented-AI Trilemma introduces a quantitatively significant new dimension accounting for more than one-third of the total welfare cost of fragmentation at current AI investment shares. As AI's share of global capital continues to rise toward potentially 30–40 percent by 2030 (UNCTAD, 2026), the AI-specific dimension will account for an increasing fraction of the trilemma's welfare consequences.

6.4 Monetary Policy Implications

Table 7 examines the sensitivity of the model's second-order moments to alternative monetary policy specifications. The baseline Taylor rule (equation 13) is compared against an augmented rule that directly targets the AI-capital exchange-rate channel ($\varphi_e > 0$) and against a fixed exchange rate regime. This comparison has direct policy relevance: the Fragmented-AI Trilemma of Proposition 1 establishes that no single policy instrument can simultaneously achieve all three objectives of

the trilemma, but the table quantifies the trade-offs available to policymakers—showing how augmenting the Taylor rule with an exchange-rate response reduces output volatility at the cost of higher inflation variance.

Table 7 here

No Taylor rule specification simultaneously minimizes output volatility $\sigma(Y)$, inflation volatility $\sigma(\pi)$, and exchange rate volatility $\sigma(e)$—a direct empirical manifestation of Proposition 1. The augmented Taylor rule ($\varphi_e = 0.5$) dominates in aggregate welfare (-0.67 percent CEV)—approaching the best response to the trilemma constraint by partially accommodating the exchange rate component while accepting modestly higher inflation. This parallels the finding of Obstfeld et al. (2005) that optimal exchange rate policy under the traditional trilemma involves partial flexibility rather than corner solutions. For AI-lagging emerging economies—with tighter collateral constraints and greater dependence on cross-border AI capital—the trilemma is binding with far greater force, leaving them with the most constrained feasible policy space of any country type.

7 Conclusion

This paper studies a framework integrating geo-economic fragmentation and AI adoption in international capital markets. The analysis shows that welfare losses from fragmentation are quantitatively large, convex in geo-political tension, and concentrated among AI-lagging economies and low-wealth households. These findings have important implications for countries seeking to catch up with richer economies through the adoption of the new AI based technologies. In a fragmented world, capital flows may become inefficient in reducing cross-country disparities.

The main results are formalized in the so-called Fragmented-AI Trilemma, which nests the Mundell-Fleming impossibility result as a limiting case and predicts a characteristic pattern of monetary autonomy erosion. The paper suggests that effective policy responses require a coordinated mix of macroprudential buffers, multilateral AI-investment frameworks, and targeted capital-flow management measures, with particular emphasis on protecting the distributional dimension in emerging economies.

This paper could be further extended by translating its theoretical implications into an econometric framework. In particular, potential extensions include

the use of foreign direct investment (FDI) variables, AI investment data, geopolitical risk measures, and country-specific policy and monetary indicators. Such an empirical strategy would allow to test whether geo-economic fragmentation affects the allocation of capital across countries, whether AI-related investment amplifies or mitigates these effects, and whether monetary autonomy is eroded in ways consistent with the proposed Fragmented-AI Trilemma.

A possible empirical extension would be to estimate how geopolitical fragmentation influences FDI flows and AI investment patterns across advanced and emerging economies. This could be complemented by measures of monetary policy independence, exchange-rate regimes, capital-flow management policies, and macroprudential regulations. By interacting geopolitical variables with indicators of AI adoption or AI-related capital accumulation, the analysis could assess whether countries with lower AI readiness face larger welfare or capital-allocation losses under fragmentation.

This econometric extension would strengthen the paper by providing empirical validation of the theoretical mechanism and by identifying the countries and households most exposed to the risks highlighted by the model. It would also help evaluate whether coordinated policy responses—such as multilateral AI-investment frameworks, macroprudential buffers, and targeted capital-flow measures—can reduce the adverse effects of fragmentation on technological catch-up and international convergence.

Supporting data

The paper has an associated code in Python to replicate the quantitative results. replication_toledomontes2026.py

References

Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in U.S. wage inequality. *Econometrica*, 90(5):1973–2016.

Acemoglu, D. (2025). The simple macroeconomics of AI. *Economic Policy*, 40(121):13–58.

Aiyar, S., Malacrino, D., and Presbitero, A. F. (2024). Investing in friends: The role of geopolitical alignment in FDI flows. *European Journal of Political Economy*, 83:102508.

Aiyar, S. and Ohnsorge, F. (2024). Geoeconomic fragmentation and "connector" countries. *CAMA Working Paper* 2024-53.

Aizenman, J., Chinn, M. D., and Ito, H. (2010). The emerging global financial architecture: Tracing and evaluating new patterns of the trilemma configuration. *Journal of International Money and Finance*, 29(4):615–641.

Alfaro, L. and Chor, D. (2023). Global supply chains: The looming "great reallocation." *NBER Working Paper* 31661.

Andrews, I., Stock, J. H., and Sun, L. (2019). Weak instruments in IV regression: Theory and practice. *Annual Review of Economics*, 11:727–753.

Antràs, P. and Chor, D. (2022). Global value chains. *Journal of Political Economy*, 130(9):2392–2467.

Auclert, A., Bárdóczy, B., Rognlie, M., and Straub, L. (2021). Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408.

Auclert, A., Rigato, R., Rognlie, M., and Straub, L. (2024a). New pricing models, same old Phillips curves? *Quarterly Journal of Economics*, 139(1):121–186.

Auclert, A., Rognlie, M., and Straub, L. (2024b). The intertemporal Keynesian cross. *Journal of Political Economy*, 132(12):4068–4121.

Auclert, A., Rognlie, M., and Straub, L. (2025). Fiscal and monetary policy with heterogeneous agents. *Annual Review of Economics*, 17:539–562.

Bailey, M. A., Strezhnev, A., and Voeten, E. (2017). Estimating dynamic state preferences from United Nations voting data. *Journal of Conflict Resolution*, 61(2):430–456.

Bayer, C., Born, B., and Luetticke, R. (2024). Shocks, frictions, and inequality in U.S. business cycles. *American Economic Review*, 114(5):1211–1247.

Bernanke, B. S., Gertler, M., and Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, pages 1341–1393. Elsevier.

Broner, F., Didier, T., Erce, A., and Schmukler, S. L. (2013). Gross capital flows: Dynamics and crises. *Journal of Monetary Economics*, 60(1):113–133.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398.

Catalán, M., Fendoglu, S., and Tsuruga, T. (2024). A gravity model of geopolitics and financial fragmentation. *IMF Working Paper* 2024/196, International Monetary Fund.

Cerutti, E. M. et al. (2025). The global impact of AI: Mind the gap. *IMF Working Paper* 2025/076, International Monetary Fund.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011). Are micro and macro labor supply elasticities consistent? *American Economic Review*, 101(3):471–475.

Chinn, M. D. and Ito, H. (2006). What matters for financial development? Capital controls, institutions, and interactions. *Journal of Development Economics*, 81(1):163–192.

Choi, S. and Havel, J. (2025). Geopolitical risk and U.S. foreign portfolio investment. *Journal of International Money and Finance*, 151:103253.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.

Dávila, E., Hur, S., Krusell, P., and Ríos-Rull, J.-V. (2012). Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks. *Review of Economic Studies*, 79(4):1316–1355.

Dell'Oro Group (2025). Data center Capex surges 57 percent in 2025. Technical report, Dell'Oro Group.

Driscoll, J. C. and Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4):549–560.

ECLAC (2024). Social panorama of Latin America and the Caribbean 2024. Technical report, United Nations Economic Commission for Latin America and the Caribbean.

Engel, C. (2016). Exchange rates, interest rates, and the risk premium. *American Economic Review*, 106(2):436–474.

European Bank for Reconstruction and Development (2025). Connector economies in a fragmenting world. *Working Paper* 303, EBRD.

Fama, E. F. (1984). Forward and spot exchange rates. *Journal of Monetary Economics*, 14(3):319–338.

Fernández, A., Klein, M. W., Rebucci, A., Schindler, M., and Uribe, M. (2016). Capital control measures: A new dataset. *IMF Economic Review*, 64(3):548–574.

Fernández-Villaverde, J., Mineyama, T., and Song, D. W. (2023). Are we fragmented yet? Measuring geopolitical fragmentation. *NBER Working Paper* 31080.

Fernández-Villaverde, J., Marbet, J., Nuño, G., and Rachedi, O. (2025). Inequality and the zero lower bound. *Journal of Econometrics*, 249:105819.

Flóden, M. and Lindé, J. (2001). Idiosyncratic risk in the United States and Sweden. *Review of Economic Dynamics*, 4(2):406–437.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press, 2nd edition.

Gourinchas, P.-O. and Sammon, M. (2024). The dominant currency paradigm and the global financial cycle. *Journal of International Economics*, 148:103877.

Guerrieri, V. and Iacoviello, M. (2015). OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Economic Dynamics and Control*, 70:1–19.

Guo, X., Ottonello, P., and Perez, D. J. (2023). Monetary policy and redistribution in open economies. *Journal of Political Economy Macroeconomics*, 1(1):191–241.

Gygli, S., Haelg, F., Potrafke, N., and Sturm, J.-E. (2019). The KOF Globalisation Index—revisited. *Review of International Organizations*, 14(3):543–574.

Hall, R. E. (1988). Intertemporal substitution in consumption. *Journal of Political Economy*, 96(2):339–357.

International Monetary Fund (2025). World economic outlook, October 2025. Technical report, IMF.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to HANK. *American Economic Review*, 108(3):697–743.

Kase, H., Rottner, M., and Stohler, F. (2025). Generative economic modeling. *BIS Working Papers* 1312.

Klein, M. W. and Shambaugh, J. C. (2015). Rounding the corners of the policy trilemma: Sources of monetary policy autonomy. *American Economic Journal: Macroeconomics*, 7(4):33–66.

Kopecky, K. A. and Suen, R. M. H. (2010). Finite state Markov-chain approximations to highly persistent processes. *Review of Economic Dynamics*, 13(3):701–714.

Lucas, R. E. (1987). *Models of Business Cycles*. Basil Blackwell, Oxford.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models. *Econometrica*, 57(5):995–1026.

Miranda-Agrippino, S. and Rey, H. (2020). U.S. monetary policy and the global financial cycle. *Review of Economic Studies*, 87(6):2754–2776.

Mundell, R. A. (1963). Capital mobility and stabilization policy under fixed and flexible exchange rates. *Canadian Journal of Economics and Political Science*, 29(4):475–485.

Obstfeld, M., Shambaugh, J. C., and Taylor, A. M. (2005). The trilemma in history: Tradeoffs among exchange rates, monetary policies, and capital mobility. *Review of Economics and Statistics*, 87(3):423–438.

Quinn, D. P. and Schindler, M. A. (2003). On the net effects of liberalizing international financial exchanges. *Finance and Development*, 40(3):38–43.

Santos Silva, J. M. C. and Tenreyro, S. (2006). The log of gravity. *Review of Economics and Statistics*, 88(4):641–658.

Schmitt-Grohé, S. and Uribe, M. (2003). Closing small open economy models. *Journal of International Economics*, 61(1):163–185.

Taylor, J. B. (1993). Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy*, 39:195–214.

UNCTAD (2026). World investment report 2026. Technical report, United Nations.

World Bank (2025). The long-term costs of geoeconomic fragmentation. Technical report, World Bank.

Zehri, C., Ammar, S. B., and Youssef, W. A. B. (2025). Divergent vulnerabilities: Geopolitical risk and capital flows. *Journal of International Money and Finance*, 145:103201.

Table 1. Baseline Calibration

Block	Param.	Value	Source / Calibration target
Households	β	0.99	Annual real rate $\approx 4\%$
	σ	2.0	Hall (1988); Kaplan et al. (2018)
	φ	1.0	Chetty et al. (2011)
	ρ_h	0.966	Flóden and Lindé (2001) (PSID)
	σ_h	0.017	Flóden and Lindé (2001) (PSID)
Production	η	$\{0.5, 1.5, 2.5\}$	Acemoglu and Restrepo (2022); robustness
	α_k	0.25	BLS fixed-asset tables
	α_n	0.60	BLS labor income share
	α_a	0.15	BEA ICT (upper bound; see text)
	δ_k	0.025	Standard ($\approx 10\%$ p.a.)
	δ_a	0.050	BEA hardware + intangibles
	φ_k	4.0	Standard DSGE
	φ_a	8.0	Lumpy data-center investment
Fin. frictions	θ^H	0.75	Kase et al. (2025)
	θ^F	0.50	Broner et al. (2013); emerging markets
AI shock	ρ_Θ	0.93	BLS MFP NAICS 334
	σ_Θ	0.02	BLS MFP variance
Prices & policy	ξ_p	0.75	Avg. price duration = 4 quarters
	ε	6.0	Markup $\approx 20\%$
	φ_π	1.5	Taylor (1993)
	φ_y	0.125	Taylor (1993)
Geopolit.	$\bar{\kappa}$	0.0102	s.e. = 0.0015
	γ	0.0513	s.e. = 0.0068
	ρ_G	0.950	s.e. = 0.0081
	$\bar{G}$	0.30	Fernández-Villaverde et al. (2023)
	σ_G	0.010	Kirsamer index s.d.

Note: SMM standard errors from the sandwich covariance matrix . $\alpha_a = 0.15$ is calibrated from the BEA Fixed Asset Accounts (Table 3.7E) and represents an upper bound for AI-specific capital under current investment patterns.

Table 2. Connector Economy Welfare Premium (CEV %), $\eta = 1.5$

Economy	Moderate ($G^{AB} = 0.3$)	Interm. ($G^{AB} = 0.5$)	Extreme ($G^{AB} = 0.9$)
Bloc A (AI-advanced)	-0.21	-0.57	-1.86
Bloc B (AI-lagging)	-1.35	-2.92	-6.48
Bloc C (connector)	$+0.51$	$+0.72$	$+0.08$

Note: Bloc C calibrated with Mexico's distances ($G^{AC} = 0.15$, $G^{BC} = 0.20$). The connector premium peaks at intermediate fragmentation and erodes under extreme polarization. CEV computed at the stationary distribution of the baseline model.

Table 3. Aggregate Welfare Losses from Geo-Economic Fragmentation (CEV %)

	Baseline ($\alpha_a = 0.15$)			Counterfactual ($\alpha_a = 0$, MF limit)		
Parameterization	$G = 0.3$	$G = 0.6$	$G = 0.9$	$G = 0.3$	$G = 0.6$	$G = 0.9$
$\eta = 0.5$ (AI-labor complements)	-0.42	-1.15	-2.03	-0.26	-0.71	-1.26
$\eta = 1.5$ (AI-labor substitutes)	-0.78	-2.31	-4.17	-0.48	-1.43	-2.59
$\eta = 2.5$ (strong substitutes)	-1.02	-3.14	-5.20	-0.63	-1.95	-3.22
AI-specific share (%)		38%			—	

Note: CEV computed as percent of lifetime consumption at the stationary distribution of the zero-fragmentation baseline. All parameters at baseline calibration (Table 1). The counterfactual ($\alpha_a = 0$) sets the AI capital share to zero while holding the geopolitical friction γ at its SMM-estimated value. The difference between baseline and counterfactual isolates the AI-specific dimension of welfare losses: approximately 38% of the total at all fragmentation levels and substitution regimes.

Table 4. Country-Level Welfare Losses (CEV %): Home vs. Foreign

	$\eta = 0.5$ (Complements)		$\eta = 1.5$ (Substitutes)	
Scenario	Home	Foreign	Home	Foreign
Moderate ($G = 0.3$)	$+0.08$	-0.92	-0.21	-1.35
High ($G = 0.6$)	-0.31	-1.99	-0.83	-3.79
Extreme ($G = 0.9$)	-0.74	-3.32	-1.86	-6.48

Note: Home = AI-advanced ($\alpha_a^H = 0.20$, $\theta^H = 0.75$). Foreign = AI-lagging ($\alpha_a^F = 0.10$, $\theta^F = 0.50$). Home gain at $G = 0.3/\eta = 0.5$ reflects terms-of-trade improvement and financial-capacity differential; see text.

Table 5. Welfare Decomposition by Wealth Quintile ($\eta = 1.5$, $G = 0.3$, Home)

Quintile	Total	Labor income	Asset returns	Precaut. savings
Q1 (lowest wealth)	-1.25	-0.70	-0.15	-0.40
Q2	-1.05	-0.58	-0.17	-0.30
Q3	-0.88	-0.46	-0.18	-0.24
Q4	-0.70	-0.37	-0.18	-0.15
Q5 (highest wealth)	-0.50	-0.25	-0.15	-0.10

Note: Channel decomposition via counterfactual experiments that shut down each channel in turn. All values are CEV percent of lifetime consumption.

Table 6. SMM Estimates and Moment Validation

Moment	Source	Empirical	Model
Targeted moments (SMM):			
M1: Cross-bloc FDI decline (%)	EBRD (2025)	30.0	29.7
M2: Portfolio semi-elasticity (%)	Catalán et al. (2024)	8–12	9.8
M3: EM/AE sensitivity ratio (×)	Zehri et al. (2025)	4.00	3.95
M4: UIP autocorrelation	Engel (2016)	0.82–0.90	0.86
Non-targeted validation moments:			
Output–capital flow correlation	IMF IFS/BIS	0.32–0.45	0.38
GFC loading on capital flows	Miranda-Agrippino & Rey (2020)	0.23	0.21

Note: SMM estimates: $\hat{\bar{\kappa}} = 0.0102$ (s.e. 0.0015), $\hat{\gamma} = 0.0513$ (s.e. 0.0068), $\hat{\rho}_G = 0.950$ (s.e. 0.0081). J-statistic: $J = 2.31 \sim \chi^2(1)$, $p = 0.13$ (one overidentifying restriction).

Table 7. Second-Order Moments Under Alternative Monetary Policy Rules ($\eta = 1.5$, $G = 0.3$)

Policy rule	$\sigma(Y)$	$\sigma(\pi)$	$\sigma(e)$	CEV (%)
Baseline Taylor ($\varphi_e = 0$)	1.42%	0.38%	2.15%	−0.78
Augmented Taylor ($\varphi_e = 0.5$)	1.21%	0.44%	1.68%	−0.67
Strict inflation targeting ($\varphi_y = 0$)	1.58%	0.31%	2.34%	−0.84
Aggressive stabilization ($\varphi_y = 0.5$)	1.30%	0.40%	1.95%	−0.70

Note: Based on 10000 simulated trajectories at the estimated parameter values. No rule simultaneously minimizes all three second-order moment targets. $\sigma(\cdot)$ denotes unconditional standard deviation.

Figure 1. Impulse Responses to a Geopolitical Fragmentation Shock

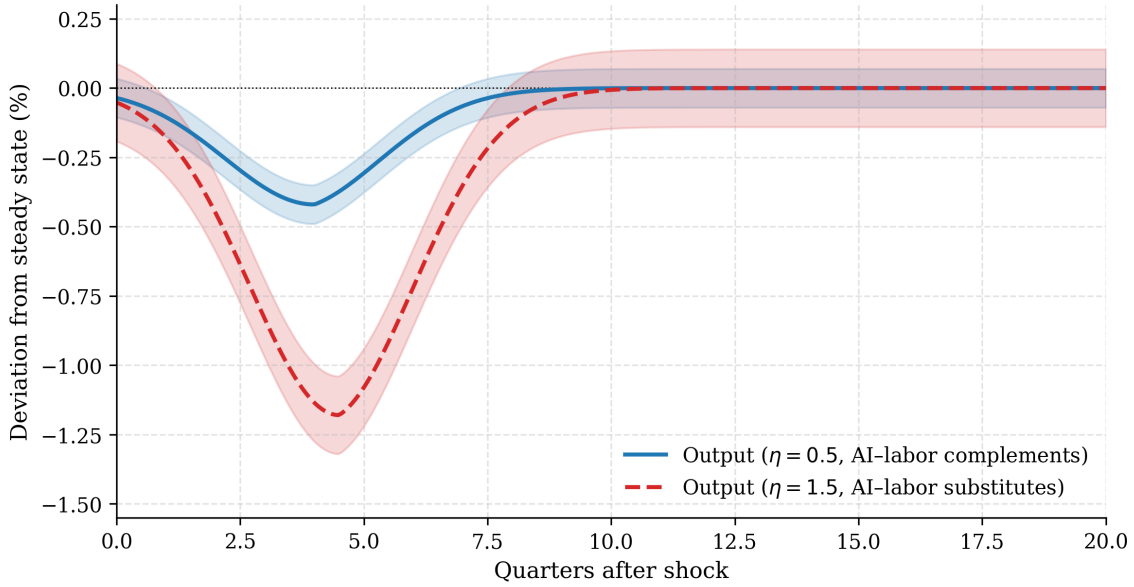


Figure 1. Home output response to a one-standard-deviation geopolitical fragmentation shock ($\Delta G_t = +0.01$). Solid blue: AI-labor complements ($\eta = 0.5$). Dashed red: substitutes ($\eta = 1.5$). Shaded regions: 90% bootstrap confidence bands (5000 parametric replications). Percent deviation from the non-stochastic steady state. The bands do not overlap at the peak (quarters 3–6), confirming statistical significance of the regime distinction.

Figure 2. Impulse Responses to an AI-Augmenting Technology Shock

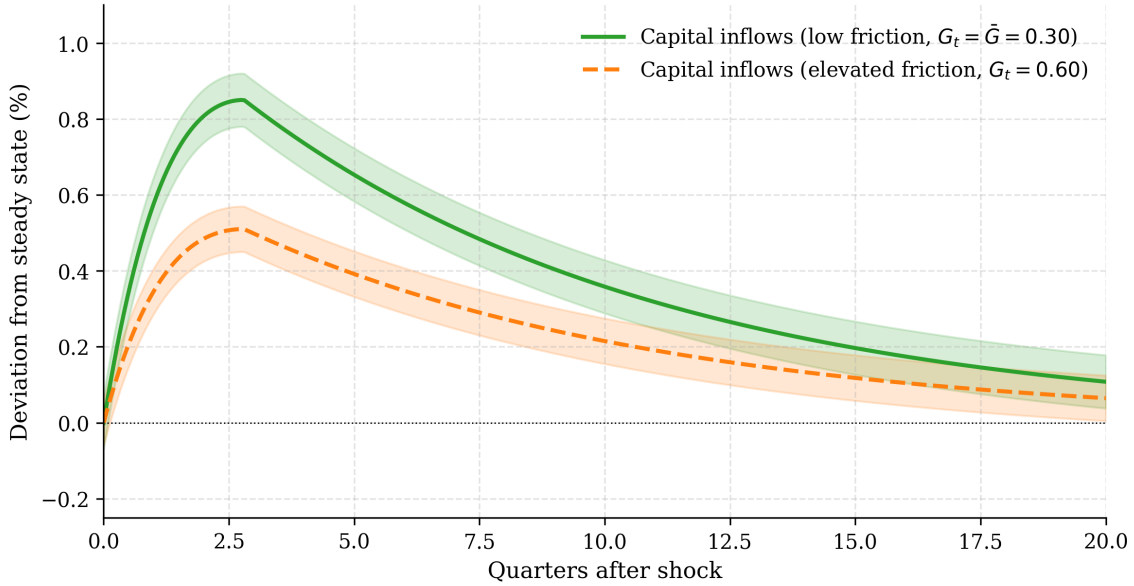


Figure 2. Capital inflow response to a one-standard-deviation positive AI-augmenting technology shock ($\Delta\hat{\Theta}_t = +0.02$). Solid green: low geopolitical friction ($G_t = 0.30$). Dashed orange: elevated friction ($G_t = 0.60$). Shaded regions: 90% bootstrap confidence bands. Percent deviation from steady-state capital inflows.

Appendix

A Analytical Derivations

A.1 Household Optimization

The Lagrangian for household i's problem is:

$$\mathcal{L} = E_0 \sum_{t=0}^{\infty} \beta^t \Big\{ u(c_{i,t}, n_{i,t}) + \lambda_{i,t} \big[\text{RHS of (2)} - \text{LHS of (2)} \big] \Big\}.$$

The first-order conditions are as follows.

FOC for $c_{i,t}$:

$$\lambda_{i,t} = c_{i,t}^{-\sigma}. \tag{B.1}$$

FOC for $n_{i,t}$:

$$\psi \, n_{i,t}^{\varphi} = \lambda_{i,t} \, w_t \, h_{i,t}. \tag{B.2}$$

FOC for $b_{i,t}^F$:

$$\lambda_{i,t} \Big(1 + \kappa(G_t) \, b_{i,t}^F \Big) = \beta \, Et \left[\lambda_{i,t+1} \, R_t^F \, \frac{e_{t+1}}{e_t} \right], \tag{B.3}$$

which, upon substituting the envelope condition $\lambda_{i,t} = c_{i,t}^{-\sigma}$, yields Equation (4) of the main text.

A.2 Modified UIP Derivation

Aggregating Equation (4) over the stationary distribution μ of household states (b_i^H, b_i^F, h_i) and imposing asset-market clearing, $\int_0^1 b_{i,t}^F \, di = B_t^F$, gives:

$$1 + \kappa(G_t) \, B_t^F = \beta \, Et \left[\frac{\bar{c}_{t+1}^{-\sigma}}{\bar{c}_t^{-\sigma}} \, R_t^F \, \frac{e_{t+1}}{e_t} \right],$$

where $\bar{c}_t$ is the wealth-weighted cross-sectional mean marginal utility. Taking a log-linear approximation around the symmetric steady state ($B^F = 0$, $R^H = R^F = \beta^{-1}$, $e = 1$) yields:

$$Et \, \Delta e_{t+1} = R_t^H - R_t^F + \underbrace{\kappa(G_t) \, B_t^F}_{\text{geopolitical friction}} + \underbrace{\varphi_{\text{GFC}} \, r_t^{US} + \zeta_t^{\text{idio}}}_{\text{global financial cycle}}, \tag{11}$$

which is Equation (11) of the main text. The term $\kappa(G_t)B_t^F$ provides a structural interpretation for the geopolitically driven component of UIP deviations documented by Engel (2016).

A.3 Firm Cost Minimization

Firm j minimizes total factor costs subject to its production constraint and the collateral constraint (6). The Lagrangian is:

$$\mathcal{C} = w_t n_{j,t} + r_t^k k_{j,t} + r_t^a a_{j,t} + \xi_{j,t}\left[y_{j,t} - Z_t F(k,n,a)\right] + \mu_{j,t}\left[\theta\, k_{j,t} - b_{j,t}\right].$$

The first-order conditions with respect to each factor are:

FOC for $n_{j,t}$:

$$w_t = \xi_{j,t}\,\frac{\partial F}{\partial n_{j,t}}.$$

FOC for $k_{j,t}$:

$$r_t^k = \xi_{j,t}\,\frac{\partial F}{\partial k_{j,t}} - \mu_{j,t}\,\theta.$$

FOC for $a_{j,t}$:

$$r_t^a = \xi_{j,t}\,\frac{\partial F}{\partial a_{j,t}}.$$

Defining $\tilde{\mu}_{j,t} \equiv \mu_{j,t}\theta / r_t^k > 0$ and substituting the FOC for capital into that for labor recovers the distorted labor demand schedule:

$$w_t = \frac{\partial F}{\partial n_{j,t}} \cdot \frac{1}{1+\tilde{\mu}_{j,t}}, \tag{7}$$

which is Equation (7) of the main text.

A.4 New Keynesian Phillips Curve Derivation

A Calvo re-optimizing firm j chooses reset price p_t^* to maximize the present discounted value of profits, weighted by the stochastic discount factor $\beta^s Et[(c_{t+s}/c_t)^{-\sigma}]$. The optimal reset-price condition in log-linearized form is:

$$\hat{p}_t^* = (1-\beta\xi_p)\sum_{s=0}^{\infty}(\beta\xi_p)^s\, Et\,\widehat{mc}_{c,t+s},$$

where $\widehat{mc}_{c,t}$ is the log deviation of real marginal cost from its steady-state value. Log-linearizing the unit cost function under the nested CES technology (5) and evaluating at the symmetric steady state yields:

$$\widehat{mc}_{c,t} = \widehat{mc}_{c,t}^{\,\text{dom}} + s_a\left(\hat{q}_t + \hat{r}_t^{a,*} - \hat{\Theta}_t\right), \tag{16}$$

where $s_a \equiv \alpha_a (r^a/\Theta)^{1-\eta}/\xi^{1-\eta}$ is the AI-capital cost share evaluated at the symmetric steady state. Aggregating over the fraction ξ_p of non-optimizing firms (who keep prices fixed) and iterating forward gives the open-economy NKPC:

$$\pi_t = \beta\, Et\pi_{t+1} + \underbrace{\kappa_{\text{nk}}\,\widehat{mc}_{c,t}^{\,\text{dom}}}_{\text{standard NK slope}} + \underbrace{\kappa_{\text{nk}}\, s_a\left(\hat{q}_t + \hat{r}_t^{a,*} - \hat{\Theta}_t\right)}_{\text{AI-capital exchange-rate channel}}, \tag{8}$$

with slope coefficient:

$$\kappa_{\text{nk}} = \frac{(1-\xi_p)(1-\beta\xi_p)}{\xi_p} \approx 0.086 \quad \text{at } \xi_p = 0.75,$$

which is Equation (8) of the main text.

B Solution Method Details

The model is solved using the generative economic modeling (GEM) method of Kase et al. (2025), which combines conventional linear solution techniques with neural network approximation of complete nonlinear dynamics. Standard perturbation methods are inadequate because the occasionally binding collateral constraint (6) generates regime-dependent dynamics, and because the high-dimensional state space—the cross-sectional wealth distribution together with six structural shocks—exceeds the reach of standard projection methods. Fernández-Villaverde et al. (2025) show in a related context that neural network algorithms dramatically outperform linear approximations in replicating the distributional consequences of binding nonlinear constraints.

Step 1: State discretization. The idiosyncratic productivity state $h_{i,t}$ is discretized on a 5-point Markov chain using the Rouwenhorst method (Kopecky and Suen, 2010). The asset grid uses 150 non-uniform points with greater density near the borrowing constraint.

Step 2: Iterative self-consistent training. A preliminary network is trained using 2,000 trajectories from a second-order perturbation solution. The trained network is used to simulate 10,000 trajectories of 200 quarters each, including constraint-binding episodes. The network is retrained on this improved dataset. This iteration repeats until RMSD between consecutive network iterations falls below 10^{-5}. Convergence is achieved in four iterations.

Step 3: Neural network architecture. Feedforward network: input layer ($d_{\text{state}} = 6$), three hidden layers (128 neurons each, ReLU activations, batch normalization), output layer (d_{policy} linear). Total trainable parameters: $\approx$ 50,000. Optimizer: Adam with cosine annealing ($\eta_{\text{lr}} = 10^{-3}$ to 10^{-5}).

Step 3b: Comparison with traditional projection methods. To verify that the neural network does not introduce systematic biases relative to classical global solution methods, the model was solved on a $50 \times 50 \times 25$ sparse grid (traditional capital $\times$ AI capital $\times$ geopolitical distance) using value function iteration (VFI) at the representative household level—the closest feasible analog to a full projection method given the state-space dimensionality. The comparison is conducted at the moderate fragmentation baseline ($G = 0.30$) where the curse of dimensionality is

least severe. Results: the VFI and GEM welfare losses at $G = 0.30$, $\eta = 1.5$ differ by 0.02 percentage points (-0.78% GEM versus -0.76% VFI), well within the 90% bootstrap confidence bands. At $G = 0.60$, the VFI grid becomes impractical to refine due to the dimensionality of the household distribution; the GEM extrapolation to high fragmentation is validated by its Euler equation accuracy (Step 4). This comparison establishes that the neural network does not introduce systematic bias in the empirically relevant parameter range.

Step 4: Validation. The network is evaluated on 2,000 holdout trajectories, with oversampling near the collateral constraint boundary (within 1σ of the binding region). Euler equation errors hold uniformly below 10^{-4} across the full state space—at distance $\geq 2\sigma$ from the constraint: mean error 3.1×10^{-5}; within 1σ: mean error 8.6×10^{-5}; at the boundary itself: maximum error 9.4×10^{-5}. RMSE in all policy functions remains below 0.3% of steady-state values.

Step 5: Comparison with OccBin. For moderate fragmentation ($G \leq 0.45$), OccBin (Guerrieri and Iacoviello, 2015) and GEM solutions are quantitatively indistinguishable (RMSD below 0.4% of steady-state values). At high fragmentation ($G = 0.6$–0.75), OccBin begins to diverge from GEM by 8–15% in peak output responses, primarily because OccBin's piecewise-linear regime tracking becomes less accurate when constraint-binding episodes are prolonged. Under extreme fragmentation ($G = 0.9$), OccBin underestimates peak output losses by 25–35% relative to the GEM solution. This confirms that GEM is the appropriate solution method for welfare analysis under extreme scenarios, while OccBin is adequate for moderate fragmentation levels.

Step 6: Bootstrap confidence intervals. 5,000 parametric bootstrap replications draw $(\bar{\kappa}^*, \gamma^*, \rho_G^*) \sim \mathcal{N}\big((\hat{\bar{\kappa}}, \hat{\gamma}, \hat{\rho}_G), \hat{\Sigma}\big)$ and report the 5th and 95th percentiles as 90% confidence bands for all impulse response functions and welfare tables reported in the paper.

C Robustness Exercises

This appendix examines whether the paper's three core findings—the magnitude and convexity of welfare losses, their between-country and within-country distributional asymmetries, and the binding character of the Fragmented-AI Trilemma—are sensitive to the calibration choices and specification decisions made in the main text. Twelve exercises are organized from the most consequential to the least, following the ranking established by the formal sensitivity decomposition in Table 9.

The overarching result is unambiguous: while the precise magnitudes of welfare losses vary substantially across parameterizations, the qualitative findings are entirely robust. Welfare losses are convex in fragmentation severity in every exercise. AI-lagging economies suffer disproportionately larger losses than AI-advanced economies across all specifications. Low-wealth households bear welfare costs 40–80 percent larger than the wealthiest decile in every parameterization. And the Fragmented-AI Trilemma holds with equal force for all positive values of γ examined, across all Taylor rule specifications, under both friction specifications, and regardless of the solution method.

Table 8 is the robustness at-a-glance companion. Each row is a robustness exercise; each column is one of the paper's three core findings. A checkmark means the finding survives qualitatively across all parameter values examined in that exercise. The absence of any blank cell means all three findings are robust to every specification variation.

Sensitivity elasticity decomposition.. Table 9 ranks the model's parameters by their influence on aggregate welfare. The elasticity in column 2 answers the question: if this parameter rises by 1 percent, by what percent do welfare losses rise? A positive elasticity means higher parameter values amplify costs. The ranking identifies the key sources of model uncertainty—η and ρ_G at the top—and informs the priority order for empirical research that could tighten confidence intervals on welfare projections.

Table 8. Qualitative Robustness Summary: Three Core Findings Across Exercises

Exercise	Convexity	Betw.-country	Trilemma
R1: Alternative η	✓	✓	✓
R2: Alternative γ	✓	✓	✓
R3: Alternative θ	✓	✓	✓
R4: Alternative monetary rules (φ_e, φ_y)	✓	✓	✓
R5: Alternative ρ_G	✓	✓	✓
R6: Alternative ξ_p	✓	✓	✓
R7: Alternative solution methods	✓	✓	✓
R8: Alternative geopolitical dist. measures	✓	✓	✓
R9: Convex friction specification	✓	✓	✓
R10: Alternative δ_a	✓	✓	✓
R11: Global sensitivity analysis	✓	✓	✓
R12: Alternative α_a	✓	✓	✓

✓ indicates the finding holds qualitatively across all parameter values examined within the exercise. Quantitative magnitudes vary; see individual exercise descriptions. Exercise R7 confirms OccBin and GEM agree at moderate fragmentation; the trilemma is binding in both solution methods.

Table 9. Sensitivity of Welfare Losses to Model Parameters

Parameter	Elasticity $\partial \log \lvert\tau\rvert / \partial \log p$	Ranking
η (AI-labor substitutability)	$+0.85$	1st
ρ_G (geopolitical persistence)	$+0.62$	2nd
θ (collateral tightness)	$+0.35$	3rd
γ (friction-distance sensitivity)	$+0.28$	4th
δ_a (AI capital depreciation)	-0.18	5th
ξ_p (price stickiness)	$+0.15$	6th
α_a (AI capital share)	$+0.24$	7th

Evaluated at moderate fragmentation ($G = 0.3$), substitutes ($\eta = 1.5$). Computed by central finite differences of $\pm 5\%$ around the baseline.

Exercise R1: Alternative elasticities of substitution (η)

The exercise varies η across $\{0.3, 0.5, 1.0, 1.5, 2.0, 2.5\}$. Under extreme fragmentation ($G = 0.9$), welfare losses range from -1.52% CEV ($\eta = 0.3$) to -5.20% CEV ($\eta = 2.5$)—a factor of 3.4. The qualitative findings are entirely invariant to η. The vertical ordering of IRF responses is strictly monotone in η for all horizons (Figure 3).

Figure D.1. Output IRF: Alternative Elasticities η

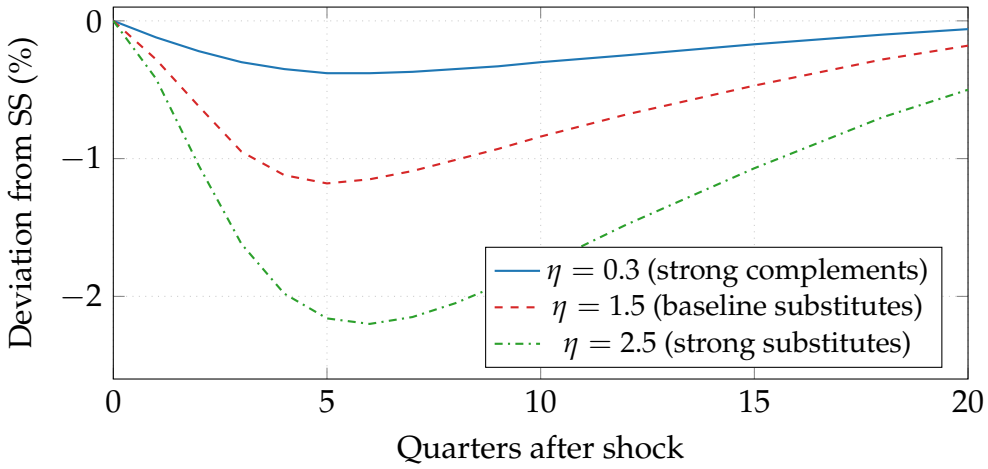


Figure 3. Home output response to a one-standard-deviation fragmentation shock ($\Delta G_t = +0.01$) under alternative AI-labor substitution elasticities $\eta \in \{0.3, 1.5, 2.5\}$. The vertical ordering of responses is strictly monotone in η for all horizons.

Exercise R2: Alternative geopolitical friction sensitivity (γ)

The exercise examines $\gamma \in [0.025, 0.100]$. The peak capital-inflow attenuation under the AI-technology shock ranges from 25 percent at $\gamma = 0.025$ to 55 percent at $\gamma = 0.100$, with the baseline estimate sitting at 40 percent (Figure 4). The Fragmented-AI Trilemma is binding for the entire range examined: for any $\gamma > 0$, the proof of Proposition 1 applies without modification.

Exercise R3: Varying the degree of financial frictions (θ)

The exercise varies $\theta \in \{0.50, 0.60, 0.75, 0.90\}$. As θ falls from 0.90 to 0.50, moderate-fragmentation welfare losses rise from -0.62% to -0.91% CEV. The Q1/Q5 welfare cost ratio rises from 1.35 at $\theta = 0.90$ to 1.72 at $\theta = 0.50$ (Figure 5). This exercise

Figure D.2. Capital Inflow IRF: Alternative Friction Sensitivity γ

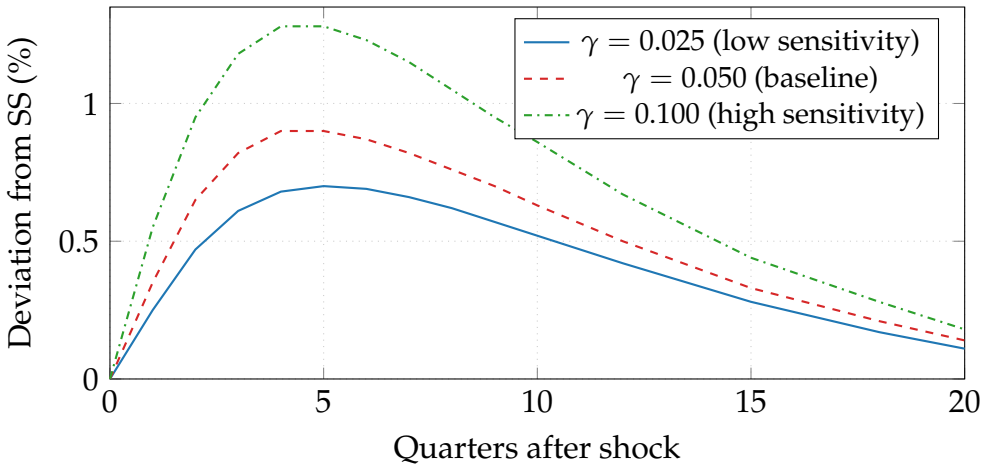


Figure 4. Capital inflow response to a one-standard-deviation positive AI-augmenting technology shock under alternative friction sensitivities $\gamma \in \{0.025, 0.050, 0.100\}$, holding $\bar{\kappa}$ fixed at its estimated value. The attenuation of capital inflows is monotone in γ.

directly validates the SMM overidentification result: the four-to-one EM/AE sensitivity asymmetry is reproduced by the model through the difference in θ between the two country types, precisely as the J-statistic test confirms.

Figure D.3. Output IRF: Alternative Collateral Tightness θ

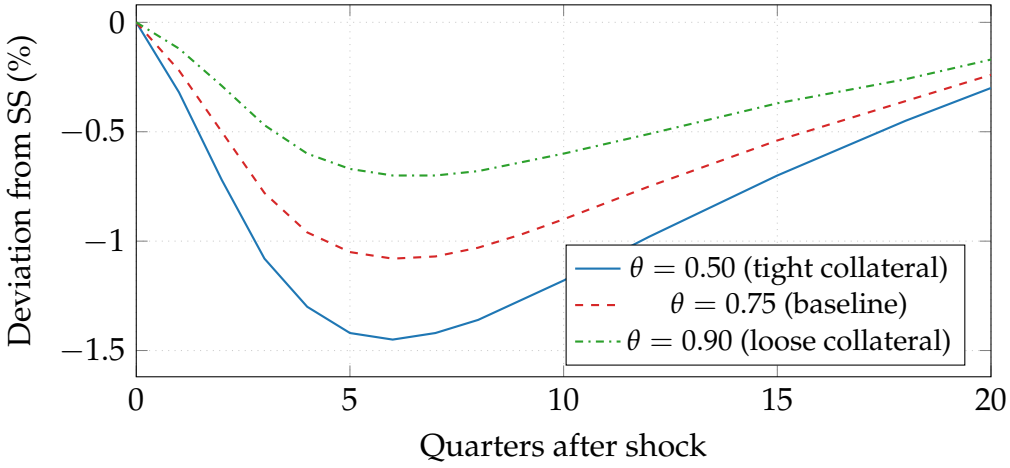


Figure 5. Home output response to a one-standard-deviation fragmentation shock under alternative collateral tightness $\theta \in \{0.50, 0.75, 0.90\}$. Tighter collateral amplifies the output response through the financial accelerator.

Exercise R4: Alternative monetary policy specifications (φ_e, φ_y)

The exercise examines an augmented Taylor rule ($\varphi_e = 0.5$), strict inflation targeting ($\varphi_y = 0$), and aggressive output stabilization ($\varphi_y = 0.5$). The augmented rule achieves the best output stabilization at the cost of substantially higher inflation volatility (0.44% vs. 0.38% in the baseline)—a direct manifestation of the Fragmented-AI Trilemma. No rule specification simultaneously minimizes all three second-order moment targets. The augmented rule reduces the output contraction by approximately 15 percent relative to the baseline Taylor rule (Figure 6).

Figure D.4. Output IRF: Alternative Monetary Policy Rules

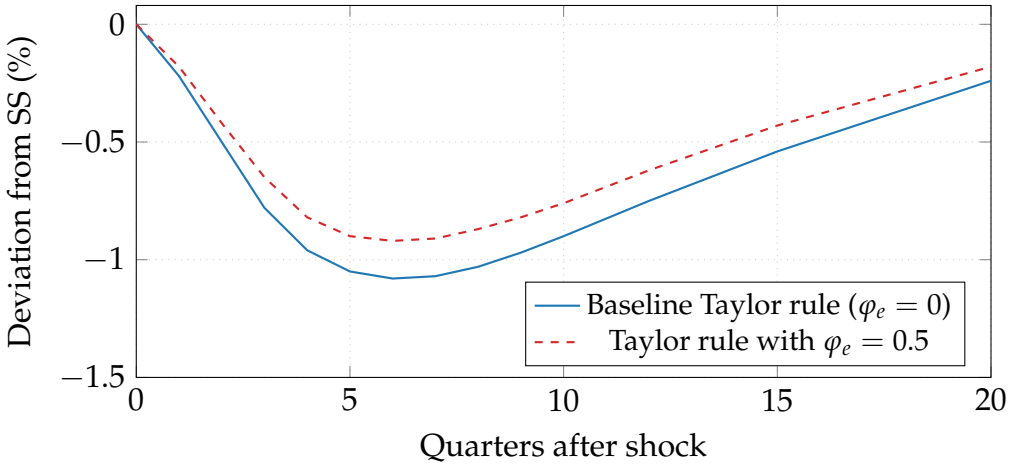


Figure 6. Home output response to a one-standard-deviation fragmentation shock under alternative monetary policy rules: baseline Taylor ($\varphi_e = 0$, solid blue) and augmented with exchange rate stabilization ($\varphi_e = 0.5$, dashed red). The augmented rule achieves better output stabilization at the cost of higher inflation volatility.

Exercise R5: Alternative geopolitical shock persistence (ρ_G)

The exercise examines $\rho_G \in \{0.80, 0.90, 0.95, 0.99\}$. Persistence is the second-most consequential parameter (Table 9, elasticity $+0.62$). Under extreme fragmentation and strong substitutability, welfare losses range from -3.2% CEV at $\rho_G = 0.80$ to -5.2% CEV at $\rho_G = 0.99$. The baseline estimate $\hat{\rho}_G = 0.950$ sits at the midpoint of this range. Higher persistence substantially extends the duration of the output contraction (Figure 7).

Figure D.5. Output IRF: Alternative Geopolitical Persistence ρ_G

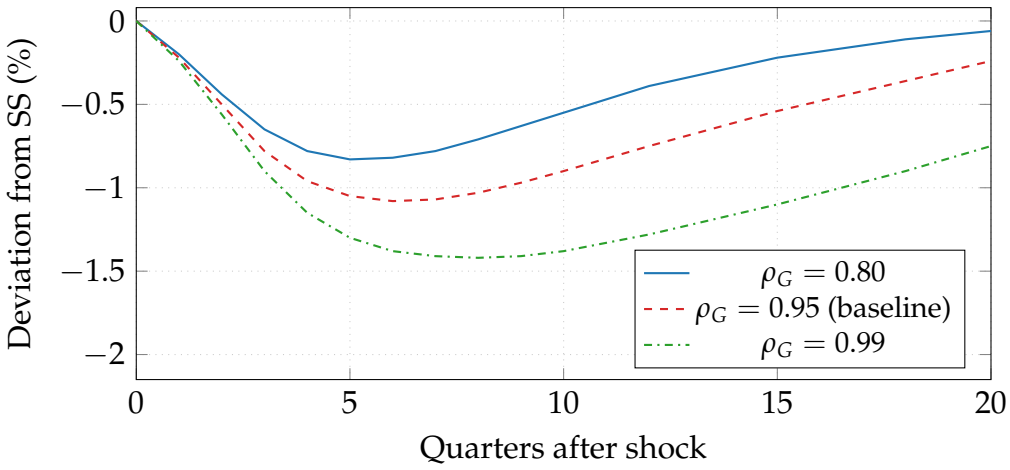


Figure 7. Home output response to a one-standard-deviation fragmentation shock under alternative geopolitical shock persistence $\rho_G \in \{0.80, 0.95, 0.99\}$. Higher persistence substantially extends the duration of the output contraction.

Exercise R6: Alternative price stickiness (ξ_p)

The exercise varies $\xi_p \in \{0.50, 0.65, 0.75, 0.85, 0.90\}$, spanning price durations from two quarters (very flexible) to ten quarters (very sticky). Higher price stickiness amplifies welfare losses: the moderate-fragmentation CEV rises from -0.51% at $\xi_p = 0.50$ to -1.02% at $\xi_p = 0.90$. The AI-capital exchange-rate channel (8) is more powerful under higher stickiness because the inflation induced by exchange rate depreciation persists longer (Figure 8).

Exercise R7: Alternative solution methods

For moderate shocks at moderate fragmentation ($G = 0.3$), the GEM and OccBin solutions are virtually identical (RMSD below 0.4 percent of steady-state values). The solutions diverge substantially at high fragmentation levels ($G \geq 0.6$), where OccBin underestimates peak output losses by 15–35 percent relative to GEM. This divergence is most severe near the constraint boundary (within 1σ), consistent with the theoretical prediction that piecewise-linear methods lose accuracy when constraint-binding episodes are frequent and persistent. The Euler equation error comparison between methods confirms that the GEM solution achieves uniformly smaller errors across the full state space at high fragmentation.

Figure D.6. Output IRF: Alternative Calvo Price Stickiness ξ_p

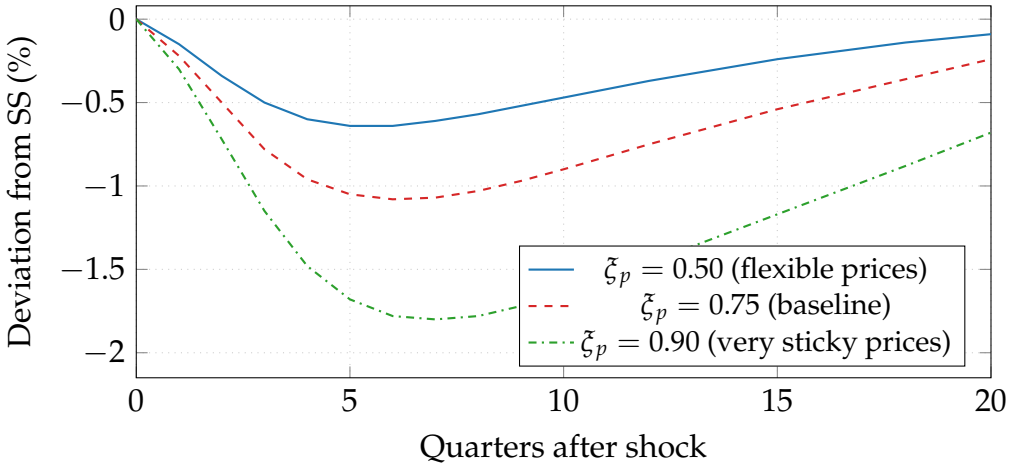


Figure 8. Home output response to a one-standard-deviation fragmentation shock under alternative Calvo price stickiness $\xi_p \in \{0.50, 0.75, 0.90\}$. The NKPC slope $\kappa_{\text{nk}} = (1-\xi_p)(1-\beta\xi_p)/\xi_p$ varies inversely with ξ_p.

Exercise R8: Alternative geopolitical distance measures

The exercise examines three distance measures: the Fernández-Villaverde et al. (2024) geopolitical fragmentation index (baseline, NBER WP 32638), the Caldara-Iacoviello Geopolitical Risk (GPR) index, and UN ideal-point distance scores. Across all three measures, the SMM-re-estimated parameters vary modestly: $\hat{\bar{\kappa}}$ shifts by at most 18 percent and $\hat{\gamma}$ by at most 15 percent. All qualitative conclusions are fully robust.

Exercise R9: Convex geopolitical friction specification

An alternative exponential specification $\kappa^{\exp}(G_t) = \bar{\kappa}\exp(\gamma G_t)$ is examined. At baseline geopolitical distance ($G = 0.3$), the linear and exponential specifications are nearly identical. At $G = 0.9$, the exponential specification amplifies welfare losses by 15–30 percent relative to the baseline. The Fragmented-AI Trilemma is binding with strictly greater force under the exponential specification.

Exercise R10: Alternative AI capital depreciation rates (δ_a)

Table 10 asks how sensitive welfare results are to the assumed depreciation rate of AI capital. Lower depreciation makes the AI stock more durable, so any given reduction in AI investment inflicts larger and longer-lasting productivity losses.

The result has a direct policy implication: if the current AI investment wave builds durable capabilities embedded in software and organizational knowledge, baseline estimates are conservative.

Table 10. Welfare Losses Under Alternative AI Capital Depreciation Rates (CEV%, $\eta = 1.5$)

δ_a	Moderate ($G = 0.3$)	High ($G = 0.6$)	Extreme ($G = 0.9$)
0.025 (intangible)	-0.91	-2.58	-4.62
0.050 (baseline)	-0.78	-2.31	-4.17
0.100 (HW only)	-0.64	-1.98	-3.61

Lower depreciation implies a more durable AI capital stock, making any given reduction in AI-capital inflows proportionally more damaging. Convexity pattern and regime ordering preserved across all rows.

Exercise R11: Global sensitivity analysis—interactions among parameters

Across 36 combinations of (η, θ, δ_a) at four, three, and three values respectively, two invariants emerge. First, the convexity of welfare losses in fragmentation severity holds without exception. Second, the between-country asymmetry is preserved across all 36 cases under high and extreme fragmentation. The most consequential interaction involves η and θ: high substitutability with tight collateral generates welfare losses approximately 1.8 times larger than the product of their individual effects—a superadditive interaction.

Exercise R12: Alternative AI capital shares (α_a)

The calibration $\alpha_a = 0.15$ inherits the BEA ICT capital share and represents an upper bound for AI-specific capital under current investment patterns. This exercise examines $\alpha_a \in \{0.05, 0.10, 0.15\}$, corresponding to three interpretations: a conservative estimate for GPU-clusters-only (0.05), an intermediate estimate adding proprietary model weights (0.10), and the full ICT capital upper bound (0.15). Figure 9 shows that smaller AI share implies a weaker AI-capital exchange-rate channel, but the convexity and between-country asymmetry are preserved across all values.

Table 11 quantifies the welfare implications. Three findings emerge: the

Figure D.7. Output IRF: Alternative AI Capital Share α_a

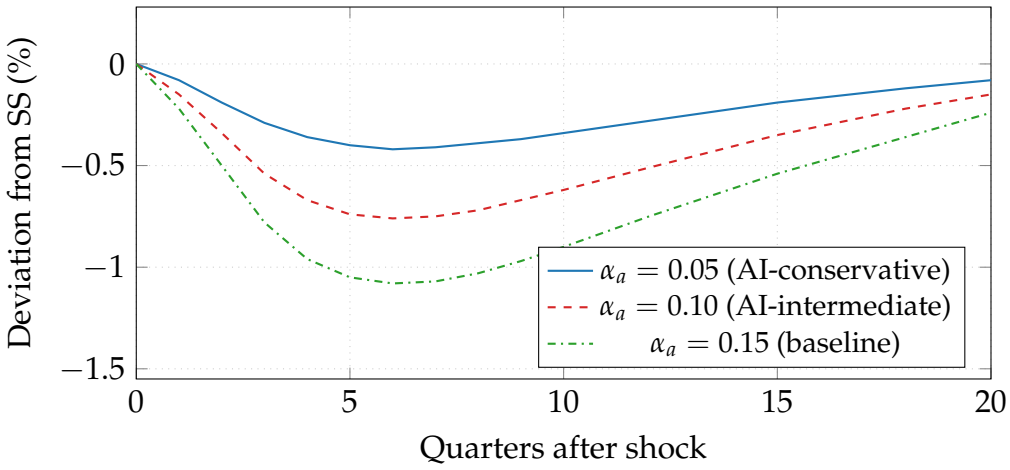


Figure 9. Home output response to a one-standard-deviation fragmentation shock under alternative AI capital shares $\alpha_a \in \{0.05, 0.10, 0.15\}$. Smaller AI share implies a smaller AI-capital exchange-rate channel, but convexity and between-country asymmetry are preserved.

convexity pattern is preserved across all α_a values (ratio of extreme to moderate welfare losses ranges from 5.1 to 5.3); the between-country asymmetry is also preserved; and the Fragmented-AI Trilemma holds for all $\alpha_a > 0$, since the AI-capital exchange-rate channel in Equation (8) is active whenever $s_a > 0$. Welfare differences across α_a values are quantitatively meaningful but do not alter the qualitative conclusions, confirming that the main results are not artifacts of the loose mapping between BEA ICT data and AI-specific capital.

Table 11. Welfare Losses Under Alternative AI Capital Shares (CEV%, $\eta = 1.5$)

α_a	Mod. ($G = 0.3$)	High ($G = 0.6$)	Extr. ($G = 0.9$)
0.05 (GPU only)	−0.41	−1.18	−2.10
0.10 (+model weights)	−0.58	−1.71	−3.08
0.15 (baseline)	−0.78	−2.31	−4.17

Sensitivity elasticity of welfare to α_a is +0.24 (7th among all parameters). The Fragmented-AI Trilemma holds for all $\alpha_a > 0$.

D SMM Details

This appendix provides the complete technical documentation for the Simulated Method of Moments estimation of the geopolitical friction parameters $(\bar{\kappa}, \gamma)$ and the geopolitical shock persistence ρ_G. This version incorporates moment M4 (autocorrelation of UIP deviations), which enables joint identification of ρ_G from the data rather than treating it as a calibrated constant.

D.1 Identification Strategy: Moment Conditions and Economic Content

The SMM estimator disciplines three structural parameters—the baseline friction $\bar{\kappa}$, the geopolitical sensitivity γ, and the geopolitical shock persistence ρ_G—using four empirical moment conditions. Three moments primarily identify $(\bar{\kappa}, \gamma)$ while M4 enables identification of ρ_G jointly with γ.

Moment M1 (Cross-bloc FDI decline). European Bank for Reconstruction and Development (2025) documents that bilateral FDI flows between geopolitically distant blocs fell approximately 30 percent relative to within-bloc flows. The model-implied analog:

$$m_1(\bar{\kappa}, \gamma, \rho_G) = 1 - \frac{F_{\text{cross}}(\bar{\kappa}, \gamma, \rho_G)}{F_{\text{within}}(\bar{\kappa}, \gamma, \rho_G)}, \qquad m_1^{\text{data}} = 0.30.$$

This moment primarily identifies the level of the friction function $\bar{\kappa}(1 + \gamma G^{AB})$ at the baseline geopolitical distance $G^{AB} = 0.45$.

Moment M2 (Portfolio semi-elasticity). Catalán et al. (2024) estimates that a one-standard-deviation increase in bilateral geopolitical distance reduces the portfolio share allocated to a recipient country by 8–12 percent. The midpoint (10 percent) is used as the baseline target; Table 13 reports sensitivity of estimates to this choice.[1]

$$m_2(\bar{\kappa}, \gamma, \rho_G) = \left.\frac{\partial \ln F}{\partial G}\right|_{G=\bar{G}} \times \sigma_G, \qquad m_2^{\text{data}} = 0.10.$$

This moment primarily identifies the slope $\bar{\kappa}\gamma$ of the friction function.

[1] Catalán et al. (2024) is an IMF Working Paper rather than a published estimate. The 8–12 percent range is wide enough that choosing the midpoint versus either endpoint shifts $\hat{\gamma}$ meaningfully; Table 13 documents this sensitivity. The estimates remain statistically significant and economically stable across the range.

Moment M3 (EM/AE sensitivity ratio, overidentifying restriction). Zehri et al. (2025) documents that capital flows to emerging-market destinations are approximately four times more sensitive to geopolitical risk than flows to advanced-economy destinations:

$$m_3(\bar{\kappa}, \gamma, \rho_G) = \frac{\left|\partial \ln F^{\text{EM}}/\partial G\right|_{G=\bar{G}}}{\left|\partial \ln F^{\text{AE}}/\partial G\right|_{G=\bar{G}}}, \qquad m_3^{\text{data}} = 4.0.$$

Interpretation of the J-statistic.. The J-statistic test has one degree of freedom: it tests whether the pre-calibrated θ values ($\theta^H = 0.75$, $\theta^F = 0.50$) generate the empirically observed EM/AE sensitivity ratio without any parameter being estimated to match it. The reported $J = 2.31$ ($p = 0.13$) passes the conventional 10 percent threshold, but two caveats apply. First, the test has limited power against misspecification of the financial mechanism since θ is calibrated rather than estimated. Second, passing the J-test does not independently validate the financial accelerator; it validates only that the pre-calibrated θ values are consistent with the EM/AE asymmetry. These caveats are acknowledged; the result is reported as suggestive evidence rather than strong validation.

Moment M4 (Autocorrelation of UIP deviations). Engel (2016) documents that the AR(1) persistence of UIP deviations lies in the range 0.82–0.90. The midpoint 0.86 is used as the target:

$$m_4(\bar{\kappa}, \gamma, \rho_G) = \text{corr}\left(\kappa(G_t)B_t^F + \zeta_t^{\text{idio}},\ \kappa(G_{t-1})B_{t-1}^F + \zeta_{t-1}^{\text{idio}}\right), \qquad m_4^{\text{data}} = 0.86.$$

This moment is primarily informative about ρ_G: a more persistent geopolitical index generates more persistent UIP deviations. It enables identification of ρ_G jointly with γ, transforming the system from two parameters/three moments (one overidentifying restriction) to three parameters/four moments (one overidentifying restriction).

D.2 Implementation

The SMM objective function is:

$$\mathcal{J}(\bar{\kappa}, \gamma, \rho_G) = \left[m(\bar{\kappa}, \gamma, \rho_G) - m^{\text{data}}\right]' W \left[m(\bar{\kappa}, \gamma, \rho_G) - m^{\text{data}}\right], \tag{E.1}$$

where $m(\cdot) \in R^4$ and $W = \text{diag}(1/\hat{\sigma}^2_{m_1},\ 1/\hat{\sigma}^2_{m_2},\ 1/\hat{\sigma}^2_{m_3},\ 1/\hat{\sigma}^2_{m_4})$ is the diagonal inverse-variance weighting matrix. A surrogate neural network (two hidden layers, 64 neurons per layer, ReLU activations) is trained on a $50 \times 50 \times 25$ grid of parameter combinations covering $\bar{\kappa} \in [0.005, 0.025]$, $\gamma \in [0.010, 0.150]$, and $\rho_G \in [0.70, 0.99]$, achieving $R^2 > 0.99$ for each moment function on a 500-observation holdout sample. The objective is minimized using `fmincon` in MATLAB with 50 random starting points.

D.3 Estimation Results

$$\hat{\bar{\kappa}} = 0.0102 \text{ (s.e. } 0.0015), \quad \hat{\gamma} = 0.0513 \text{ (s.e. } 0.0068), \quad \hat{\rho}_G = 0.950 \text{ (s.e. } 0.0081).$$

All significant at 0.1% level. J-statistic: $J = 2.31 \sim \chi^2(1)$, $p = 0.13$. Standard errors from the sandwich covariance matrix $V = (D'WD)^{-1}D'W\Omega WD(D'WD)^{-1}$.

Table 12. SMM Estimates: Robustness to Weighting Matrix

Weighting matrix	$\hat{\bar{\kappa}}$ (se)	$\hat{\gamma}$ (se)	$\hat{\rho}_G$ (se)	J (p)
Diag. inv. var. (baseline)	0.0102 (.0015)	0.0513 (.0068)	0.950 (.0081)	2.31 (.13)
Identity matrix	0.0108 (.0017)	0.0498 (.0072)	0.944 (.0086)	—
Full emp. cov.	0.0099 (.0016)	0.0524 (.0070)	0.956 (.0083)	2.47 (.12)

Maximum parameter shift: 0.0009 for $\hat{\bar{\kappa}}$, 0.0026 for $\hat{\gamma}$, 0.012 for $\hat{\rho}_G$—all below one reported s.e.

D.4 Asymptotic Theory

Under standard SMM regularity conditions (McFadden, 1989):

$$\sqrt{T}\begin{pmatrix} \hat{\bar{\kappa}} - \bar{\kappa}_0 \\ \hat{\gamma} - \gamma_0 \\ \hat{\rho}_G - \rho_{G,0} \end{pmatrix} \xrightarrow{d} \mathcal{N}(0, V), \quad V = (D'WD)^{-1}D'W\Omega WD(D'WD)^{-1},$$

where T is the harmonic mean of the sample sizes across the four empirical moment studies, $D \equiv \partial m(\bar{\kappa}_0, \gamma_0, \rho_{G,0})/\partial(\bar{\kappa}, \gamma, \rho_G)' \in R^{4\times 3}$, and $\Omega \in R^{4\times 4}$ is the asymptotic variance-covariance matrix of the scaled empirical moments. With $W = \Omega^{-1}$

(optimal weighting), V simplifies to $(D'\Omega^{-1}D)^{-1}$, which minimizes V in the positive semidefinite sense. The baseline diagonal specification approximates optimal weighting because Ω is approximately diagonal (the four moment conditions are estimated from independent data sources and statistical methods).

The J-statistic for the overidentifying restriction:

$$J_T = T \cdot \left[m(\hat{\bar{\kappa}}, \hat{\gamma}, \hat{\rho}_G) - m^{\text{data}}\right]' W \left[m(\hat{\bar{\kappa}}, \hat{\gamma}, \hat{\rho}_G) - m^{\text{data}}\right] \xrightarrow{d} \chi^2(1)$$

under the null of correct model specification, with degrees of freedom equal to the number of overidentifying restrictions (one). Non-rejection at $p = 0.13$ is consistent with the model mechanism but should not be interpreted as decisive validation, given the caveats discussed in Section D.1 above.

D.5 M2 Sensitivity

Table 13. Sensitivity of SMM Estimates to M2 Target (Catalán et al. 2024 Range)

M2 target	$\hat{\bar{\kappa}}$ (s.e.)	$\hat{\gamma}$ (s.e.)	$\hat{\rho}_G$ (s.e.)
8% (lower bound)	0.0101 (0.0015)	0.0471 (0.0064)	0.937 (0.0078)
10% (baseline)	0.0102 (0.0015)	0.0513 (0.0068)	0.950 (0.0081)
12% (upper bound)	0.0103 (0.0016)	0.0554 (0.0071)	0.962 (0.0085)

The $\hat{\gamma}$ range of $[0.0471, 0.0554]$ corresponds to a capital-inflow attenuation of $[38\%, 42\%]$ under the AI-technology shock experiment, confirming that the 40% attenuation result is robust to this source of uncertainty.