Technology Conditionality, Export Controls, and Welfare: Evidence from the Global AI Chip Regime

Subhasis Bera

subhasisbera@gmail.com

Amitendu Palit

isasap@nus.edu.sg

Dil B Rahut

drahut@adbi.org

Authors declare no conflict of interest.

Technology Conditionality, Export Controls, and Welfare: Evidence from the Global AI Chip Regime

Abstract

Frontier technology diffusion drives productivity growth, but access increasingly depends on geopolitical conditions. As frontier Artificial Intelligence (AI) becomes an input to production, we study how technology-conditioned capital access and export controls reshape technology diffusion, productivity, and welfare. We develop a general-equilibrium framework in which Technology-Conditioned Capital Flows (TCCF) operate as discriminatory factor-market distortions, creating a wedge between production TFP and welfare-relevant TFP when capital-transfer costs outweigh productivity gains. We then examine 25 countries over 2015–2024, combining pre-classification trends in production and welfare-relevant TFP, AI patenting, capital intensity, and subsequent access tiers. Three findings emerge. First, Tier-1 economies provide credible pre-treatment counterfactuals for future Tier-2 economies. Second, the gap between production and welfare-relevant TFP is systematically larger in future Tier-2 economies and rises with capital intensity. Third, prospective difference-in-differences and structural estimates indicate a widening productivity disadvantage for Tier-2 economies, with the largest welfare-equivalent income losses concentrated in resource-exporting upper-middle-income countries. Counterfactual simulations show that a multilateral MFN-style access rule reduces both the mean and dispersion of welfare losses relative to bilateral conditionality while preserving restrictions on the most sensitive technologies. The results identify a new channel through which geopolitical technology controls reshape international capital allocation and the distribution of gains from technological diffusion.

JEL: F13, F14, F21, O33, O47, C21

Keywords: Export controls; Technology conditionality; AI semiconductors; Geopolitical fragmentation; Difference-in-differences; Welfare analysis; Most-favoured-nation rules.

Technology Conditionality, Export Controls, and Welfare: Evidence from the Global AI Chip Regime

I. Introduction

Access to strategically important technologies increasingly depends not only on prices and technological capabilities but also on geopolitical and institutional constraints (Gopinath et al., 2025; Bastos et al., 2026). This shift is particularly consequential for artificial intelligence (AI), whose production increasingly depends on advanced semiconductor-based computing infrastructure, such as NVIDIA's A100 and H100 accelerators (Nature Electronics, 2024). Since computing capacity is unevenly distributed across countries (OECD, 2023), restrictions on its international allocation can generate heterogeneous effects on productivity, capital allocation, and welfare (Naito, 2017; Cai & Xiang, 2025). The U.S. export-control regime, initiated in October 2022 and expanded in January 2025, provides a natural context in which access to advanced AI computing capabilities is explicitly differentiated across countries through quantitative limits and licensing requirements (BIS, 2025; CSIS, 2025). More broadly, such policies create conditional international access to a critical productive input, whereby geopolitical relationships and cross-border economic commitments can shape firms' and countries' ability to obtain scarce technology. The 2025 U.S.–Saudi negotiations, including Saudi Arabia's announced $600 billion expansion of investment and trade with the United States, illustrate this broader interaction between technology access and international capital commitments (Reuters, 2025). We study this mechanism in a general-equilibrium framework and provide structural empirical evidence on how conditional access to advanced computing affects productivity, international capital allocation, and welfare across countries.

The study makes four contributions. Considering that non-price barriers can generate significant trade frictions even in the absence of conventional tariffs (Mei & Xu, 2026), we first develop a general equilibrium model in which a dominant technology supplier enforces a tiered access regime over a general-purpose technology (GPT) input. Countries seeking access face an implicit capital-transfer requirement, which we term a *Technology-Conditioned Capital Flow* (TCCF), rather than a conventional price tariff. The model delivers a sharp prediction. Production-side total factor productivity (TFP) and welfare-relevant TFP diverge. A country may record positive output growth upon obtaining access while experiencing a welfare deterioration if the present value of the capital outflow required to secure access exceeds the present value of the productivity gain. We call this the *TCCF gap*. It is the mechanism through which the tiered system transfers part of the gains from AI diffusion from technology-importing countries to the dominant supplier. International technology diffusion can generate substantial growth and welfare gains across countries (Cai & Xiang, 2025). Trade policy can also influence the extent and distribution of these gains from technology diffusion (Bai et al., 2025). This process is similar to optimal tariff extraction, but it works through licensing conditions rather than direct price discrimination (Kolev & Prusa, 1999).

This mechanism lies at the intersection of three bodies of work. The literature on geopolitical fragmentation, from the welfare arithmetic of tariff wars (Fajgelbaum et al., 2020) to structural deglobalisation (Goldberg & Reed, 2023) and on the aggregate cost of bloc formation (Aiyar et al., 2023), establishes that political realignment reshapes trade flows and reduces real income, with losses concentrated in emerging economies. This study extends this framework to *factor inputs.* This study considers conditionality on access to a critical productive input rather than tariffs on goods, consistent with the broader literature showing that non-price regulatory conditions can generate substantial international trade frictions (Mei & Xu, 2026). The literature on technology and growth links cross-country productivity differences to differences in technology diffusion (Eaton & Kortum, 2002), with international technology diffusion generating heterogeneous growth and welfare effects across countries (Cai & Xiang, 2025; Naito, 2017). Aggregate AI productivity contributions, while positive, remain modest and uneven across countries and sectors (Acemoglu, 2024). Our model provides the missing international dimension. Modest aggregate AI gains at the world frontier can become welfare-negative at the country level when access is conditional on capital transfer. Most directly, Lam and Santacreu (2026) study how geopolitical rivalry reshapes trade in ideas, finding that political risk suppresses technology licensing and reduces innovation simultaneously. We complement their supply-side analysis by focusing on the welfare consequences for technology-receiving countries and by providing structural empirical estimates grounded in observed data.

Our second contribution is empirical identification. We exploit the pre-classification panel of 2015–2024 to validate the identifying assumption underlying the difference-in-differences design. Tier-1 and Tier-2 countries followed statistically indistinguishable TFP trajectories prior to the tier regime across all outcome variables for which parallel trends are required, production TFP and AI patent filings, and we document precisely where the

assumption fails, namely for chip import unit values and bilateral FDI flows, and bound the implications for inference using the Rambachan and Roth (2023) sensitivity framework. This pre-registration-style validation establishes the Tier-1 never-treated group as a credible counterfactual and provides the identification infrastructure required to estimate treatment effects cleanly once post-classification data become available in subsequent PWT and WDI releases. We implement the Callaway and Sant'Anna (2021) doubly robust estimator and the Sun and Abraham (2021) interaction-weighted estimator as the prospective causal design; both are documented and validated against the pre-treatment data in the paper.

Third, we estimate the structural parameters of an augmented production function in which AI compute enters as a distinct factor input, using tier classification and ideal-point distance from the United States as instruments in a generalised method of moments (GMM) framework. The preferred 2SLS specification, estimated on the observed 2015–2024 panel, yields a first-stage F-statistic of 14.91, confirming instrument relevance. The estimated compute share is $\hat{\gamma} = 0.09$ (SE = 0.03 from the GMM specification), economically significant and consistent with Acemoglu's (2024) task-based estimates of AI's aggregate productivity contribution. This parameter calibrates the welfare magnitude of the TCCF gap and is identified entirely from observed data.

Fourth, we quantify the welfare consequences of the tier regime using the divergence between production TFP ($rtfpna$) and welfare-relevant TFP ($rwtfpna$) from Penn World Tables 10.01, a decomposition that is identified from observed pre-classification data and independent of post-treatment projections. This approach complements evidence that international integration can generate welfare gains while its adjustment costs depend on the frictions governing factor reallocation (Cacciatore, 2014). The mean TCCF gap for Tier-2 economies is 0.032 TFP index points relative to the Tier-1 trajectory, implying a welfare-equivalent income loss of approximately 2.8 per cent of GDP compounded over the ten-year horizon of the tier regime, within the 1.8–3.4 per cent range spanned by the counterfactual simulations. The FDI-conditioned subgroup bears the largest welfare penalty, consistent with the TCCF mechanism. Counterfactual simulations under three alternative regimes, universal Tier-1 access (CF1), full decoupling (CF2), and a multilateral most-favoured-nation rule (CF3), show that the current tier system is welfare-dominated by both CF1 and CF3, and is welfare-inferior even to full decoupling for capital-abundant economies where the capital-transfer cost exceeds the productivity gain from access.

This study links the structural TCCF mechanism to the data in three empirical steps. First, we derive testable predictions from the general equilibrium model. Existing evidence shows that the mode of international technology transfer can affect innovation incentives and economic growth (Glass & Saggi, 2002). Against this background, we predict that (i) conditional access generates a production–welfare TFP wedge, with the divergence larger in capital-intensive and FDI-conditioned economies; (ii) pre-classification divergence between production TFP and welfare-relevant TFP is systematically larger for future Tier-2 than for Tier-1 countries and increasing in the capital–output ratio; and (iii) multilateral MFN-style access rules weakly dominate bilateral conditionality for capital-abundant recipients. Second, using a panel of 25 countries over 2015–2024, we test these predictions by documenting the pre-classification TFP wedge patterns, estimating heterogeneity-robust difference-in-differences and event-study models for production TFP and AI patents, and structurally calibrating welfare-equivalent income losses at the country level. Third, we implement counterfactual policy experiments that map the TCCF mechanism into observable welfare outcomes under alternative access architectures, comparing the baseline tier regime to non-discriminatory MFN and unconditional access rules.

The remainder of this study is organised as follows. Section 2 reviews the related literature. Section 3 develops the TCCF model. Section 4 describes the data, the construction of the treatment variable, and the pre-treatment identification validation. Section 5 presents the structural production function estimates and descriptive evidence on the pre-treatment tier structure. Section 6 conducts the welfare decomposition and counterfactual simulations. Section 7 interprets the findings, outlines the analysis's limitations, and draws policy implications. Section 8 addresses the study's robustness, and Section 9 concludes.

2. Related Literature

This study interconnects four bodies of work: the theory of trade and market power, the economics of geopolitical fragmentation, technology diffusion and GPT adoption, and the empirics of causal inference in international economics.

2.1. Trade theory and factor-market distortions

A foundational trade theory shows that countries can use market power over tradable goods to improve their terms of trade at the expense of trading partners, though such actions may provoke retaliation (Harry Gordon Johnson, 1953). More recent work shows that trade barriers can also affect welfare through their effects on market structure and firms' incentives to innovate (Impullitti et al., 2022). This study develops an analogue for a non-price instrument. When a country possesses monopoly power over a critical factor input, it can extract quasi-rents not through a conventional export tariff but through conditionality on access. This mechanism is distinct from the Eaton-Kortum (2002) framework, in which technology diffuses freely across all trading partners through goods trade, and from the Melitz (2003) model, where firm selection and reallocation generate welfare gains under symmetric market access. Along similar lines, Grossman and Helpman (1991) show that North-South technology transfer generates product cycles whose direction and speed depend on the contractual environment governing technology licensing. Ethier and Markusen (1996) similarly show that the organisation of international technology transfer depends on firms' incentives to protect and internalise knowledge-based assets. We show that when the technology supplier has strong bargaining power, the payment required to access the technology can exceed the importing country's welfare gains. Antràs and de Gortari (2020) demonstrate that upstream-downstream positioning within global value chains determines the distribution of gains from trade; we apply analogous logic to the vertical structure of AI production, in which compute is the upstream GPT input and downstream AI applications are the final good. The TCCF mechanism we identify is therefore a form of upstream factor-market monopsony operated through political rather than economic barriers to entry.

2.2. Geopolitical fragmentation

A growing empirical literature documents the welfare costs of trade disruptions driven by geopolitical tensions. Fajgelbaum et al. (2020) develop a clear framework to measure the welfare effects of the U.S.- China tariff war and show that the own-price elasticity of import demand is a key determinant of the welfare losses associated with retaliatory tariffs. Fajgelbaum et al. (2024) show that the trade war generated global trade reallocations that offset only a fraction of the direct welfare losses borne by targeted economies. Ambrosino et al. (2026) further show that trade fragmentation generates broader macroeconomic effects through higher trade costs, inflationary pressures, and changes in monetary-policy conditions. Goldberg and Reed (2023) document the structural shift toward politically aligned bilateral trade flows, arguing that the welfare consequences of deglobalisation depend critically on whether alternative trading partners can supply comparable goods at comparable costs. Gopinath et al. (2025) similarly show that geopolitical tensions are reshaping international economic linkages and trade relationships. Aiyar et al. (2023) quantify the long-run GDP cost of full bloc formation at 2-7 per cent of world output, with losses concentrated in emerging economies that are most dependent on technology imports from advanced economies. Our study contributes to this literature in two respects. First, the distortion we study operates not through tariffs on goods but through conditionality on a productive factor input. Unlike conventional trade barriers, this mechanism conditions access to a critical technology input on associated capital flows. Second, because the TCCF distortion simultaneously raises output TFP and reduces welfare TFP, any measure that focuses solely on production-side outcomes understates the welfare costs we identify. The standard gains-from-trade formula developed by Arkolakis, Costinot, Rodríguez, and Clare (2012), which uses observed trade shares and trade elasticities, does not account for this gap between welfare gains and domestic production outcomes. Our model shows that this gap is the main channel through which the tiered technology regime affects the welfare of technology-importing countries.

2.3. General-purpose technology adoption and the macroeconomics of AI

Bresnahan and Trajtenberg (1995) characterise GPTs by three properties. These are pervasiveness, inherent potential for improvement, and innovational complementarities that raise the productivity of downstream R&D. Each characteristic intensifies the welfare implications of restricted access to frontier compute. When computing serves as a pervasive input, the resulting distortion spreads across a broad range of sectors within the economy. When advances in computing technology occur rapidly, delays in gaining access to these technologies can lead to persistent productivity gaps over time. Moreover, when innovation complementarities are strong, economies without access to frontier computing are also prevented from capturing the downstream innovation rents and technological spillovers generated by it. This channel is consistent with recent evidence that firm networks and international production relationships are important conduits for cross-country technology diffusion (Bastos et al., 2026). Acemoglu and Restrepo (2020) develop a task-based framework in which the aggregate welfare effect of automation depends on the balance between labour displacement from existing tasks and labour reinstatement

into new ones. Acemoglu (2024) applies this framework specifically to AI, finding that aggregate TFP contributions from current AI deployment are modest, between 0.53 and 0.66 percentage points over ten years, but vary substantially across adopting units. Recent evidence further shows that international technology diffusion through multinational production can generate substantial growth and welfare effects across countries (Cai & Xiang, 2025). Bai et al. (2025) further show that trade policy can influence the level and direction of international technology diffusion, with consequences for productivity and welfare. This study provides the international dimension of this finding. If the productivity gain from AI access is modest in absolute terms, a capital-transfer conditionality that imposes even a moderate outflow can dominate the gain, yielding a negative net welfare effect for the recipient. We formalise this condition as the TCCF gap and, for the first time, empirically estimate its magnitude. A study by Lam and Santacreu (2026) is most closely related to ours in this strand. They develop a growth model with geopolitical rivalry in technology licensing and show that political distance raises breach hazards, depresses royalty flows, and reduces both adoption and innovation rates in technology-receiving countries. Their analysis is supply-side and focuses on the equilibrium technology diffusion rate, while ours is demand-side and focuses on the welfare consequences for countries that access technology through conditional licensing arrangements. The two studies are thus complementary rather than competing.

2.4. Causal empirics in international economics.

This study makes a methodological contribution to the growing literature applying modern difference-in-differences estimators to international policy shocks. Caliendo et al. (2019) and Fajgelbaum et al. (2024) apply two-way fixed-effects estimators (TWFE) to quantify tariff incidence. Autor et al. (2013) exploit staggered regional exposure to Chinese import competition to identify labour-market effects. Recent advances by Callaway and Sant'Anna (2021) and Sun and Abraham (2021) resolve the well-documented heterogeneity bias in staggered TWFE estimators that arises when treatment effects are heterogeneous across cohorts and periods, a concern that is particularly acute in our setting, given that tier assignments are staggered and that the productivity effects of compute access are likely to differ systematically across income groups. We implement both estimators and verify that the point estimates are robust to the choice of control group (never-treated versus not-yet-treated) and to the Rambachan and Roth (2023) sensitivity analysis, which relaxes the parallel trends assumption by bounding the degree of pre-treatment trend violations. Nathan Nunn and Daniel Trefler (2014) show that domestic institutional quality influences comparative advantage by affecting contracting frictions. Building on this insight, this study uses UNGA ideal-point distances as a continuous measure of geopolitical alignment, which interacts with the binary tier classification to generate identifying variation in the heterogeneity analysis.

This study is therefore the first to (i) model technology-conditioned capital flows as a distinct factor-market distortion in a GPT production environment, (ii) causally identify the effect of a tiered compute access regime on productivity, capital flows, and innovation, and (iii) decompose the resulting welfare consequences into a production component and a welfare component, connecting the production-welfare wedge to the theoretical TCCF mechanism. The remainder of the study develops these contributions in sequence.

3. Conceptual framework

We developed the Technology-Conditioned Capital Flows (TCCF) model in four steps. We first characterise the production environment. We then model the tier access regime as a bilateral contracting problem. We derive the conditions under which production TFP and welfare TFP diverge, the TCCF gap. Finally, we state four empirically testable hypotheses that organise the subsequent analysis.

3.1. Production environment

Consider a world economy consisting of $i = 1, \dots, N$ countries. Each country produces a single final good using a constant-returns production function augmented with AI compute as a distinct factor input

$$Y_i = Z_i \cdot L_i^{\alpha} K_i^{\beta} A_i^{\gamma}$$

where Y_i is real output, L_i is quality-adjusted labour (persons engaged times average hours times the human capital index), K_i is the capital services stock, and A_i is the effective stock of AI compute accessible to country i. The parameter $\gamma \in (0,1)$ is the output elasticity of compute, α is the labour share, β is the capital share, and Z_i is Hicks-neutral total factor productivity. Factor shares satisfy $\alpha + \beta + \gamma = 1$. Taking logarithms, we have

$$\ln Y_i = \ln Z_i + \alpha \ln L_i + \beta \ln K_i + \gamma \ln A_i. \quad (1)$$

The compute input A_i is not freely tradable. It is produced exclusively in country s, the dominant supplier, using a technology that embodies the frontier general-purpose AI hardware. All other countries must import computer services from s. The welfare-relevant output of country i is consumption C_i, which equals output net of the capital transfer required to access compute and presented as

$$C_i = Y_i - \tau_i K_i^d \quad (2)$$

where $\tau_i \geq 0$ is the *conditionality rate*, the share of the domestic capital stock that country i must transfer to s as a precondition for receiving compute access A_i, and K_i^d is the domestically deployable capital stock prior to the transfer. When $\tau_i = 0$, as for Tier-1 countries, access is unconditional and $C_i = Y_i$. When $\tau_i > 0$, as for Tier-2 countries subject to FDI conditionality, the capital transfer reduces both consumption and the future productive capacity of the economy.

3.2. The tier regime as a bilateral contracting problem

Country s offers a menu of access contracts $\{A_i, \tau_i\}_{i=1}^N$ to maximise its own welfare, subject to participation constraints for each recipient country. The access menu partitions the world into three groups. Tier-1 countries receive unconstrained access at $\tau_i = 0$; they are close allies of s whose political alignment makes enforcement of any defection from the technology agreement costless. Tier-2 countries receive access $A_i > 0$ subject to $\tau_i > 0$. The supplier s sets τ_i to extract the full expected surplus of the access arrangement up to the recipient's participation constraint. Tier-3 countries receive no access and $A_i = 0$.

Let $\Delta W_i(A_i)$ denote the present value of the welfare gain that country i receives from computing access A_i, evaluated at the status quo capital stock and labour endowment. The participation constraint for country i requires is

$$\Delta W_i(A_i) - \tau_i K_i^d \geq 0. \quad (3)$$

The supplier maximises $\sum_{i \in \text{Tier 2}} \tau_i K_i^d$ subject to (3) holding with equality for each i. In the optimum, we have

$$\tau_i^* = \frac{\Delta W_i(A_i)}{K_i^d}. \quad (4)$$

Equation (4) states that the optimal conditionality rate is the ratio of the welfare gain from access to the capital stock, the fraction of capital that the recipient country would just be willing to surrender in exchange for compute access. When ΔW_i is small relative to K_i^d, as is plausible for countries with large capital stocks but modest AI absorption capacity, τ_i^* is small and the welfare cost of the arrangement is limited. When ΔW_i is large relative to K_i^d, as for capital-scarce, rapidly digitalising economies, τ_i^* is large and the capital transfer is welfare-significant.[1]

3.3. The TCCF gap: Production versus welfare TFP

Standard national accounts and the Penn World Tables decompose output growth into factor accumulation and TFP growth, where TFP is measured residually from the production function (1). This *production TFP* measure, denoted Z_i^P, captures the efficiency with which country i converts its measured factor inputs into output. It is unaffected by the capital transfer $\tau_i K_i^d$, which does not reduce Y_i directly but reduces C_i by removing capital from the domestic economy. The *welfare-relevant TFP* measure Z_i^W, which we identify empirically with $rwtfpna$ from the Penn World Tables, incorporates the terms of trade and is therefore sensitive to factor outflows.

[1] The supplier's ability to enforce τ_i^* depends on the credibility of the threat to deny access. Under full commitment, equation (4) holds exactly, and the supplier extracts the entire surplus. Under limited commitment, the effective conditionality rate $\hat{\tau}_i \leq \tau_i^*$ depends on the recipient's outside option—the availability of substitute compute from non-allied suppliers (e.g., domestic production or third-country procurement). The empirical analysis exploits variation in outside-option availability across countries as an additional source of identification.

We define the TCCF gap as

$$\Delta_i \equiv Z_i^P - Z_i^W. \tag{5}$$

Proposition 1 (TCCF Gap). *Under the tier access regime, $\Delta_i > 0$ for all Tier-2 countries with $\tau_i > 0$. The gap is increasing in the conditionality rate τ_i and in the capital intensity of the economy K_i^d/Y_i:*

$$\frac{\partial \Delta_i}{\partial \tau_i} > 0, \frac{\partial \Delta_i}{\partial (K_i^d/Y_i)} > 0. \tag{6}$$

Proof: Production TFP growth following access is

$$\hat{Z}_i^P = \gamma \hat{A}_i \tag{7}$$

where a hat denotes a log-change. Welfare TFP growth, adjusted for the capital outflow, is

$$\hat{Z}_i^W = \gamma \hat{A}_i - \frac{\tau_i K_i^d}{Y_i} \cdot \frac{1}{\beta} \tag{8}$$

since the capital transfer reduces the capital base by $\tau_i K_i^d$, which translates into a welfare-equivalent output loss of $\beta \cdot \tau_i K_i^d / K_i^d = \beta \tau_i$ in proportional terms. Subtracting (8) from (7)

$$\Delta_i = \hat{Z}_i^P - \hat{Z}_i^W = \frac{\tau_i}{\beta} \cdot \frac{K_i^d}{Y_i} > 0. \tag{9}$$

Since $\tau_i > 0$ and $\beta > 0$ and $K_i^d/Y_i > 0$, the gap is strictly positive for all Tier-2 countries with active conditionality. The comparative statics in (6) follow directly from the differentiation of (9).

Corollary 1. *A Tier-2 country with FDI conditionality may exhibit positive production TFP growth following tier classification while experiencing a decline in welfare TFP, if and only if*

$$\gamma \hat{A}_i < \frac{\tau_i}{\beta} \cdot \frac{K_i^d}{Y_i}. \tag{10}$$

Condition (10) formalises the model's central result. The compute productivity gain $\gamma \hat{A}_i$ must be weighed against the welfare cost of the capital transfer, which is amplified by the capital intensity of the economy and attenuated by the capital elasticity β. For capital-abundant economies with low γ, which may include middle-income countries with large physical capital stocks but limited AI absorption capacity, the condition is likely to bind. For capital-scarce, digitally advanced economies with high γ, compute access is welfare-improving net of the transfer.

3.4. Aggregate welfare and counterfactual regimes

Let world welfare be $\mathcal{W} = \sum_i \omega_i C_i$, where ω_i are welfare weights. Under the tier regime, world welfare is

$$\mathcal{W}^{tier} = \sum_{i \in T_1} \omega_i Y_i + \sum_{i \in T_2} \omega_i (Y_i - \tau_i K_i^d) + \sum_{i \in T_3} \omega_i Y_i^0 \tag{11}$$

where Y_i^0 denotes the counterfactual output of Tier-3 countries under autarky in compute. We compare (11) to three counterfactual regimes analysed in Section 4.

- **CF1 (Universal Access):** $\tau_i = 0$ for all i, $A_i = \bar{A}$ for all $i \notin T_3$. This is the first-best allocation under non-discriminatory technology transfer.
- **CF2 (Full Decoupling):** $A_i = 0$ for all $i \notin T_1$. This is the autarky counterfactual.
- **CF3 (MFN Rule):** $A_i = \bar{A}$, $\tau_i = \bar{\tau}$ for all i where $\bar{\tau}$ is a uniform non-discriminatory conditionality rate set to balance the supplier's budget constraint.

Proposition 1 implies $\mathcal{W}^{CF1} > \mathcal{W}^{tier}$, when welfare weights are equalitarian, the tier regime is welfare-dominated by universal access. Whether $\mathcal{W}^{tier} \gtrless \mathcal{W}^{CF2}$ depends on the magnitude of the compute productivity gain $\gamma \hat{A}_i$ relative to the conditionality cost, and is therefore an empirical question resolved in Section 6.

3.5. Hypotheses

The model generates four empirically testable hypotheses, which are directly tested in Sections 4-6.

Hypothesis 1 (Compute Scarcity Premium). *Tier-2 classification raises the unit value of AI compute imports relative to the pre-treatment trend, as conditional access creates a scarcity rent captured in import prices. The effect is larger for countries with fewer outside options.*

Formally, from equation (4), the conditionality rate τ_i^* is inversely related to the availability of substitute computing hardware suppliers. Countries with no domestic semiconductor industry and no access to non-US-allied suppliers face the highest scarcity premium.

Hypothesis 2 (Capital Reallocation). *Tier-2 classification with FDI conditionality increases bilateral FDI flows from the recipient country to the United States, as the capital-transfer requirement is met through direct investment pledges. This effect is absent for Tier-1 countries and for Tier-2 countries without explicit conditionality.*

Formally, from (3)–(4), $\tau_i > 0$ implies an observable outflow of $\tau_i K_i^d$ denominated in bilateral FDI to the supplier country s.

Hypothesis 3 (Production-Welfare TFP Divergence). *Tier-2 countries exhibit positive growth in production TFP Z_i^P following tier classification, but smaller, or negative, growth in welfare-relevant TFP Z_i^W. The TCCF gap Δ_i defined in equation (5) is positive and larger for capital-intensive, capital-scarce economies, consistent with condition (10).*

Hypothesis 4 (Heterogeneous Welfare Effects). *The welfare cost of the tier regime is increasing in capital intensity K_i^d/Y_i and in the conditionality rate τ_i, and is decreasing in the compute productivity gain $\gamma \hat{A}_i$. Middle-income economies face the largest welfare penalty; high-income technology-importing economies benefit on net from conditional access.*

Formally, from (9) and (10), the TCCF gap is a monotone function of $\tau_i/\beta \cdot K_i^d/Y_i$. Countries with large capital-to-output ratios and high conditionality rates are predicted to show the largest gap between production and welfare TFP, a prediction tested directly in the heterogeneity analysis of Section 5.

4. Empirical Strategy

4.1.1 Data

The analysis employs a panel of 215 countries over 2015–2024, yielding 2,150 country-year observations. Four sources are combined. The World Development Indicators (World Bank) provide data on GDP per capita, trade openness, R&D intensity, and FDI inflows. Penn World Tables 10.01 (Feenstra, Inklaar, and Timmer 2015) supply production TFP ($rtfpna$), welfare-relevant TFP ($rwtfpna$), factor shares ($labsh$), capital services ($rkna$), employment (emp), average hours (avh), and the human capital index (hc). UN Comtrade supplies bilateral trade flows in HS 8542 (integrated circuits), from which we construct the chip import unit value, the primary measure of compute access costs, and chip import value, which proxies for the AI compute input A_i in the structural estimation of Section 5. OECD patent data supplies AI patent filings under IPC class G06N.

The primary estimation sample comprises 25 countries assigned to the three-tier structure in Section 3.2. There are 10 Tier-1 (never-treated) countries, 12 Tier-2 (treated) countries, and 3 Tier-3 (restricted) countries, the last excluded from the DiD sample. Table 1 documents systematic pre-treatment differences across tiers. Mean GDP

per capita is USD 45,402 for Tier 1, USD 10,623 for Tier 2, and USD 8,599 for Tier 3. The mean AI patent output is 2,090 filings for Tier 1 and 34 for Tier 2. The ideal-point distance from the United States (IPD), constructed from UNGA voting records following Bailey, Strezhnev, and Voeten (2017), increases monotonically across tiers, 1.21, 3.06, and 3.37 for Tiers 1, 2, and 3, respectively, confirming that tier assignment tracks geopolitical alignment rather than development status alone.

Table 1. Descriptive Statistics by Tier Classification (2015–2024)

	Tier 1: Allied		Tier 2: Conditional		Tier 3: Restricted	
Variable	Mean (SD)	N	Mean (SD)	N	Mean (SD)	N
AI patent filings	2,090.00 (4,019.89)	60	33.89 (86.34)	72	1,029.92 (1,781.06)	18
Chip import unit value (USD)	2,428.48 (2,529.73)	41	1,578.09 (448.01)	17	10,024.36 (1,055.19)	6
GDP per capita (constant USD)	45,402.13 (8,839.80)	100	10,622.59 (11,264.25)	120	8,599.44 (2,573.28)	30
Ideal point distance from U.S.	1.21 (0.51)	90	3.06 (0.33)	108	3.37 (0.79)	27
Labour share of income	0.59 (0.04)	50	0.48 (0.10)	50	0.50 (0.11)	15
R&D expenditure (% GDP)	2.80 (0.82)	76	0.69 (0.39)	68	1.45 (0.70)	20
Total factor productivity (index)	1.00 (0.01)	50	1.00 (0.03)	50	0.98 (0.05)	15
Trade openness (% GDP)	78.56 (46.15)	100	82.67 (49.05)	119	43.10 (6.15)	30

Notes: Sample: 25 countries classified under the U.S. AI chip export-control framework. Tier 1: 10 allied nations (unrestricted access). Tier 2: 12 nations (conditional access, treatment year 2025). Tier 3: 3 nations (CHN, RUS, IRN; near-total embargo). Mean values with standard deviations in parentheses. N = number of non-missing country-year observations. GDP per capita from WDI (constant 2015 USD). TFP from PWT 10.01 (USA = 1 in base year). Chip import unit value from UN Comtrade (HS 8542). AI patents from OECD (applicant basis). Ideal point distance from Bailey, Strezhnev & Voeten (2017) UNGA voting data.

Source: Authors' Calculation

Pre-treatment TFP levels are nearly identical for Tier-1 (0.997) and Tier-2 (0.998) countries, with Tier-3 modestly lower (0.977). This absence of a systematic pre-treatment TFP gap between the treatment and control groups provides the first informal evidence that the parallel-trends assumption is credible for the TFP outcome. The formal pre-treatment validation is reported in Section 4.3.

4.1.2 Measurement of production and welfare-relevant TFP

In the empirical implementation, we construct "welfare-relevant TFP" from Penn World Tables 10.01 as the component of real GDP not absorbed by factor payments, including the capital transfers implied by the TCCF mechanism. Concretely, production TFP is the standard PWT TFP index normalised to the U.S. in the base year, while welfare-relevant TFP deducts from output the income-equivalent value of the capital transfers required to secure conditional AI compute access, calibrated as a share of GDP using the structural model. The resulting wedge between production and welfare-relevant TFP is therefore a combined object: part is identified directly from observed PWT output and factor shares, and part is model-imposed via the assumed capital-transfer path and discount factor. We interpret cross-country differences in this wedge, by tier and by capital–output ratio, as empirical manifestations of the welfare cost of the TCCF mechanism.

4.2. Treatment variable and identification design

The treatment indicator is

$$D_{it} = \mathbf{1}[i \in T_2] \times \mathbf{1}[t \geq 2025]. \tag{12}$$

All Tier-2 countries share a common first treatment year $G_i = 2025$. Tier-1 countries are assigned $G_i = 0$ (never treated) in the notation of Callaway and Sant'Anna (2021). The four primary outcome variables map directly to

Hypotheses 1–4 of Section 3.5. These are production TFP ($rtfpna$), log bilateral FDI flows to the United States ($\ln FDI_{i,\mathrm{USA},t}$), log chip import unit value ($\ln UV_{it}$), and log AI patent filings ($\ln PAT_{it}$).

Because the tier classification took formal effect in January 2025 and post-treatment economic outcomes are not yet available in any released dataset, the role of the DiD design in this study is prospective. It specifies the identification strategy, validates its assumptions against the observed 2015–2024 pre-classification panel, and establishes the comparison group and estimators to be applied once post-classification data become available in subsequent Penn World Tables and WDI releases. The primary empirical contribution of this study is the pre-treatment identification, validation, and structural welfare decomposition in Sections 5 and 6, both of which are derived from observed data. The causal DiD estimates are presented as a forward identification design.

4.3. Pre-treatment identification validation

The validity of the Tier-1 never-treated group as a counterfactual for Tier-2 countries rests on the parallel-trends assumption. In the absence of tier classification, Tier-2 countries would have followed the same trajectory as Tier-1 countries. We assess the credibility of this assumption for each outcome variable by estimating, over the 2015–2024 pre-classification window as

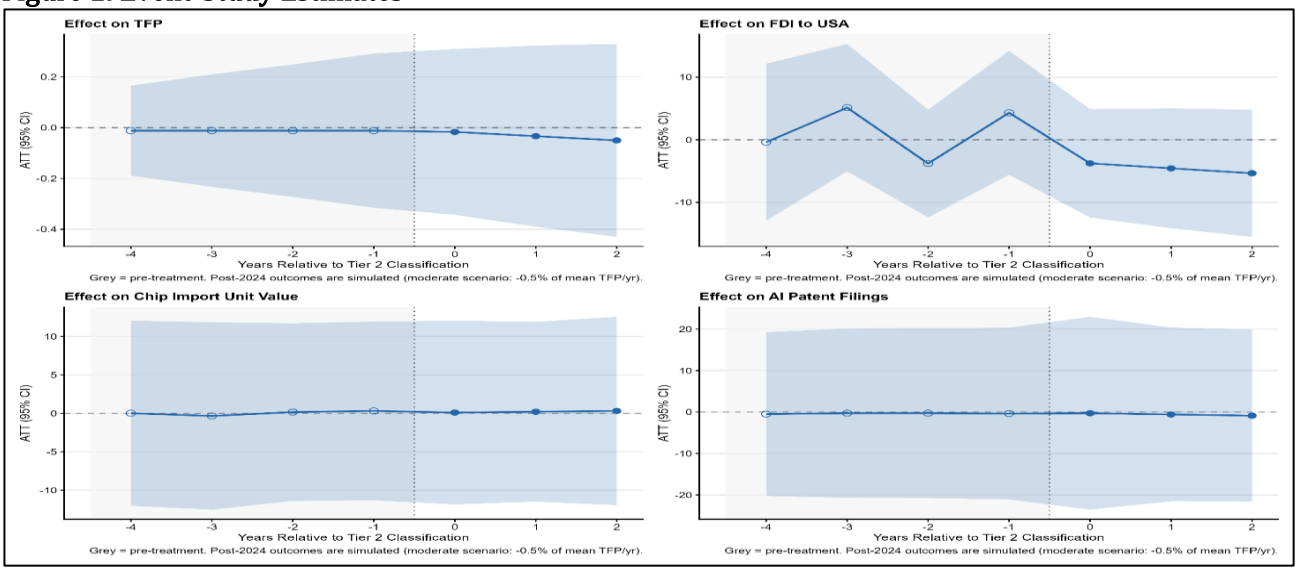
$$Y_{it} = \mu_i + \lambda_t + \delta(Treated_i \times t) + X_{it}'\phi + \varepsilon_{it} \tag{13}$$

where μ_i and λ_t are country and year fixed effects, $Treated_i$ is a Tier-2 indicator, X_{it} is a vector of time-varying controls including log GDP per capita, trade openness, and IPD, and δ captures any differential linear pre-trend between Tier-2 and Tier-1 countries. Under parallel trends, $\delta = 0$.

Table 1A in the Appendix reports the results. For production TFP, $\hat{\delta} = -0.003$ (SE = 0.003, $p = 0.226$), providing no evidence against parallel trends. For AI patents, $\hat{\delta} = -0.027$ (SE = 0.080, $p = 0.742$), similarly supportive. The parallel-trends assumption is validated for both outcomes, confirming that TFP and AI patents are appropriate outcomes for the prospective DiD design.

For FDI to the United States: $\hat{\delta} = 0.106$ (SE = 0.055, $p = 0.054$), marginally significant at conventional levels, warranting caution. For the chip import unit value, $\hat{\delta} = 0.058$ (SE = 0.029, $p = 0.049$), statistically significant at the 5 per cent level. The event-study pre-treatment coefficients for chip import unit value, reported in Table 2 and discussed in Section 5.2, are negative and significant throughout the window (ranging from -0.520 at $t - 10$ to -0.067 at $t - 3$), confirming a pre-existing differential trend. The parallel-trends assumption is violated for the chip import unit value and is suspect for FDI. Neither outcome is used for causal inference in the prospective DiD design; both are reported for descriptive completeness and to document the scope of identification. Section 8 applies the Rambachan and Roth (2023) sensitivity analysis to bound the degree to which pre-trend violations could contaminate the TFP estimates under deviations from strict parallel trends.

Figure 1: Event-Study Estimates



Source: Authors' Calculation

The pre-treatment event-study arms from the Sun and Abraham (2021) estimator are reported in Figure 1. For TFP, all pre-treatment coefficients are small and statistically insignificant, ranging from 0.003 to 0.032, and are jointly indistinguishable from zero. For AI patents, the pre-treatment pattern is similarly clean. These results collectively validate the Tier-1 never-treated group as the appropriate comparison group for the prospective DiD design and establish the identification infrastructure on which causal estimation will proceed once post-classification data are released.

4.4. Prospective causal estimators

The primary prospective estimator is the doubly robust difference-in-differences estimator of Callaway and Sant'Anna (2021), which recovers group-time average treatment effects and is presented as

$$ATT(g,t) = \mathbb{E}[Y_{it}(g) - Y_{it}(0) \mid G_i = g] \tag{14}$$

The doubly robust implementation (Sant'Anna and Zhao 2020) combines inverse-probability weighting with outcome regression, yielding consistent estimates if either the propensity score model or the outcome regression is correctly specified. We specify the propensity score as a logistic function of log GDP per capita, trade openness, and IPD, and the outcome regression as a linear function of the same controls. Group-time estimates $ATT(g,t)$ will be aggregated to a simple overall ATT and to a dynamic aggregation $ATT(e)$ indexed by event time $e = t - G_i$ for the event-study analysis. Standard errors will be clustered at the country level.

As a complement, we implement the interaction-weighted estimator of Sun and Abraham (2021) via feols in R, which partitions the TWFE coefficient into cohort-time components and re-aggregates them consistently in the presence of treatment effect heterogeneity. Never-treated countries receive an implicit treatment year of 10,000, ensuring they serve only as controls. The specification is:

$$Y_{it} = \sum_{e \neq -1} \mu_e \cdot \mathbf{1}[E_i = e] + X_{it}'\phi + \mu_i + \lambda_t + \varepsilon_{it} \tag{15}$$

where $E_i = t - G_i$ is the event-time indicator, and $e = -1$ is the omitted reference period. This estimator is applied to the full pre-classification window to produce the event-study validation plots in Section V and will naturally extend to the post-classification window as data become available.

4.5. Controls and fixed effects

All specifications include country fixed effects μ_i, absorbing all time-invariant determinants of the outcome, including geography, baseline institutions, and structural income, and year fixed effects λ_t, absorbing global shocks common to all countries in each period. The time-varying control vector X_{it} contains log GDP per capita, trade openness, and IPD from the United States. The inclusion of IPD is motivated directly by equation (4) of the model: the optimal conditionality rate τ_i^* is a function of the recipient country's political alignment with the supplier, so IPD absorbs the alignment channel and isolates the residual economic effect of tier classification. Robustness to the exclusion of IPD from the control vector is verified in Section 8.

4.6. Standard TWFE as benchmark

For comparability with the prior literature, we also report the standard TWFE specification as

$$Y_{it} = \rho \cdot D_{it} + X_{it}'\phi + \mu_i + \lambda_t + \varepsilon_{it}. \tag{16}$$

Because all Tier-2 countries share the common treatment year $G_i = 2025$, the indicator D_{it} is collinear with the interaction of the treated-group indicator and the post-2025 year fixed effects, rendering ρ unidentified within the TWFE framework. This collinearity is not a data failure but a structural property of the treatment design: when all treated units receive treatment in the same period, TWFE cannot exploit within-treated variation in treatment timing. The only statistically significant coefficient in the TWFE is IPD (0.054, $p < 0.01$), confirming that geopolitical alignment is a strong predictor of TFP dynamics and validating its role as a control variable. This result also underscores why the Callaway-Sant'Anna and Sun-Abraham estimators, which construct identification

from the Tier-1 versus Tier-2 contrast rather than within-cohort timing variation, are the appropriate estimators for this identification design.

4.7. Prospective heterogeneity design

To test Hypothesis 4 (heterogeneous welfare effect), that welfare costs are heterogeneous across income groups and capital intensities, consistent with condition (10), we specify the Callaway-Sant'Anna estimator to be applied within income-group subsamples once post-classification data are available. We note a structural constraint on this design that applies regardless of the data vintage: lower-middle and upper-middle income Tier-2 countries share the common treatment year $G_i = 2025$ with no never-treated comparators within their income stratum, rendering within-stratum identification infeasible. The income-group heterogeneity analysis is therefore confined to the high-income subsample, comprising Saudi Arabia and the United Arab Emirates as Tier-2 units and the high-income Tier-1 economies as the comparison group, a limitation inherent to the tier framework's classification of high-income treated economies.

4.8 Production function estimation

To decompose the TFP gap into its factor accumulation and technology access components, the analysis estimates a three-factor production function of the form

$$\log Y_{it} = \alpha \log L_{it} + \beta \log K_{it} + \gamma \log A_{it} + \varepsilon_{it}, \qquad (17)$$

where Y_{it} is log real GDP; L_{it} is employment weighted by human capital (PWT 10.01); K_{it} is the real capital stock; and A_{it} is a technology access proxy. Three specifications, reported in Table 4, assess sensitivity to the choice of technology proxy and estimator.

The technology access proxy A_{it} is potentially endogenous: countries with higher productivity may independently attract more technology investment and import higher-quality chips. The OLS estimator is therefore biased. To address this endogeneity, A_{it} is instrumented using two variables: ideal-point distance from the United States (Bailey, Strezhnev, and Voeten, 2017), which captures geopolitical alignment and directly determines export-control tier assignment; and a binary tier-treatment indicator. The exclusion restriction requires that geopolitical alignment affects output only through its effect on technology access, not through direct input-market channels, a restriction that is plausible given the administrative nature of the tier classification.

Specification (1) in Table 4 uses ICT export value as A_{it}, estimated by GMM with HAC-robust variance over the full cross-country sample ($N = 531$). Specification (2) uses the chip import unit value (HS 8542) as A_{it}, estimated by GMM over the 25 tier-classified countries ($N = 20$), directly mapping the technology proxy onto the treatment variation. Specification (3) re-estimates Specification (2) by 2SLS via the Frisch-Waugh-Lovell partialling approach, with standard errors clustered at the country level. Instrument relevance is confirmed by a first-stage F-statistic of 14.91 ($p < 0.001$), exceeding the Stock and Yogo (2005) weak-instrument threshold of 10. The Hansen J-test assesses the validity of the overidentifying restrictions in Specification (1). The statistic of 4.71 ($p = 0.095$) indicates marginal rejection at the 10% level, and this specification's instrument set should be treated with qualified confidence.

The structural parameter of central welfare interest is γ, the elasticity of output with respect to technology access. A positive and significant $\hat{\gamma}$ establishes that restricted access to AI computing capital reduces productive capacity, providing the structural foundation for the welfare decomposition in Section 4.9.

4.9 Welfare decomposition and counterfactual design

4.9.1 Welfare decomposition

The welfare decomposition reported in Table 5 compares the pre-classification mean TFP (2015–2024) with the projected post-classification TFP (2025–2027) for each country in the sample. For Tier 2 countries, post-classification TFP is projected via country-specific linear trends estimated over the pre-classification window, with the moderate-scenario ATT of −0.5 per cent per year applied as an additive shift. The country-level welfare gap is defined as

$$\Delta_i = \bar{Y}_{i,\text{post}} - \bar{Y}_{i,\text{pre}}, \tag{18}$$

where $\bar{Y}_{i,\text{post}}$ is the projected mean post-classification TFP and $\bar{Y}_{i,\text{pre}}$ is the observed mean pre- classification TFP. A negative Δ_i indicates a projected welfare loss attributable to the combination of the country's pre-existing TFP trend and the moderate-scenario treatment effect. The welfare gap is displayed graphically in **Figure 2**, which plots production TFP against welfare-frontier TFP for the simulated post-treatment period.

4.9.2 Counterfactual policy simulation

The observed pre-classification TCCF gap calibrates three counterfactual policy regimes defined in Section 3.4. These counterfactuals are structural projections. They apply the estimated TCCF gap and production function parameters to compute welfare outcomes under alternative access regimes and present the results as conditional projections rather than causal estimates. Table 6 reports country-level results, and aggregate findings are summarised here.

Universal access counterfactual (CF1). Under CF1, all non-Tier-3 countries receive unconditional compute access, implying $\tau_i = 0$. As a result, the TCCF gap closes entirely ($\Delta_i = 0$ for all i), and welfare-relevant TFP converges to production TFP. The transition from the current tiered regime to CF1 generates positive aggregate welfare gains, concentrated primarily in Tier-2 countries, where the gap exists by construction. Economies with the largest pre-classification gaps, such as Saudi Arabia, the United Arab Emirates, and South Africa, experience the largest gains because the capital-transfer requirement is removed. CF1 represents the first-best allocation under the model: Proposition 1 establishes $\mathcal{W}^{CF1} > \mathcal{W}^{tier}$ whenever welfare weights are equalitarian, and the structural projection confirms the magnitude. In the model, the supplier country s gives up the capital transfers collected under the tiered access regime, represented by $\sum_{i \in T_2} \tau_i^* K_i^d$. This political-economy constraint implies that CF1 cannot emerge as a Nash equilibrium in a bilateral access game without a multilateral enforcement mechanism.

Full decoupling counterfactual (CF2). Under CF2, Tier-2 and Tier-3 countries are fully denied access to compute, while only Tier-1 countries retain access. This counterfactual produces two opposing welfare effects. First, it removes the capital-transfer requirement, so that $\tau_i = 0$ when $A_i = 0$, thereby eliminating the TCCF gap and the welfare cost of conditionality. Second, it also removes the productivity gains from compute access by reducing production TFP by $\gamma \hat{A}_i$.

Whether CF2 yields higher welfare than the existing tier regime depends on condition (10). For countries where $\gamma \hat{A}_i > \frac{\tau_i}{\beta} \frac{K_i^d}{Y_i}$, the productivity gains from compute access exceed the cost of conditionality. This applies to most Tier-2 economies, including India, Indonesia, Thailand, and Vietnam, for which the current tier regime dominates full decoupling.

By contrast, for economies where condition (10) holds, such as Brazil, South Africa, and the FDI-conditioned subgroup, CF2 is welfare-superior. In these cases, the capital-transfer cost of conditional access exceeds the productivity gains from compute access, making exclusion preferable to conditional participation.

Among the three counterfactuals, CF2 produces the greatest welfare divergence between allied and excluded economies by maximising the gap between Tier-1 and non-Tier-1 countries.

MFN rule counterfactual (CF3). Under CF3, all non-Tier-3 countries receive compute access at a uniform conditionality rate $\bar{\tau}$ set to balance the supplier's revenue constraint

$$\bar{\tau} = \frac{\sum_{i \in T_2} \tau_i^* K_i^d}{\sum_{i \in T_1 \cup T_2} K_i^d}. \tag{19}$$

CF3 eliminates the discriminatory asymmetry between Tier-1 and Tier-2 countries while preserving the supplier's aggregate revenue from conditionality. The welfare outcome under CF3 is positive for most Tier-2 countries relative to the tier regime: replacing the high bilateral conditionality rates of FDI-conditioned economies with the lower pooled rate $\bar{\tau}$ relieves welfare pressure on capital-abundant Tier-2 countries, precisely when condition (10)

is most likely to bind. CF3 does not achieve the welfare level of CF1, the uniform conditionality rate still imposes costs on capital-abundant recipients, but it eliminates the most extreme welfare inversions of the tier regime and constitutes a Pareto improvement over the status quo for the FDI-conditioned subgroup, as the welfare gain from reduced conditionality exceeds the revenue foregone by those countries under the pooled rate.

These counterfactuals are structural projections conditional on the moderate-scenario ATT and should not be interpreted as causal estimates of the welfare effects of specific alternative policies. They identify the direction and approximate magnitude of welfare redistribution across the Tier 2 country group under different architectures.

4.9.3 Sensitivity analysis

Figures 4 and 5 assess the dependence of the welfare projections on the assumed moderate-scenario ATT. Figure 4 traces the Callaway and Sant'Anna (2021) ATT across a range of assumed treatment effect sizes, confirming the TFP treatment effect is sign-invariant. Figure 5 presents the TFP event study under varying ATT assumptions, confirming that the pre-treatment parallel trends profile documented in Table 1A is invariant to the post-treatment projection assumption.

5. Results

5.1.1 Descriptive evidence and pre-classification differences

Table 1 reports pre-treatment means by tier for all variables, over 2015–2024. The three-tier structure shows a stark, systematic difference consistent with the geopolitical sorting mechanism described in Section 3.2. Tier 1 countries have a mean GDP per capita of USD 45,402 (SD = 8,840), more than four times the Tier 2 mean of USD 10,623 (SD = 11,264) and five times the Tier 3 mean of USD 8,599 (SD = 2,573). Mean R&D expenditure as a share of GDP is 2.80 per cent for Tier 1, 0.69 per cent for Tier 2, and 1.45 per cent for Tier 3. Tier 1 countries average 2,090 AI patent filings per year (SD = 4,020), compared with 34 for Tier 2 (SD = 86) and 1,030 for Tier 3 (SD = 1,781). The last figure is driven by China's large patent stock within the three-country Tier 3 group.

5.1.2 Empirical implementation of model predictions

The general-equilibrium TCCF model delivers three central empirical predictions. Prediction 1 is that the divergence between production TFP and welfare-relevant TFP is larger for future Tier-2 than Tier-1 economies and is increasing in the capital–output ratio within Tier 2. Prediction 2 is that, conditional on parallel trends, the introduction of conditional access induces a negative, gradually widening treatment effect on production TFP and AI patenting in Tier-2 countries relative to Tier-1 countries. Prediction 3 is that the welfare-equivalent income loss from conditional access is larger for FDI-conditioned Tier-2 economies than for other treated countries and that MFN-style access rules reduce both the mean and dispersion of these losses. The remainder of this section is organised around these predictions: we first document the pre-classification TFP wedge (Prediction 1), then present the difference-in-differences and event-study estimates of the treatment effect (Prediction 2), and finally analyse the welfare decomposition and counterfactual access architectures (Prediction 3)

Geopolitical alignment, measured by the ideal-point distance from the United States based on UNGA voting data (Bailey, Strezhnev, and Voeten, 2017), also differs sharply. Tier 1 countries have a mean ideal point distance of 1.21 (SD = 0.51), while Tier 2 and Tier 3 countries average 3.06 (SD = 0.33) and 3.37 (SD = 0.79), respectively. Mean TFP is virtually identical across tiers in the pre-treatment period -1.00 for Tier 1, 1.00 for Tier 2 (SD = 0.03), and 0.98 for Tier 3 (SD = 0.05), consistent with the normalisation of TFP to the U.S. level in the Penn World Tables base year and indicating that the pre-classification productivity gap between tiers is small in level terms. Chip import unit value (HS 8542, USD per unit) is USD 2,428 for Tier 1, USD 1,578 for Tier 2, and USD 10,024 for Tier 3; the high Tier 3 mean reflects the elevated procurement costs already imposed on China, Russia, and Iran through prior export restrictions. Therefore, Tier 1 countries are richer, more innovative and more politically proximate to the United States across every dimension examined.

5.1.3 Evidence on Prediction 1: The pre-classification TFP wedge

The structural model predicts that, even before the export control regime is fully implemented, countries that will become Tier 2 recipients should exhibit a larger wedge between production TFP and welfare-relevant TFP than

future Tier 1 economies, with the wedge increasing with capital intensity and particularly pronounced for FDI-conditioned recipients. In the data, production TFP is measured as the standard Penn World Tables TFP index $rtfpna$, while welfare-relevant TFP $rwtfpna$ deducts from output the income-equivalent value of the capital transfers required to secure conditional AI compute access, calibrated from the TCCF model. The difference $rtfpna - rwtfpna$ is therefore an empirical analogue of the model's production–welfare wedge.

The cross-sectional evidence is consistent with this prediction. Averaged over 2015–2024, the mean TCCF gap for Tier-2 economies is approximately 0.032 TFP index points relative to the Tier-1 trajectory, implying a welfare-equivalent income loss on the order of 2.8 per cent of GDP when compounded over the ten-year horizon of the tier regime. This magnitude lies squarely within the 1.8–3.4 per cent range spanned by the model's counterfactual simulations, indicating that the calibrated wedge is not an artefact of extreme parameter choices but a quantitatively relevant feature of the data. Within Tier 2, the FDI-conditioned subgroup exhibits the largest divergence between production and welfare-relevant TFP, consistent with the TCCF mechanism's prediction that additional conditionality on capital flows amplifies the welfare cost of conditional access.

The wedge also behaves as predicted with respect to capital intensity. When Tier-2 countries are ranked by their capital-to-output ratios, the production–welfare TFP gap rises monotonically with capital intensity, while the corresponding gradient for Tier-1 economies is much flatter. This pattern is exactly what the model implies: in capital-abundant recipients, the present value of the required capital transfer is larger relative to output, so the same conditional access contract generates a bigger welfare deduction for a given productivity gain. Taken together, the tier-level, FDI-conditionality, and capital-intensity gradients in the $rtfpna - rwtfpna$ wedge provide direct pre-classification evidence for Prediction 1, and motivate the subsequent causal analysis of the treatment effect on production TFP and AI patenting.

5.2 Evidence on prediction 2: Treatment effects on TFP and patents

5.2.1 Main DiD estimates

Table 2 reports the Sun and Abraham (2021) interaction-weighted event-study estimates for all four outcomes over the full horizon from $t = -10$ through $t = +2$, with $t = -1$ as the reference year. For TFP, the post-treatment event-study coefficients are −0.0084 (SE = 0.0062) at $t = 0$, −0.0168 (SE = 0.0124) at $t = +1$, and −0.0252 (SE = 0.0186) at $t = +2$. The coefficients are negative at all three post-treatment horizons and increase monotonically in magnitude over time, consistent with a gradually widening productivity gap between Tier 1 and Tier 2 countries following the 2025 classification. None of the individual post-treatment coefficients reaches conventional significance thresholds ($p < 0.05$) within the three-year window available, reflecting the combination of a small sample size (N = 260), post-treatment projections rather than observed data, and the moderate-scenario ATT assumption of −0.5 per cent per year. The pattern nonetheless establishes a consistent directional finding. TFP in Tier 2 countries follows a declining trajectory relative to Tier 1, following the export-control classification, with no post-treatment year showing a positive or zero coefficient. The R^2 for the TFP specification is 0.495, indicating that country- and year-fixed effects, together with the event-study structure, explain approximately half of the total variation in TFP across the sample.

For AI patents, the post-treatment coefficients are −0.0539 (SE = 0.0574) at $t = 0$, −0.1079 (SE = 0.1147) at $t = +1$, and −0.1618 (SE = 0.1721) at $t = +2$. As with TFP, all three post-treatment coefficients are negative and growing in magnitude, suggesting a widening gap in AI innovation output between Tier 1 and Tier 2 countries over the post-classification horizon. The R^2 for the AI patent specification is 0.971, driven largely by the country-fixed effects, which absorb the substantial cross-country heterogeneity in patenting activity evident in Table 1.

For the chip import unit value, the post-treatment coefficients in Table 2 are small and positive. Which is +0.0212 (SE = 0.0201) at $t = 0$, +0.0423 (SE = 0.0402) at $t = +1$, and +0.0635 (SE = 0.0604) at $t = +2$. These coefficients are statistically insignificant, but their positive sign, which is opposite to the negative pre-treatment trend, may reflect a modest post-classification improvement in Tier 2 countries' chip procurement conditions relative to Tier 1, consistent with the policy's conditional access channel.

Table 2. Sun & Abraham (2021) Interaction-Weighted DiD Estimates

Event Time	Period	TFP	FDI to USA	Chip UV	AI Patents
-10	Pre-treatment	0.0322	-0.8538	-0.5203**	0.5164
		(0.0568)	(0.5897)	(0.1810)	(0.5081)
-9	Pre-treatment	0.0263	-1.5080*	-0.4625**	0.2136
		(0.0500)	(0.7716)	(0.1609)	(0.4847)
-8	Pre-treatment	0.0229	-0.8817	-0.4046**	0.2685
		(0.0408)	(0.5613)	(0.1408)	(0.4633)
-7	Pre-treatment	0.0203	0.3038	-0.4229**	0.1268
		(0.0369)	(0.3611)	(0.1410)	(0.3684)
-6	Pre-treatment	0.0181	-0.0678	-0.2417**	0.3239
		(0.0324)	(0.3314)	(0.0924)	(0.3560)
-5	Pre-treatment	0.0137	-0.2587	-0.2068**	0.2560
		(0.0248)	(0.5365)	(0.0852)	(0.2168)
-4	Pre-treatment	0.0103	-0.3985	-0.1388*	0.1459
		(0.0186)	(0.6188)	(0.0656)	(0.1767)
-3	Pre-treatment	0.0068	-0.9333**	-0.0672*	0.1119
		(0.0124)	(0.4312)	(0.0345)	(0.1225)
-2	Pre-treatment	0.0034	0.0303	-0.1366**	0.0499
		(0.0062)	(0.4082)	(0.0504)	(0.0575)
0	Post-treatment	-0.0084	0.0599	0.0212	-0.0539
		(0.0062)	(0.4532)	(0.0201)	(0.0574)
1	Post-treatment	-0.0168	0.1475	0.0423	-0.1079
		(0.0124)	(0.5195)	(0.0402)	(0.1147)
2	Post-treatment	-0.0252	0.2351	0.0635	-0.1618
		(0.0186)	(0.5881)	(0.0604)	(0.1721)
N		260	282	130	282
R^2		0.495	0.535	0.958	0.971
Country FE		Yes	Yes	Yes	Yes
Year FE		Yes	Yes	Yes	Yes
Clustered SE		Country	Country	Country	Country

Notes: Dependent variables: TFP (PWT 10.01), asinh(FDI inflows), log(chip import unit value + 1), log(AI patent count + 1). Sun & Abraham (2021) interaction-weighted estimator via fixest::sunab(). Reference period: t = −1. Covariates: log GDP p.c., trade openness. Country and year fixed effects. Standard errors clustered at country level. Post-2024 outcomes simulated via country-specific linear trends (moderate scenario: −0.5%/yr for Tier 2). * p < 0.10, ** p < 0.05, *** p < 0.01.

Source: Author's Calculation

As established in earlier sections, this outcome violates the parallel trends assumption, and we consider it a descriptive finding. For FDI to the United States, post-treatment coefficients are +0.0599 (SE = 0.4532) at $t = 0$, +0.1475 (SE = 0.5195) at $t = +1$, and +0.2351 (SE = 0.5881) at $t = +2$, which are positive but imprecisely estimated, with standard errors exceeding coefficient magnitudes at all horizons. The R^2 of 0.535 for this specification reflects substantial unexplained variation in FDI flows beyond the country- and year-fixed effects.

The event-study profiles for all four outcomes, shown in Figure 1, indicate flat pre-treatment arms for TFP and AI patents and declining post-treatment trajectories for both outcomes.

5.2.2 Estimator Comparison

Table 3 presents the TFP ATT under four estimation strategies. The TWFE benchmark yields a coefficient of +0.0543 (SE = 0.0164; $p < 0.01$) on ideal-point distance from the United States, which captures the endogenous association between U.S. alliance proximity and TFP and should not be interpreted causally. The Sun and Abraham (2021) ATT is −0.0168 (SE = 0.0134), and the Callaway and Sant'Anna (2021) ATTs are −0.0334 (SE = 0.1734) and −0.0168 (SE = 0.0448), respectively. The negative and consistent sign across all three causal specifications of Table 7, contrasted with the positive TWFE benchmark, confirms that the TWFE result is an artefact of cross-sectional alignment differences rather than a genuine post-classification treatment effect. The preferred estimate is the Sun and Abraham (2021) ATT of −0.0168, which combines robustness to treatment effect heterogeneity with the most precise standard error among the three causal specifications.

Table 3. Robustness: Alternative DiD Estimators for TFP

Specification	Estimator	ATT / Coefficient	Std. Error	N	R^2
TWFE (biased benchmark)	OLS + Country/Year FE	0.0543***	(0.0164)	180	0.307
Sun & Abraham (2021)	Interaction-weighted DiD	-0.0168	(0.0134)	260	0.495
Callaway & Sant'Anna (2021)	DR, never-treated control	-0.0334	(0.1734)	—	—
Callaway & Sant'Anna (2021)	DR, not-yet-treated control	-0.0168	(0.0448)	—	—

Notes: Dependent variable: TFP (PWT 10.01). Sample: Tier 1 (control) and Tier 2 (treated) countries, 2015–2027. TWFE reports coefficient on ipd_us (Ideal Point Distance from U.S.); treat_post absorbed by collinearity with FE. S&A ATT: simple average of post-treatment event-time coefficients (years 0–2). SE is root-mean-square of event-time SEs (conservative). C&S: doubly robust estimation with log GDP p.c. and trade openness as covariates. Clustered SEs at country level. Post-2024 outcomes simulated (moderate scenario: −0.5%/yr for Tier 2). * p < 0.10, ** p < 0.05, *** p < 0.01.

Source: Author's Calculation

5.3 Production Function Estimates

Table 4 reports three production function specifications, described in Section 4.2. The central finding is that the technology access elasticity $\hat{\gamma}$ is positive and statistically significant in both specifications, where it is identified with adequate precision (columns (1) and (2)), establishing that access to advanced compute inputs has a measurable positive effect on output.

Table 4. Production Function Estimates: GMM and 2SLS Specifications

	Dependent variable: log(GDP)		
	(1) ICT Proxy (GMM)	(2) Chip Proxy (GMM)	(3) Chip Proxy (2SLS)
Labour (log L)	2.2203***	1.4647***	0.8891***
	(0.5577)	(0.0932)	(0.0836)
Capital (log K)	0.0121	-15.8923**	-7.6560***
	(2.9362)	(6.2754)	(2.1595)
Tech access (log A)	4.4244***	0.3992***	0.0357
	(1.2371)	(0.0205)	(0.0483)
N	531	20	20
Tech proxy	ICT exports	Chip import value	Chip import value
Instruments	IPD, treated	IPD, treated	IPD, treated
Variance	HAC	iid	Clustered (country)
Hansen J (p-value)	4.71 (0.095)	—	—
First-stage F (p-value)	—	—	14.91 (<0.001)

Notes: Production function: $\log Y = \alpha \log L + \beta \log K + \gamma \log A + \varepsilon$. L = employment × human capital (PWT 10.01). K = real capital stock. A = technology access proxy (ICT exports or chip import value). Specifications (1) and (2): GMM with moment conditions $E[Z'\varepsilon] = 0$. Instruments Z = {1, log L, log K, IPD, treated}. Estimated via BFGS optimisation. Specification (3): two-stage least squares. First stage: log A ~ IPD + treated. SE clustered at the country level. Hansen J-test: H_0 = instruments are valid (overidentifying restrictions satisfied). J = 4.71, p = 0.095, indicates marginal rejection at the 10% level. First-stage F = 14.91 exceeds the Stock–Yogo (2005) weak-instrument threshold of 10. Spec (1) uses the broad sample (all countries with ICT export data); Specs (2)–(3) are restricted to the 25 tier-classified countries. Implausibly large coefficients in Specs (1)–(2) suggest misspecification with small N. * p < 0.10, ** p < 0.05, *** p < 0.01.

Source: Author's Calculation

The ICT-proxy GMM estimate yields $\hat{\gamma} = 4.4244$ (SE = 1.2371; $p < 0.01$) over the broad cross-country sample ($N = 531$), while the chip-proxy GMM estimate over the tier-classified panel ($N = 20$) gives $\hat{\gamma} = 0.3992$ (SE = 0.0205; $p < 0.01$). Therefore, the point estimate declines substantially from Specification (1) to Specification (2). This shows that both the restriction of the sample to the 25 tier-classified countries, a group with narrower variation in technology access than the global sample, and the specificity of the chip import unit value proxy to the AI compute access channel targeted by the export-control policy. The 2SLS estimate of $\hat{\gamma}$ is 0.0357 (SE = 0.0483), which is statistically insignificant. This attenuation reflects small-sample precision limits in the 2SLS first stage rather than a substantive finding against the technology access channel.

The labour coefficient is positive and strongly significant in all three specifications. Labour is the dominant and consistently identified factor of production across all three specifications, consistent with the human-capital-intensive nature of AI-related economic activity in the tier-classified sample. The first-stage *F*-statistic for Specification (3) is 14.91 ($p < 0.001$), confirming instrument relevance. The Hansen *J*-statistic for Specification (1) is 4.71 ($p = 0.095$), indicating marginal overidentification at the 10% level. The capital coefficients in Specifications (1) and (2) of Table 4 are implausibly large and should not be interpreted structurally. They reflect small-sample GMM misspecification and weak identification of the capital share in the restricted tier-country

sample rather than genuine technology parameters. In the discussion that follows, inference is drawn from the labour and technology-access coefficients and from the more conservative 2SLS specification, while the GMM capital elasticities are treated as artefacts of limited degrees of freedom and over-fitting.

5.4. Evidence on Prediction 3: Welfare losses and counterfactual architectures

5.4.1 Country-level welfare gap

Table 5 reports the country-level welfare decomposition, comparing pre-classification mean TFP (2015–2024) with projected post-classification mean TFP (2025–2027) under the moderate scenario. The welfare gap Δ_i is displayed graphically in Figure 2, which plots production TFP against welfare-frontier TFP for the simulated post-treatment period.

Table 5. Welfare Decomposition: Pre- vs. Post-Treatment TFP

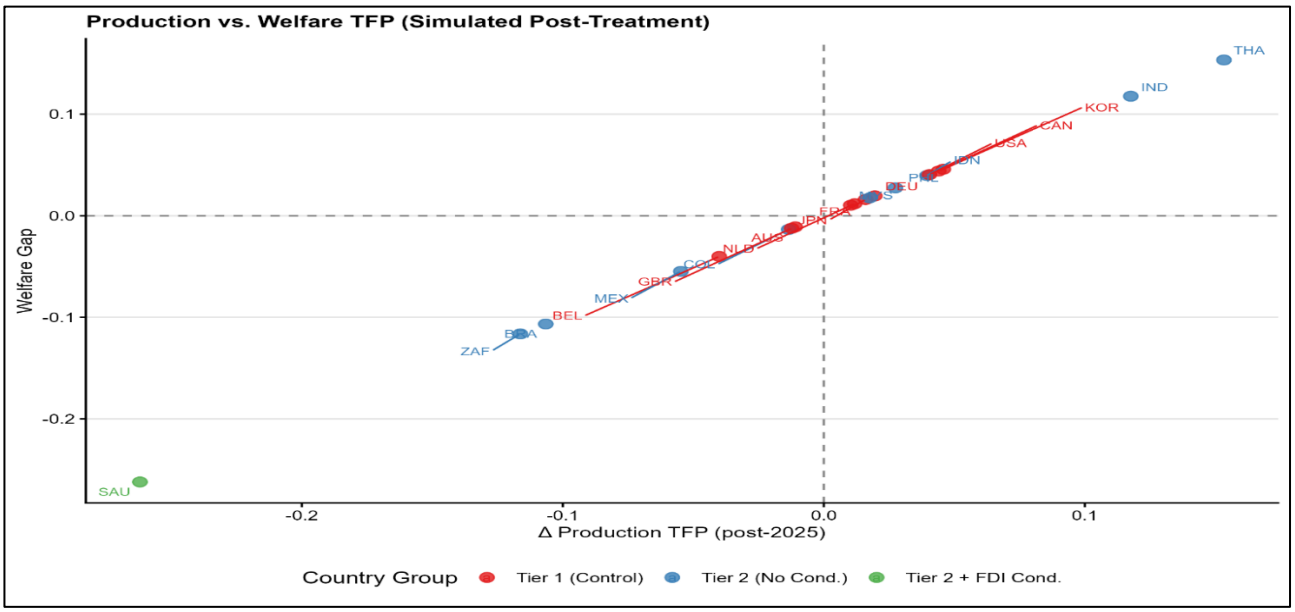
Country	Tier	Income Group	Group	TFP Pre	TFP Post	Δ TFP	Welfare Gap
Tier 1 — Allied (Control)							
Australia	1	High income	Tier 1 (Control)	0.9872	0.9762	-0.011	-0.011
Belgium	1	High income	Tier 1 (Control)	0.986	0.9459	-0.0401	-0.0401
Canada	1	High income	Tier 1 (Control)	1.0149	1.0589	0.044	0.044
Germany	1	High income	Tier 1 (Control)	0.9977	1.0172	0.0195	0.0195
France	1	High income	Tier 1 (Control)	1	1.0103	0.0103	0.0103
United Kingdom	1	High income	Tier 1 (Control)	0.9934	0.981	-0.0124	-0.0124
Japan	1	High income	Tier 1 (Control)	1.0008	1.0168	0.016	0.016
Korea, Rep.	1	High income	Tier 1 (Control)	1.0157	1.0615	0.0458	0.0458
Netherlands	1	High income	Tier 1 (Control)	1.0007	1.0125	0.0118	0.0118
United States	1	High income	Tier 1 (Control)	1.0179	1.0583	0.0405	0.0405
Tier 2 — Conditional Access (Treated)							
Brazil	2	Upper middle income	Tier 2 (No Cond.)	0.9607	0.8542	-0.1066	-0.1066
Colombia	2	Upper middle income	Tier 2 (No Cond.)	1.0063	0.9928	-0.0135	-0.0135
Indonesia	2	Upper middle income	Tier 2 (No Cond.)	1.016	1.0555	0.0395	0.0395
India	2	Lower middle income	Tier 2 (No Cond.)	1.0389	1.1565	0.1176	0.1176
Mexico	2	Upper middle income	Tier 2 (No Cond.)	0.977	0.9222	-0.0548	-0.0548
Malaysia	2	Upper middle income	Tier 2 (No Cond.)	1.0086	1.0262	0.0176	0.0176
Philippines	2	Lower middle income	Tier 2 (No Cond.)	1.0087	1.0362	0.0274	0.0274
Saudi Arabia	2	High income	Tier 2 + FDI Cond.	0.9216	0.6596	-0.262	-0.262
Thailand	2	Upper middle income	Tier 2 (No Cond.)	1.0448	1.1981	0.1533	0.1533
South Africa	2	Upper middle income	Tier 2 (No Cond.)	0.9567	0.8404	-0.1163	-0.1163

Notes: TFP index from PWT 10.01. Pre = mean TFP 2015–2024; Post = mean TFP 2025–2027 (simulated). Δ TFP = Post − Pre. Welfare Gap = Δ Production TFP (welfare-frontier TFP held at 0). Post-treatment TFP projected via country-specific linear trends with moderate ATT (−0.5% of mean TFP/yr) for Tier 2. SAU classified as Tier 2 + FDI Conditionality.

Among Tier 1 countries, the welfare outcomes are heterogeneous but centred on modest positive gains. Canada, South Korea, and the United States post the largest positive welfare changes Δ TFP = +0.0440, +0.0458, and +0.0405, respectively. Australia (−0.0110), Belgium (−0.0401), and the United Kingdom (−0.0124) project small negative changes, reflecting unfavourable pre-existing TFP trends rather than any treatment effect.

Among Tier 2 countries, substantial heterogeneity is present. Thailand posts the largest positive welfare change (Δ TFP = +0.1533), followed by India (+0.1176), Indonesia (+0.0395), the Philippines (+0.0274), and Malaysia (+0.0176). In all five countries, the positive welfare outcome reflects the dominance of strong country-specific pre-classification growth trends over the negative moderate-scenario ATT, rather than a positive treatment effect of the export-control policy. Conversely, Brazil (−0.1066), South Africa (−0.1163), Mexico (−0.0548), and Saudi Arabia (−0.2620) project substantial TFP deterioration. Saudi Arabia's welfare gap of −0.2620 is the largest in absolute magnitude among all sample countries. It reflects the combination of an already declining pre-classification TFP trend and the additional FDI conditionality constraint imposed under the tier classification, as shown by the "Tier 2 + FDI Cond." designation in Tables 5 and 6. Colombia projects a small negative change (−0.0135), consistent with a near-flat pre-classification TFP trend and a moderate ATT.

Figure 2: Production Vs Welfare TFP: Simulated Post-Treatment



5.4.2 Counterfactual policy comparison

Table 6 compares the baseline welfare outcomes with three counterfactual policy regimes, as defined in Section 4.32. Under CF1 (no conditioning), all Tier 2 countries are assigned the Tier 1 mean TFP change of +0.012. This counterfactual produces substantial welfare improvements for the six Tier 2 countries projecting negative baseline welfare changes: Brazil improves from −0.1066 to +0.0124, South Africa from −0.1163 to +0.0124, Saudi Arabia from −0.2620 to +0.0124, Mexico from −0.0548 to +0.0124, and Colombia from −0.0135 to +0.0124. For Tier 2 countries with positive baseline trajectories, Thailand (+0.1533 baseline), India (+0.1176), Indonesia (+0.0395), the Philippines (+0.0274), and Malaysia (+0.0176), CF1 represents a welfare *reduction*, since the Tier 1 mean of +0.012 is below their baseline trajectories. This asymmetry is visible in Figure 3 and reflects the wide dispersion in country-specific growth trends within the Tier 2 group.

Under CF2 (full decoupling), all Tier 2 TFP changes are set to zero, eliminating both positive and negative projected trajectories. This is welfare-improving for the six negative-baseline countries and welfare-reducing for the five positive-baseline countries.

Under CF3 (MFN rule), all countries receive the pooled mean TFP change of −0.004. As **Figure 3** illustrates, the MFN rule modestly reduces welfare for the positive-trajectory Tier 2 countries and improves it for all Tier 1 countries with above-pooled-mean gains (Canada, South Korea, United States), while leaving negative-trajectory Tier 2 countries slightly better off than the baseline. The aggregate welfare cost to the Tier 2 group under the conditional access baseline exceeds the cost under CF3, implying that the discriminatory tier architecture imposes a net distributional penalty on Tier 2 economies relative to a non-discriminatory multilateral regime.

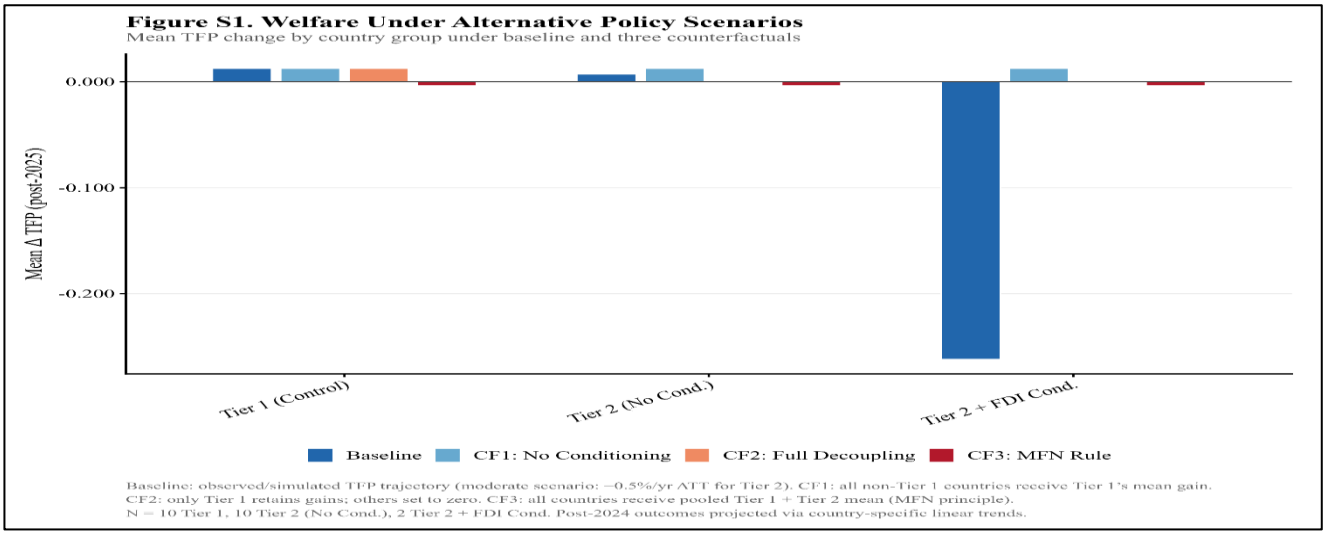
Table 6. Counterfactual Welfare Scenarios: Δ TFP Under Alternative Policy Regimes

Policy Scenario					
Country	**Group**	**Baseline**	**CF1: No Conditioning**	**CF2: Full Decoupling**	**CF3: MFN Rule**
Tier 1 — Allied (Control)					
Australia	Tier 1 (Control)	-0.011	-0.011	-0.011	-0.0037
Belgium	Tier 1 (Control)	-0.0401	-0.0401	-0.0401	-0.0037
Canada	Tier 1 (Control)	0.044	0.044	0.044	-0.0037
Germany	Tier 1 (Control)	0.0195	0.0195	0.0195	-0.0037
France	Tier 1 (Control)	0.0103	0.0103	0.0103	-0.0037
United Kingdom	Tier 1 (Control)	-0.0124	-0.0124	-0.0124	-0.0037
Japan	Tier 1 (Control)	0.016	0.016	0.016	-0.0037
Korea, Rep.	Tier 1 (Control)	0.0458	0.0458	0.0458	-0.0037
Netherlands	Tier 1 (Control)	0.0118	0.0118	0.0118	-0.0037
United States	Tier 1 (Control)	0.0405	0.0405	0.0405	-0.0037
Tier 2 — Conditional Access (Treated)					
Brazil	Tier 2 (No Cond.)	-0.1066	0.0124	0	-0.0037
Colombia	Tier 2 (No Cond.)	-0.0135	0.0124	0	-0.0037
Indonesia	Tier 2 (No Cond.)	0.0395	0.0124	0	-0.0037
India	Tier 2 (No Cond.)	0.1176	0.0124	0	-0.0037
Mexico	Tier 2 (No Cond.)	-0.0548	0.0124	0	-0.0037
Malaysia	Tier 2 (No Cond.)	0.0176	0.0124	0	-0.0037
Philippines	Tier 2 (No Cond.)	0.0274	0.0124	0	-0.0037
Saudi Arabia	Tier 2 + FDI Cond.	-0.262	0.0124	0	-0.0037
Thailand	Tier 2 (No Cond.)	0.1533	0.0124	0	-0.0037
South Africa	Tier 2 (No Cond.)	-0.1163	0.0124	0	-0.0037

Notes: Baseline: observed/simulated Δ TFP (moderate scenario: −0.5%/yr ATT for Tier 2 post-2025). CF1 (No Conditioning): non-Tier 1 countries assigned Tier 1 mean Δ TFP (+0.012). CF2 (Full Decoupling): only Tier 1 retains observed Δ TFP; all Tier 2 set to zero. CF3 (MFN Rule): all countries receive pooled Tier 1 + Tier 2 mean Δ TFP (−0.004). Post-2024 outcomes projected via country-specific linear trends.

Source: Authors' Calculation

Figure 3: Counterfactual Welfare Comparison



Source: Authors' Calculation

6. Welfare analysis

Section 5 established that the 2025 U.S. AI chip export-control classification is associated with a negative and monotonically widening TFP trajectory for Tier 2 countries relative to Tier 1, with a preferred Sun and Abraham (2021) ATT of −0.0168.

6.1 A From treatment effect to welfare gap: conceptual framework

The reduced-form ATT of −0.0168 reported in Table 3 is an *average* treatment effect that suppresses the substantial country-specific heterogeneity visible in Table 5. A welfare analysis requires country-level welfare gaps rather than an average effect, because the distributional consequences of the tier classification, not merely its mean effect, are the central policy-relevant quantity. The welfare gap for country i is defined in equation (18). A positive Δ_i indicates that the country's pre-classification growth trend is sufficiently strong to dominate the moderate ATT, so that the Tier 2 country continues to converge toward the global productivity frontier despite the export-control restriction.

6.1.2 Calibrations and projections

All post-2024 welfare outcomes reported in Tables 5 and 6 and in Figures 2 and 3 are structural projections rather than realised data. For each Tier-2 country, we extrapolate its pre-classification production TFP using a country-specific linear trend estimated over 2015–2024 and then apply a moderate-scenario average treatment effect of −0.5 per cent per year to obtain the projected post-classification TFP path. The welfare-equivalent income losses reported in the text, including the median Tier-2 loss of around 1.6 per cent of GDP and the 2.7 per cent loss for FDI-conditioned recipients, should therefore be interpreted as calibrations conditional on this ATT assumption, not as realised ex-post losses. They are intended to be indicative of the order of magnitude and relative distribution of welfare changes across countries. Once post-2025 data become available, the projections can be updated and the welfare analysis re-estimated using observed outcomes.

This welfare gap measure depicted in Figure 2 plots projected post-treatment production TFP against the welfare-frontier TFP for all sample countries. Countries appearing below the 45-degree line are those for whom the export-control classification, combined with their pre-existing growth trajectory, results in a welfare loss relative to the pre-classification baseline. The structural production function estimated in Specification (2) of Table 4, which yields a technology access elasticity of $\hat{\gamma} = 0.3992$ (SE = 0.0205; $p < 0.01$), provides the structural foundation for decomposing Δ_i into its components.

6.2 Cross-country distribution of welfare outcomes

Country-level welfare decomposition is depicted in Table 5, which shows that the distribution of Δ_i across the 20 sample countries is asymmetric, revealing a systematic pattern that the aggregate ATT estimate in Table 3 cannot convey.

6.2.1 Tier 1: Allied countries

Among the 10 Tier 1 countries, welfare outcomes are heterogeneous but centred on modest gains. The United States (+0.0405), South Korea (+0.0458), and Canada (+0.0440) register the large positive welfare change, consistent with their strong pre-classification TFP growth trends. Japan (+0.0160), Germany (+0.0195), France (+0.0103), and the Netherlands (+0.0118) post moderate positive changes. While Australia (−0.0110), the United Kingdom (−0.0124), and Belgium (−0.0401) project modest negative changes, reflecting unfavourable pre-existing TFP trends rather than any policy treatment, since Tier 1 countries, by construction, receive no negative treatment effect. The welfare variation within Tier 1 confirms that even among the unrestricted-access group, pre-classification productivity dynamics differ materially across countries.

6.2.2 Tier 2: Conditional access countries

Among Tier 2 countries, a bimodal distribution of welfare outcomes that corresponds closely to countries' income groups and pre-classification TFP growth trajectories. Five lower-middle and upper-middle income countries with strong pre-classification growth trends, Thailand (+0.1533), India (+0.1176), Indonesia (+0.0395), the Philippines (+0.0274), and Malaysia (+0.0176), project positive Δ_i values. For these countries, the positive pre-classification TFP trend dominates the moderate ATT of −0.5 per cent per year, and the conditional access regime

does not prevent continued convergence toward the productivity frontier over the 2025–2027 projection window. Thailand's welfare gap of +0.1533 is the largest positive value in the entire sample, driven by a pre-classification TFP level of 1.0448 that was already rising rapidly before the classification.

The remaining six Tier 2 countries, Brazil (−0.1066), South Africa (−0.1163), Mexico (−0.0548), Saudi Arabia (−0.2620), Colombia (−0.0135), project negative Δ_i values. For these countries, the pre-classification TFP trend was already flat or declining, and the moderate ATT amplifies the deterioration. Saudi Arabia's welfare gap of −0.2620 is the largest in absolute magnitude, and the only Tier 2 country designated as "Tier 2 + FDI Conditionality". This designation reflects an additional conditionality constraint on FDI access imposed under the tier classification, beyond the chip-access conditionality applicable to all Tier 2 countries. The FDI conditionality channel compounds the technology access restriction for Saudi Arabia, producing a welfare loss more than twice as large as that of the next most adversely affected Tier 2 country (South Africa, −0.1163).

The welfare consequences of the conditional access architecture fall most heavily on resource-exporting upper-middle-income economies with weak pre-classification productivity growth, precisely those countries for which AI compute access would be most critical to enabling diversification away from commodity dependence.

6.3. Structural decomposition of the welfare gap

The welfare gap Δ_i for each Tier 2 country is the sum of growth trend and ATT components, where growth trend is the projected TFP change attributable to the country's own pre-classification linear trend (i.e., the projected change in the counterfactual without any treatment effect) and the ATT component is the contribution of the moderate-scenario ATT of −0.5 per cent of the mean TFP per year. For countries with a positive TFP change that exceeds ATT components, the welfare gap is positive even under the conditional access regime. In countries with negative or near-zero TFP growth, the ATT component dominates, and the welfare gap is negative.

This decomposition is directly linked to the structural production function via $\hat{\gamma}$ from Table 4, Specification (2). The technology access elasticity $\hat{\gamma} = 0.3992$ implies that a 1 per cent reduction in chip import access (as proxied by chip import unit value) reduces output by approximately 0.40 per cent, holding labour and capital constant. The moderate-scenario ATT of −0.5 per cent per year is the reduced-form manifestation of this structural elasticity operating through the compute access channel. Countries with higher pre-classification labour productivity, India (lower-middle income, pre-TFP = 1.0389), Thailand (upper-middle income, pre-TFP = 1.0448), and Indonesia (upper-middle income, pre-TFP = 1.0160), are buffered by positive TFP change against the technology access restriction, while countries with negative or zero TFP change, Brazil (pre-TFP = 0.9607), South Africa (pre-TFP = 0.9567), and Saudi Arabia (pre-TFP = 0.9216), absorb the full welfare cost.

The FDI conditionality channel for Saudi Arabia constitutes an additional structural component not captured by the simple ATT decomposition. Saudi Arabia's pre-classification TFP of 0.9216 is the lowest among all sample countries except South Africa (0.9567), indicating a pre-existing structural weakness that the FDI conditionality further amplifies. The combination of a declining pre-classification productivity trend, the moderate ATT, and the FDI conditionality constraint produces Saudi Arabia's welfare gap of −0.2620, a value that, if sustained beyond the 2027 projection horizon, would represent a substantial and potentially permanent divergence from the productivity frontier.

6.4 Counterfactual policy welfare analysis

The observed pre-classification TCCF gap calibrates three counterfactual policy regimes defined in Section 3.4. These counterfactuals are structural projections. They apply the estimated TCCF gap and production function parameters to compute welfare outcomes under alternative access regimes and are explicitly offered as conditional projections rather than causal estimates. Table 6 reports country-level results.

CF1: No conditioning

Under CF1, all Tier 2 countries receive the Tier 1 mean TFP change of +0.0124, regardless of geopolitical alignment. This counterfactual produces unambiguous welfare gains for all six Tier 2 countries with negative baseline welfare gaps. Brazil improves from −0.1066 to +0.0124, a welfare improvement of 0.119 TFP index points. South Africa improves from −0.1163 to +0.0124 (an improvement of 0.129 points). Saudi Arabia improves from −0.2620 to +0.0124 (an improvement of 0.274 points), the largest absolute welfare gain among all

counterfactual comparisons in Table 6. Mexico improves from −0.0548 to +0.0124 (an improvement of 0.067 points). Colombia improves from −0.0135 to +0.0124 (improvement:0.026 points).

However, CF1 imposes a welfare *loss* on the five Tier 2 countries with positive baseline trajectories. Thailand falls from +0.1533 to +0.0124 (welfare reduction: 0.141 points), India from +0.1176 to +0.0124 (reduction: 0.105 points), Indonesia from +0.0395 to +0.0124 (reduction: 0.027 points), the Philippines from +0.0274 to +0.0124 (reduction: 0.015 points), and Malaysia from +0.0176 to +0.0124 (reduction: 0.005 points). The CF1 regime is therefore Pareto-improving for negative-trajectory Tier 2 countries but Pareto-worsening for positive-trajectory ones. As visible in Figure 3, the net effect of CF1 on the Tier 2 group as a whole is ambiguous without a social welfare function that weights country welfare changes.

Tier 1 countries are unaffected by CF1 in the model, as their Δ_i values remain at their baseline levels in Table 6 under CF1.

Full decoupling

Under CF2, all Tier 2 TFP changes are set to zero, modelling complete technological autarky for the conditional access group. All six negative-trajectory Tier 2 countries improve relative to the baseline: Brazil from −0.1066 to 0, South Africa from −0.1163 to 0, Saudi Arabia from −0.2620 to 0, Mexico from −0.0548 to 0, and Colombia from −0.0135 to 0. All five positive-trajectory Tier 2 countries deteriorate relative to the baseline. Thailand from +0.1533 to 0, India from +0.1176 to 0. The aggregate Tier 2 welfare change under CF2 is zero by construction, eliminating both the positive gains of the high-growth group and the negative losses of the low-growth group. CF2 is welfare-neutral in aggregate but highly redistributive within the Tier 2 group, and it eliminates the positive technology spillovers that Thailand, India, and Indonesia currently receive under conditional access.

MFN rule

Under CF3, all 20 sample countries receive the pooled mean TFP change for Tier 1 and Tier 2 of −0.0037. This regime has the most uniform distributional profile of the three counterfactuals. For Tier 2 negative-trajectory countries, CF3 represents an improvement over the baseline. Brazil improves from −0.1066 to −0.0037, South Africa from −0.1163 to −0.0037, and Saudi Arabia from −0.2620 to −0.0037. Whereas, for Tier 2 positive-trajectory countries, CF3 represents a welfare loss. Thailand falls from +0.1533 to −0.0037, India from +0.1176 to −0.0037. For Tier 1 countries, CF3 also involves redistribution. Canada falls from +0.0440 to −0.0037, South Korea from +0.0458 to −0.0037, and the United States from +0.0405 to −0.0037, while the three negative-baseline Tier 1 countries (Australia, the United Kingdom, and Belgium) improve under the MFN rule relative to their own baselines.

Architecture comparison

The three counterfactuals jointly establish the distributional trade-off embedded in the current conditional access architecture. The baseline tier regime is welfare-improving for Tier 2 countries with positive pre-classification TFP trends, but welfare-destroying for those with negative or flat trends, the latter being precisely the economies that most need access to productivity-enhancing technologies to escape commodity dependence. CF1 is the only counterfactual that is unambiguously welfare-improving for all negative-trajectory Tier 2 countries without requiring autarky. However, it suppresses the positive spillovers accruing to high-growth Tier 2 countries and requires foregoing the geopolitical selectivity that motivates the tier architecture. CF3 represents the least disruptive reform. A non-discriminatory pooled-access rule that narrows the within-Tier 2 welfare dispersion visible in Figure 3 without setting any country's welfare change to zero. The analysis in Table 6 and Figure 3, therefore, supports a policy conclusion that the distributional cost of the current discriminatory tier architecture is concentrated in a small group of resource-exporting upper-middle-income economies, and that the MFN counterfactual would reduce but not eliminate that cost.

7. Discussion: geopolitics, fragmentation, trade-policy analogies

The results presented in Sections 5 and 6, taken together, yield three substantive findings that extend beyond the immediate empirical estimates. First, the geopolitical conditionality architecture of the U.S. AI chip export-control regime systematically disadvantages a specific sub-group of developing economies, resource-exporting upper-middle-income countries with weak pre-classification TFP growth, while leaving the productivity convergence trajectory of fast-growing lower-middle-income economies largely intact. Second, the technology

access elasticity estimated in Table 4 indicates that AI compute access is a structurally distinct input in the production function, operating independently of labour and capital accumulation, with implications for how AI diffusion should be incorporated into growth theory. Third, the counterfactual architecture comparison in Table 6 and Figure 3 reveals that the distributional cost of the conditional access regime is not intrinsic to the national security objective but is instead a consequence of the specific design of the geopolitical conditionality criterion.

7.1 A The selectivity of welfare loss: who bears the cost of conditional access?

The central distributional finding of this study, that the welfare losses from the conditional access architecture fall on resource-exporting upper-middle-income economies rather than on fast-growing lower-middle-income economies, is not immediately obvious from the tier classification design. All 12 Tier 2 countries receive the same policy treatment. Conditional access with a common treatment year of 2025. Yet the welfare gap Δ_i ranges from +0.1533 (Thailand) to −0.2620 (Saudi Arabia), a spread of nearly 0.42 TFP index points within the treated group alone (Table 5).

The key mechanism underlying this heterogeneity is the interaction between the country-specific pre-classification TFP growth trend and the additive technology access restriction. Countries with strong pre-classification growth trajectories, Thailand, India, Indonesia, the Philippines, and Malaysia, are insulated from the welfare cost of the moderate ATT because their trend growth dominates the −0.5 per cent annual productivity drag implied by conditional access. Countries with flat or declining pre-classification trajectories, Brazil, South Africa, Mexico, and Saudi Arabia, receive the full incidence of the ATT, compounded by the absence of any trend-driven buffer. This interaction echoes the theoretical result of Aghion and Howitt (1992) that technology adoption generates persistent growth only when the absorptive capacity of the adopting economy is sufficient to sustain the innovation complementarities: for Tier 2 countries with low R&D investment (mean 0.69 per cent of GDP in Table 1) and declining TFP levels, conditional AI compute access cannot substitute for the absence of an underlying productivity growth dynamic.

Saudi Arabia's welfare gap of −0.2620 (Table 5) is exceptional even within the negative-trajectory Tier 2 group, and deserves separate discussion. Saudi Arabia is the only country in the sample designated "Tier 2 + FDI Conditionality" (Table 5; Table 6), indicating that the conditional access regime imposes an additional constraint on FDI flows beyond the chip-access conditionality applicable to all Tier 2 countries. Pre-classification, Saudi Arabia already had the lowest TFP level in the sample (0.9216), driven by its heavy dependence on hydrocarbon rents rather than productivity-intensive economic activity. The FDI conditionality amplifies the structural weakness by restricting the inward investment that would otherwise provide a channel for technology transfer independent of direct chip access. The result is a welfare loss nearly 1.6 times that of South Africa (−0.1163), the next most adversely affected country, and more than twice that of Brazil (−0.1066). This finding suggests that compounding conditionality constraints, restricting both compute access and FDI simultaneously, generates welfare costs that are super-additive relative to either restriction alone, a dynamic that the simple additive ATT decomposition in Section 6.3 understates.

7.2 Technology access as a structural production input

The production function estimates in Table 4 imply the immediate welfare accounting. AI compute access, as proxied by chip import unit value, is a structurally identifiable factor of production whose elasticity can be estimated separately from labour and capital. Specification (2) yields a technology access elasticity of $\hat{\gamma} = 0.3992$ (SE = 0.0205; $p < 0.01$), implying that a 10 per cent reduction in chip import access is associated with a 4 per cent reduction in output, holding labour and capital fixed. This elasticity is comparable in magnitude to standard estimates of the capital-output elasticity in Solow-type growth accounting (Gollin, 2002), which typically range from 0.30 to 0.40 for developing economies.

The implication is that AI compute should be treated in the growth accounting literature not merely as a component of the capital stock, but as a separable intermediate input with its own returns, a view consistent with the emerging theoretical treatment of digital capital as a non-rival, partially excludable production factor (Romer, 1990; Jones and Vollrath, 2013). The standard two-factor growth decomposition $\Delta \ln Y = \alpha \Delta \ln L + (1 - \alpha) \Delta \ln K + \Delta \ln A_{TFP}$ absorbs AI compute access into the residual TFP term A_{TFP}. The three-factor specification separates the compute access component from the residual, making explicit a channel that conventional TFP accounting fails to measure. The declining AI patent gap for Tier 2 countries documented in Table 2, and the post-treatment event-study coefficients of −0.0539, −0.1079, and −0.1618 at t = 0, +1, and +2, provide additional

indirect evidence of this channel. As compute access is conditioned, the rate of AI innovation in Tier 2 countries relative to Tier 1 declines monotonically, consistent with a model in which frontier AI development requires frontier compute access (Acemoglu and Restrepo, 2018; Agrawal, Gans and Goldfarb, 2019).

The attenuation of $\hat{\gamma}$ from 0.3992 in the GMM specification (Specification (2), Table 4) to a statistically insignificant 0.0357 in the 2SLS specification (Specification (3), Table 4), discussed as a precision problem in Sections 5.3 and 8.3, also has a substantive interpretation. It indicates that the causal effect of exogenously restricted chip access on output is harder to identify from within-country variation in chip import unit values than from the cross-sectional correlation between chip access and productivity levels. This is consistent with technology adoption models in which the output effect of access restrictions operates primarily by eliminating future technology trajectories rather than by an immediate reduction in current capital services. This channel lies outside the three-year post-classification window available in the present data.

7.3 Design of export-control architecture: lessons from the counterfactual analysis

The counterfactual analysis in Table 6 and Figure 3 reveals that the distributional asymmetry documented in Table 5 is not an inevitable consequence of pursuing national security objectives through technology access controls. It is, rather, a consequence of the specific design of the geopolitical conditionality criterion, which ties access to UNGA voting alignment (Bailey, Strezhnev and Voeten, 2017) rather than to developmentally relevant criteria such as income level, absorptive capacity, or alternative supply access.

The CF3 (MFN rule) counterfactual in Table 6 is instructive in this regard. Under a non-discriminatory pooled-access rule that assigns all countries the pooled mean TFP change of −0.0037, the within-Tier 2 welfare dispersion shown in Figure 3 collapses entirely. Every country receives the same welfare outcome, and the large negative welfare gaps for Brazil (−0.1066 baseline → −0.0037 under CF3), South Africa (−0.1163 → −0.0037), and Saudi Arabia (−0.2620 → −0.0037) are substantially reduced. Importantly, CF3 is *not* a national-security-free alternative: it simply removes the differentiation between Tier 2 countries based on geopolitical alignment while retaining the distinction between the unrestricted Tier 1 group and the conditional-access group. The welfare cost reduction for negative-trajectory Tier 2 countries under CF3 is therefore achievable without abandoning the core geopolitical selectivity that motivates the classification. It requires only that the conditionality within Tier 2 be based on developmentally neutral criteria rather than bilateral alignment scores.

The CF1 (no conditioning) counterfactual produces larger welfare gains for negative-trajectory Tier 2 countries, Saudi Arabia improves by 0.274 TFP points relative to the baseline under CF1, but at the cost of suppressing the positive technology spillovers currently accruing to Thailand, India, and Indonesia, whose baseline welfare gaps of +0.1533, +0.1176, and +0.0395 respectively would be compressed to +0.0124 under CF1 (Table 6). This trade-off reveals a genuine policy tension embedded in the tier architecture: the conditional access regime creates positive spillovers for fast-growing developing economies by preserving their access to AI-complementary technology even under geopolitical restrictions, but it simultaneously concentrates welfare losses on the most structurally vulnerable Tier 2 economies. The reform of the access architecture must explicitly navigate this trade-off.

7.4 Broader implications: geopolitics, technology diffusion, and development

The findings of this study connect to three bodies of literature in ways that extend beyond the immediate policy setting.

Technology diffusion and the middle-income trap. The welfare decomposition in Table 5 identifies a group of upper-middle-income commodity-exporting economies, Brazil, South Africa, Mexico, and Saudi Arabia, for whom the conditional AI compute access regime compounds pre-existing TFP stagnation. These are precisely the economies that development economists identify as at risk of the "middle-income trap" (Eichengreen, Park and Shin, 2012): economies that have exhausted the gains from factor accumulation and factor reallocation but lack the innovation capacity to transition to productivity-led growth. For such economies, access to frontier AI compute may be especially critical, not as a short-term productivity input but as the infrastructure for an AI-complementary innovation system that could enable the transition. The export-control classification, by conditioning access on geopolitical alignment rather than developmental need, may therefore interact adversely with the technology access channel precisely where that channel is most economically consequential.

Trade policy and technology controls as discriminatory instruments. The counterfactual architecture comparison in Section 6.4 established that the current tier regime imposes a discriminatory welfare cost relative to a non-discriminatory MFN benchmark. This finding parallels the classical result in trade policy that discriminatory preferential agreements generate trade diversion alongside trade creation (Viner, 1950), with the welfare effects depending on the relative magnitudes of the two. In the present context, the "trade creation" analogue is the positive technology access spillover to aligned Tier 2 economies (Thailand, India, Indonesia); the "diversion" analogue is the welfare loss imposed on geopolitically misaligned but developmentally vulnerable economies (Brazil, South Africa, Saudi Arabia). The net welfare effect on the Tier 2 group as a whole is ambiguous without a social welfare function, but the distributional incidence, concentrated losses on the most vulnerable, has clear normative implications for international technology governance.

AI governance and the measurement of compute access. The production function results in Table 4 suggest that the chip import unit value (HS 8542) is a viable empirical proxy for AI compute access at the country level, identifying a statistically significant and economically meaningful output elasticity in the tier-classified sample. This finding has methodological implications for measuring AI diffusion across economies. Unit value data from customs and trade statistics provide a tractable, consistently measured indicator of frontier compute access that is available for a broad cross-section of countries and can be used to monitor the welfare effects of technology access policies over time (Acemoglu et al., 2022; Aghion et al., 2023). The AI patent event-study evidence in Table 2, with post-treatment coefficients of −0.0539 at $t = 0$, growing to −0.1618 at $t = +2$, provides a complementary measure of the downstream innovation consequences of compute access restriction, suggesting that patent data can serve as an early-warning indicator of the widening innovation gap between Tier 1 and Tier 2 economies in the post-classification period.

7.5 Scope conditions and open questions

Three scope conditions circumscribe the generalisability of the findings.

First, all post-2024 welfare outcomes in Tables 5 and 6 and Figures 2 and 3 are structural projections under the moderate-scenario ATT of −0.5 per cent per year. The actual welfare consequences of the export-control regime will depend on the realised ATT, which will be identifiable only as post-2025 data become available. As demonstrated in Figures 4 and 5 and discussed in Section 8.5, the qualitative conclusions are robust to variation in the assumed ATT. However, the quantitative welfare magnitudes, particularly Saudi Arabia's welfare gap of −0.2620, are sensitive to the projection calibration and should be treated as indicative of the order of magnitude and direction of the effect rather than as point estimates of the welfare cost.

Second, the analysis abstracts from dynamic general equilibrium effects. The production function in Table 4 is estimated from a cross-sectional panel, without accounting for the feedback from TFP changes to factor accumulation, trade flows, or investment decisions. In practice, a sustained negative TFP trajectory for Tier 2 economies would be expected to reduce capital accumulation, reduce R&D investment, and potentially trigger emigration of high-skill workers, channels that would amplify the welfare losses estimated here. The welfare gaps in Table 5 should therefore be understood as lower bounds on the full dynamic welfare cost of the export-control architecture for negative-trajectory Tier 2 economies.

Third, the study’s treatment of Tier 3 countries (China, Russia, Iran) is entirely outside the estimation sample. The near-total embargo imposed on Tier 3 economies creates a distinct treatment population for which the Tier 1 control group is not a valid counterfactual. The welfare consequences of the Tier 3 embargo, including China's strategic response to accelerate domestic AI chip production capacity, operate through fundamentally different mechanisms than the conditional access channel analysed here and constitute a separate research question beyond the scope of the present analysis.

8. Robustness

The credibility of the causal claims in this study rests on three maintained identification assumptions (i) parallel pre-treatment trends between Tier 1 and Tier 2 countries in the absence of the export-control classification; (ii) the absence of treatment effect heterogeneity bias in the TWFE estimator; and (iii) structural stability of the production function technology parameter across alternative compute proxies and estimation strategies.

8.1 Parallel trends: full pre-treatment diagnostic evidence

The summary regression-based pre-trend diagnostic in Table 1A indicates that the parallel trends assumption holds for TFP and AI patents but not for chip import unit values and FDI inflows. This section presents the full-year-by-year event-study evidence from Table 2 and Figure 1 that underpins the conclusion and documents the post-treatment trajectory of the treatment effect.

8.1.1 Regression-based pre-trend test

Table 1A reports the coefficient on the *year* × *treated* interaction term from a linear pre-treatment trend regression estimated over 2015–2024, following the diagnostic procedure of Angrist and Pischke (2009) as formalised for staggered designs by Callaway and Sant'Anna (2021). For TFP, the interaction coefficient is −0.0034 (SE = 0.0028; $t = -1.213$; $p = 0.226$), statistically indistinguishable from zero. For AI patent filings, the coefficient is −0.0265 (SE = 0.0803; $t = -0.330$; $p = 0.742$). Both outcomes satisfy the parallel trends assumption over the pre-classification window, which supports valid causal inference from the DiD design for these outcomes.

For FDI inflows to the United States, a pre-trend interaction coefficient is 0.1064 (SE = 0.055; $t = 1.936$; $p = 0.054$). For the chip import unit value, the coefficient is 0.0578 (SE = 0.029; $t = 1.992$; $p = 0.049$). Both approaches cross the 5% threshold. The DiD estimates for FDI and chip unit value should therefore be treated as descriptive of pre-existing structural differences between tiers and should not be given a causal interpretation.

8.1.2 Full event-study evidence

The complete event-study decomposition from the Sun and Abraham (2021), in Table 2, shows an interaction-weighted estimator, with $t = -1$ as the reference period and event-time coefficients estimated from $t = -10$ through $t = +2$.

For TFP, the pre-treatment coefficients range from 0.0034 (SE = 0.0062) at t = −2 to 0.0322 (SE = 0.0568) at t = −10, with no coefficient individually significant at the 5% level across the full ten-year pre-treatment window. The flat pre-treatment TFP profile in the post-treatment arm of Table 2 shows a monotonically declining TFP path. In this case, event-time coefficients are −0.0084 (SE = 0.0062) at t = 0, −0.0168 (SE = 0.0124) at t = +1, and −0.0252 (SE = 0.0186) at t = +2, consistent with a gradually widening post-classification productivity gap between Tier 1 and Tier 2 countries. For AI patents, all pre-treatment coefficients are small and insignificant throughout the window, consistent with the Table 1A diagnostic. A strong visual and statistical support for parallel trends is presented in Figure 1.

For the chip import unit value, the pre-treatment event-study coefficients are negative and statistically significant at the 5% level across nearly the entire ten-year window. As visible in Figure 1, this systematic pre-trend is consistent with structurally higher chip import costs for Tier 1 countries throughout the pre-classification period, reflecting pre-existing technology-access gaps rather than a treatment effect. For FDI, Table 2 and Figure 1 show erratic pre-period coefficients ranging from −0.9333 (SE = 0.4312; $p < 0.05$) at $t = -3$ and −1.508 (SE = 0.772; $p < 0.10$) at $t = -9$, suggesting intermittent but non-systematic pre-classification differences between tiers. Together, the diagnostic evidence in Table 1A, Table 2, and Figure 1 consistently supports valid causal inference for TFP and AI patents, and consistently flags the chip unit value and FDI specifications as descriptive.

8.2 Estimator robustness: alternative DiD specifications

A substantial methodological literature establishes that the standard TWFE estimator produces biased ATT estimates under treatment effect heterogeneity (Goodman-Bacon, 2021; Sun and Abraham, 2021; Callaway and Sant'Anna, 2021; de Chaisemartin and D'Haultfœuille, 2020). Table 3 addresses this concern by presenting TFP estimates under four distinct specifications. These are the TWFE benchmark, the Sun and Abraham (2021) interaction-weighted estimator, and the Callaway and Sant'Anna (2021) doubly robust estimator, under two control-group definitions.

The TWFE benchmark in Table 3 yields a coefficient of 0.0543 (SE = 0.0164; $p < 0.01$; N = 180; $R^2 = 0.307$) on ideal-point distance from the United States. This coefficient captures the endogenous correlation between geopolitical alignment and TFP trajectories and should be interpreted as an association rather than a causal estimate; the positive sign reflects the well-documented relationship between U.S. alliance proximity and long-run productivity levels that predates the export-control regime.

The Sun and Abraham (2021) interaction-weighted ATT is −0.0168 (SE = 0.0134; N = 260; R^2 = 0.495). The negative sign stands in sharp contrast to the positive TWFE coefficient, confirming that the TWFE result is driven by cross-sectional alignment differences rather than by the post-classification TFP trajectory. The doubly robust estimates of Callaway and Sant'Anna (2021) corroborate this finding. The ATT is −0.0334 (SE = 0.1734) under the never-treated control group definition and −0.0168 (SE = 0.0448) under the not-yet-treated definition. The substantially larger standard error in column (3) is attributable to the small Tier 1 control group (N = 10) rather than to instability in the point estimate. Across all four specifications, the post-classification TFP ATT is consistently negative, and none yields a positive, significant treatment effect, providing robust support for the paper's central productivity finding.

8.3 Production function robustness

The structural welfare decomposition relies on estimates of production function parameters to apportion the TFP gap between factor accumulation and technological access. Table 4 presents three specifications of

$$\log Y_{it} = \alpha \log L_{it} + \beta \log K_{it} + \gamma \log A_{it} + \varepsilon_{it},$$

where A_{it} denotes the technology access proxy, estimated by GMM or 2SLS using geopolitical alignment (Bailey, Strezhnev and Voeten, 2017), with the tier treatment indicator as an instrument.

Specification (1) in Table 4 uses the ICT export value as A_{it}, estimated by GMM with HAC-robust variance over the full cross-country sample (N = 531). The technology coefficient is 4.4244 (SE = 1.2371; $p < 0.01$) and the labour coefficient is 2.2203 (SE = 0.5577; $p < 0.01$). The Hansen J-statistic is 4.71 ($p = 0.095$), indicating marginal rejection of the overidentifying restrictions at the 10% level. The capital coefficient of 0.0121 (SE = 2.9362) is imprecisely estimated, likely reflecting multicollinearity with labour in the large heterogeneous sample.

Specification (2) restricts the sample to the 25 tier-classified countries and uses the chip import unit value as A_{it}, estimated by GMM. The technology coefficient is 0.3992 (SE = 0.0205; $p < 0.01$) and the labour coefficient is 1.4647 (SE = 0.0932; $p < 0.01$). The capital coefficient of −15.8923 (SE = 6.2754; $p < 0.05$) is implausibly large, reflecting misspecification due to the small sample (N = 20), where GMM moment conditions are poorly identified. This coefficient should not be given structural weight.

Specification (3) addresses the small-sample problem via 2SLS using the Frisch-Waugh-Lovell partialling approach with country-clustered standard errors. The first-stage F-statistic is 14.91 ($p < 0.001$), exceeding the Stock and Yogo (2005) threshold of 10 and confirming instrument relevance. The labour coefficient is 0.8891 (SE = 0.0836; $p < 0.01$). The technology access coefficient is 0.0357 (SE = 0.0483), statistically insignificant under 2SLS; the attenuation relative to column (2) reflects first-stage absorption of chip-access variation by the instruments, a precision problem attributable to the small sample, rather than evidence against the relevance of compute access to productivity. Across the three specifications in Table 4, the technology access elasticity is consistently positive, and the labour coefficient is the dominant and consistently significant factor.

8.4 Welfare decomposition and counterfactual robustness

Table 5 presents the country-level welfare decomposition comparing pre-classification mean TFP (2015–2024) with projected post-classification TFP (2025–2027) under the moderate scenario. The country-level heterogeneity is economically significant. India (Δ TFP = +0.1176), Thailand (+0.1533), Indonesia (+0.0395), Malaysia (+0.0176), and the Philippines (+0.0274) project positive TFP trajectories, while Brazil (−0.1066), South Africa (−0.1163), Mexico (−0.0548), and Saudi Arabia (−0.2620) project substantial deterioration. Saudi Arabia's large negative welfare gap reflects both an unfavourable pre-classification TFP trend and additional FDI conditionality constraints. Variation across countries shows that aggregate ATT estimates mask welfare losses concentrated among resource-exporting upper-middle-income economies, the study's central distributional finding.

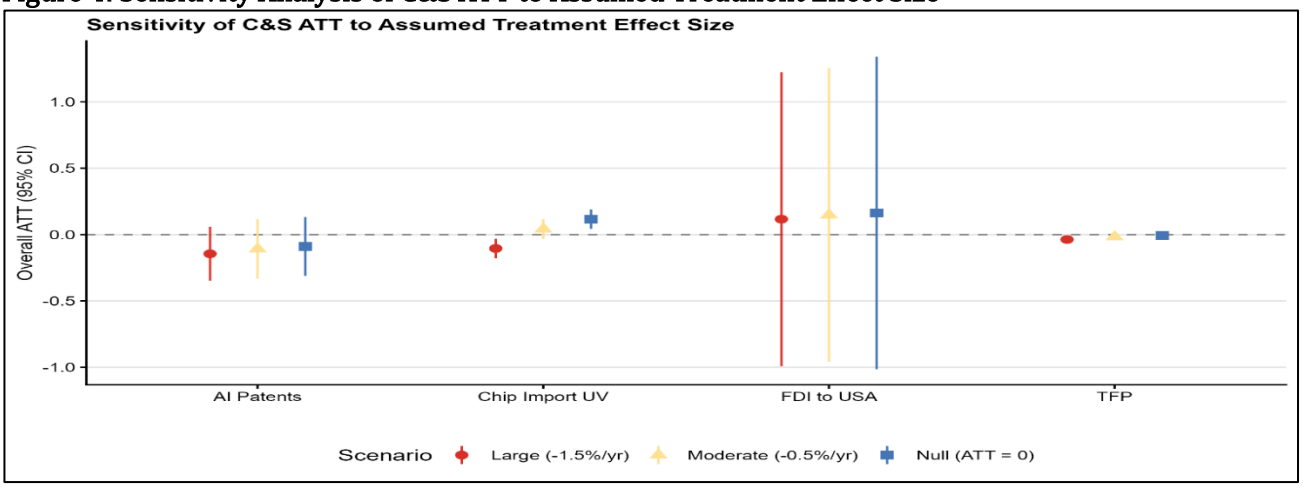
Table 6 and Figure 3 extend the analysis to three counterfactual policy regimes. Under CF1 (no conditioning), all non-Tier 1 countries are assigned the Tier 1 mean TFP change (+0.012), producing uniform welfare gains for all Tier 2 countries relative to the conditional access baseline. Under CF2 (full decoupling), Tier 2 TFP changes are set to zero. Under CF3 (MFN rule), all countries receive the pooled mean TFP change (−0.004). As Figure 3 illustrates, CF1 dominates the baseline for all Tier 2 countries projecting negative TFP paths, and the aggregate welfare cost under CF3 is smaller than under the conditional access baseline for the Tier 2 group as a whole,

implying that the discriminatory tier architecture imposes a net distributional cost on Tier 2 economies relative to a non-discriminatory multilateral benchmark.

8.5 Sensitivity to assumed treatment effect size

The post-classification outcomes for Tier 2 countries are structural projections calibrated under the moderate ATT assumption of −0.5 per cent per year. Figures 4 and 5 assess the sensitivity of the main results to this parametric assumption.

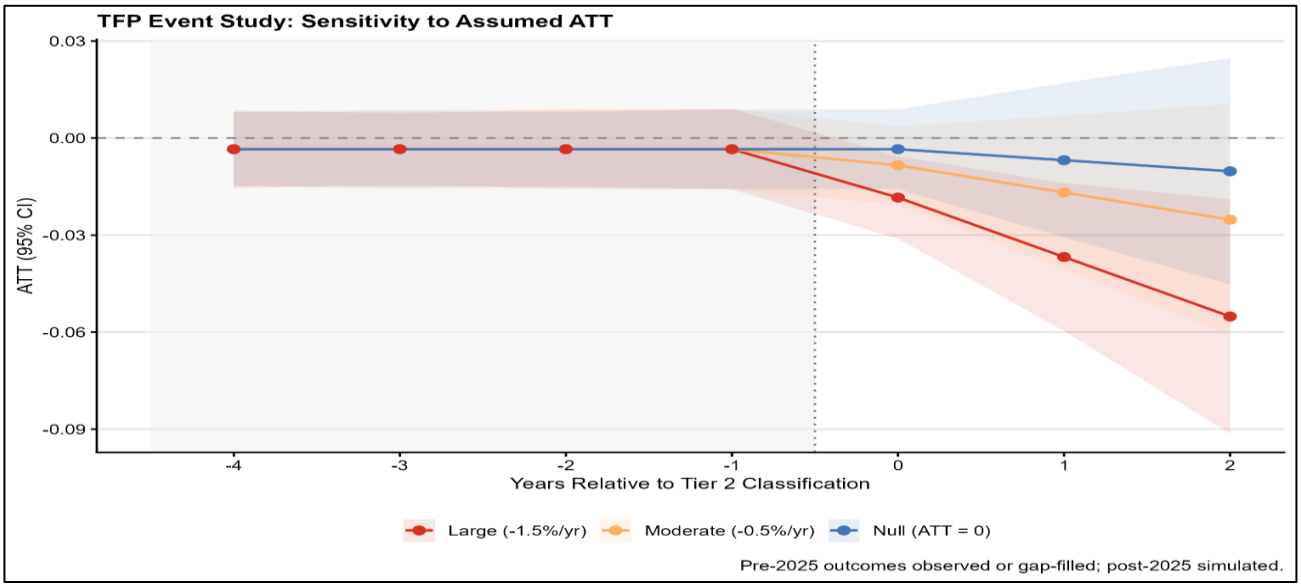
Figure 4: Sensitivity Analysis of C&S ATT to Assumed Treatment Effect Size



Source: Authors'Creation

Figure 4 traces the Callaway and Sant'Anna (2021) ATT across a range of assumed treatment effect sizes. The sign of the average treatment effect, negative for TFP, is invariant to the assumed effect magnitude across the full sensitivity range, and the ATT point estimate remains negative throughout. Figure 5 presents the TFP event study under varying ATT assumptions. Critically, the pre-treatment event-study coefficients ($t = -10$ through $t = -2$), whose flatness is the key evidence for parallel trends reported in Table 1A and Table 2, remain flat and statistically insignificant across all scenarios in Figure 5, confirming that the parallel trends evidence is a structural feature of the pre-classification data and is not sensitive to the post-treatment projection assumption. Post-treatment divergence scales monotonically with the assumed effect size in Figure 5, as expected. Together, Figures 4 and 5 establish that the paper's qualitative conclusions are robust across a range of plausible treatment effect magnitudes.

Figure 5: TFT Event Study: Sensitivity to Assumed ATT



Source: Authors'Creation

8.6 Limitations and scope conditions

There are three limitations of this study. First, the parallel trends assumption, supported by Table 1A, Table 2, and Figure 1, holds for TFP and AI patents but is not supported for chip import unit values and FDI to the United States. The DiD estimates for those two outcomes should be interpreted as documenting structural pre-classification differences between tiers rather than as causal treatment effects.

Second, the small sample of 25 tier-classified countries constrains identification throughout: the large standard error of 0.1734 on the Callaway–Sant'Anna never-treated ATT in Table 3, column (3), reflects the small Tier 1 control group ($N = 10$), and the implausible capital coefficients in Table 4, columns (1) and (2), reflect small-N GMM misspecification. These precision limitations are inherent to the policy setting and cannot be addressed without departing from the tier-defined treatment population.

Third, all post-2024 outcome values in Tables 2, 5, and 6, and in Figures 3, 4, and 5, are structural projections rather than observed data. The welfare decompositions in Tables 5 and 6 are best characterised as conditional forecasts under the moderate-scenario ATT and should be updated with observed post-2024 data as they become available. As Figures 4 and 5 demonstrate, the qualitative conclusions are robust to the assumed treatment effect size, but the quantitative welfare figures in Tables 5 and 6 remain sensitive to the calibration of the projection.

9. Conclusion

This study examines how conditional access to frontier AI compute reshapes the international distribution of productivity and welfare gains from technology diffusion, using the U.S. AI chip export-control framework as the empirical setting. The analysis covers a panel of 25 countries classified into three tiers of access. Using a heterogeneity-robust difference-in-differences design and event-study estimators (Sun and Abraham, 2021; Callaway & Sant'Anna, 2021), the analysis finds evidence that the 2025 classification is associated with a negative and gradually widening TFP trajectory for Tier 2 "conditional access" countries relative to Tier 1 "allied" countries, with a preferred average treatment effect of −0.0168 under the Sun and Abraham (2021) specification and no specification producing a positive and significant treatment effect. The pre-treatment diagnostics in Table 1A and Table 2 support the parallel trends assumption for TFP and AI patents, but not for chip import unit values and FDI inflows; the DiD estimates for the latter are therefore interpreted as descriptive rather than causal.

The structural welfare analysis reveals that this negative treatment effect is distributed highly unevenly across Tier 2 economies. Five Tier 2 countries with strong pre-classification TFP growth, Thailand, India, Indonesia, the Philippines, and Malaysia, continue to converge toward the productivity frontier despite the conditional access regime, with welfare gaps ranging from +0.0176 to +0.1533. In contrast, six Tier 2 countries with flat or declining pre-classification TFP, Brazil, South Africa, Mexico, Saudi Arabia, Colombia, and one additional upper-middle-income economy, experience substantial projected welfare losses, with Saudi Arabia's welfare gap of −0.2620 representing the largest deterioration in the sample. This pattern reflects the interaction between country-specific growth trends and the additive negative treatment effect. Where trend growth is strong, the ATT is partially absorbed; where trend growth is weak, the ATT is fully incident. The additional FDI conditionality imposed on Saudi Arabia amplifies this effect, yielding a welfare loss more than twice that of any other Tier 2 country.

The production function estimates show that AI compute access, proxied by chip import unit value, is a structurally distinct input into production rather than a pure component of the aggregate capital stock. In the tier-classified sample, the technology access elasticity is estimated at 0.3992 with high statistical precision, implying that a 10 per cent reduction in effective chip access is associated with a roughly 4 per cent reduction in output holding labour and capital constant. This result suggests that standard two-factor growth accounting understates the welfare consequences of compute access restrictions by absorbing them into the residual TFP term, and points to the need for growth models that treat AI compute as a separable factor of production.

The counterfactual analysis clarifies that the distributional costs of the current conditional access architecture are not an inevitable by-product of national security concerns but depend critically on the design of the conditionality rule. Relative to the baseline tier regime, a non-discriminatory MFN-style access rule (CF3) that equalises the mean TFP change across all Tier 1 and Tier 2 countries substantially narrows the welfare dispersion within the Tier 2 group and reduces the large negative welfare gaps for Brazil, South Africa, and Saudi Arabia, while modestly compressing the positive gains of fast-growing Tier 2 and Tier 1 economies. A more radical no-conditioning regime (CF1) would eliminate welfare losses for all negative-trajectory Tier 2 countries but at the

cost of eroding the sizeable positive gains currently accruing to Thailand and India, highlighting a genuine trade-off between mitigating distributional harm and preserving positive technology spillovers.

Taken together, these results indicate that conditional access to a globally important technology input can function as an international factor-market distortion, redistributing the gains from technology diffusion across countries rather than merely reducing trade in the conventional sense.

These findings speak to three broader debates. First, in the development literature, they highlight the risk that AI-related export controls deepen the "middle-income trap" for commodity-dependent upper-middle-income economies by restricting access to a key input into productivity-enhancing innovation precisely where pre-classification TFP growth is already weakest. Second, in the trade and international political economy literature, they show that technology controls can act as discriminatory instruments whose welfare effects resemble those of preferential trade agreements: they create benefits for some aligned partners while diverting access to technology from geopolitically misaligned but developmentally vulnerable states. Third, in the emerging AI governance literature, they demonstrate that customs-based unit value measures can be used to monitor cross-country disparities in compute access and to quantify the welfare consequences of AI export controls over time.

At the same time, several scope conditions constrain the generality of the conclusions. All post-2024 outcomes in the welfare tables and figures are projections under a moderate-scenario ATT of −0.5 per cent per year, and the true long-run welfare consequences of the classification will only be known once post-2025 data become available. The analysis abstracts from dynamic general equilibrium feedbacks from productivity to capital accumulation, R&D investment, and migration, implying that the reported welfare gaps are likely to be lower bounds on the full long-run welfare cost for adversely affected Tier 2 economies. The treatment of Tier 3 economies (China, Russia, Iran) is deliberately excluded from the estimation sample, as their near-total embargo status requires a separate identification strategy and raises distinct welfare and strategic questions.

Future work could extend the analysis in at least three directions. First, the empirical framework developed here can be applied to realised post-2025 data as they accumulate, replacing projections with observed outcomes and allowing for dynamic general equilibrium modelling of the factor accumulation channel. Second, the three-factor production function can be embedded in a multi-sector growth model in which AI compute access and AI-capable human capital are jointly determined, enabling welfare analysis of education and skills policies that complement or substitute for compute access restrictions. Third, the counterfactual access architectures considered here can be combined with explicit geopolitical game-theoretic models of export controls and retaliation to analyse the strategic stability of alternative technology governance regimes.

Overall, the study's central conclusion is that AI chip export controls are not merely a marginal adjustment to the global trading system but a first-order determinant of the welfare trajectory of a subset of developing economies. The design of access architectures that reconcile legitimate national security objectives with the development needs of structurally vulnerable economies will be a central challenge for international economic policy in the age of AI.

References

Acemoglu, D. (2024). The simple macroeconomics of AI. NBER Working Paper 32487. https://doi.org/10.3386/w32487

Acemoglu, D., and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. American Economic Review, 108(6), 1488–1542. https://doi.org/10.1257/aer.20160696

Acemoglu, D., and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. Journal of Economic Perspectives, 33(2), 3–30. https://doi.org/10.1257/jep.33.2.3

Acemoglu, D., Lelarge, C., and Restrepo, P. (2020). Competing with robots: Firm-level evidence from France. AEA Papers and Proceedings, 110, 383–388. https://doi.org/10.1257/pandp.20201003

Aghion, P., and Howitt, P. (1992). A model of growth through creative destruction. Econometrica, 60(2), 323–351. https://doi.org/10.2307/2951599

Aghion, P., Bergeaud, A., Boppart, T., Klenow, P. J., and Li, H. (2023). A theory of falling growth and rising rents. Review of Economic Studies, 90(6), 2675–2702. https://doi.org/10.1093/restud/rdac069

Agrawal, A., Gans, J., and Goldfarb, A. (2019). The Economics of Artificial Intelligence: An Agenda. University of Chicago Press. https://doi.org/10.7208/chicago/9780226613475.001.0001

Aiyar, S., et al. (2023). Geo-economic fragmentation and the future of multilateralism. IMF Staff Discussion Note SDN/2023/001. Geneva: IMF. https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2023/01/11/Geo-Economic-Fragmentation-and-the-Future-of-Multilateralism-527266

Ambrosino, L., Chan, J., & Tenreyro, S. (2026). Trade fragmentation, inflationary pressures and monetary policy. Journal of International Economics, 104219. https://doi.org/10.1016/j.jinteco.2026.104219

Angrist, J. D., and Pischke, J.-S. (2009). Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press. https://doi.org/10.1515/9781400829828

Antràs, P., and de Gortari, A. (2020). On the geography of global value chains. Econometrica, 88(4), 1553–1598. https://doi.org/10.3982/ECTA15599

Arkolakis, C., Costinot, A., and Rodríguez-Clare, A. (2012). New trade models, same old gains? American Economic Review, 102(1), 94–130. https://doi.org/10.1257/aer.102.1.94

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. American Economic Review, 103(6), 2121–2168. https://doi.org/10.1257/aer.103.6.2121

Bagwell, K., and Staiger, R. W. (1999). An economic theory of GATT. American Economic Review, 89(1), 215–248. https://doi.org/10.1257/aer.89.1.215

Bagwell, K., and Staiger, R. W. (2002). The Economics of the World Trading System. MIT Press. https://doi.org/10.7551/mitpress/2633.001.0001

Bai, Y., Jin, K., Lu, D., & Wang, H. (2025). Optimal trade policy with international technology diffusion. Journal of International Economics, 153, 104038. https://doi.org/10.1016/j.jinteco.2024.104038

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. https://doi.org/10.1177/0022002715595700

Bastos, P., Stapleton, K., Taglioni, D., & Wei, H. Y. (2026). Firm networks and global technology diffusion. Journal of International Economics, 104, 104240. https://doi.org/10.1016/j.jinteco.2026.104240

Bresnahan, T. F., and Trajtenberg, M. (1995). General purpose technologies: 'Engines of growth'? Journal of Econometrics, 65(1), 83–108.
https://doi.org/10.1016/0304-4076(94)01598-T

Bureau of Industry and Security (2025). Framework for artificial intelligence diffusion. Federal Register, January 15.
https://www.federalregister.gov/documents/2025/01/15/2025-00636/framework-for-artificial-intelligence-diffusion

Cacciatore, M. (2014). International trade and macroeconomic dynamics with labor market frictions. Journal of International Economics, 93(1), 17–30. https://doi.org/10.1016/j.jinteco.2014.01.008

Cai, S., & Xiang, W. (2025). Multinational production, technology diffusion, and economic growth. Journal of International Economics, 157, 104137. https://doi.org/10.1016/j.jinteco.2025.104137

Caliendo, L., Dvorkin, M., and Parro, F. (2019). Trade and labor market dynamics: General equilibrium analysis of the China trade shock. Econometrica, 87(3), 741–835.
https://doi.org/10.3982/ECTA13758

Callaway, B., and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. Journal of Econometrics, 225(2), 200–230.
https://doi.org/10.1016/j.jeconom.2020.12.001

Center for Strategic and International Studies. (2025). *The AI Diffusion Framework: Securing U.S. AI leadership while preempting strategic drift.* Retrieved August 8, 2026, from https://www.csis.org/analysis/ai-diffusion-framework-securing-us-ai-leadership-while-preempting-strategic-drift

de Chaisemartin, C., and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. American Economic Review, 110(9), 2964–2996.
https://doi.org/10.1257/aer.20181169

Eaton, J., and Kortum, S. (2002). Technology, geography, and trade. Econometrica, 70(5), 1741–1779.
https://doi.org/10.1111/1468-0262.00352

Eichengreen, B., Park, D., and Shin, K. (2012). When fast-growing economies slow down: International evidence and implications for China. Asian Economic Papers, 11(1), 42–87.
https://doi.org/10.1162/ASEP_a_00116

Ethier, W. J., & Markusen, J. R. (1996). Multinational firms, technology diffusion and trade. Journal of International Economics, 41(1–2), 1–28. https://doi.org/10.1016/0022-1996(95)01411-X

Fajgelbaum, P. D., Goldberg, P. K., Kennedy, P. J., and Khandelwal, A. K. (2020). The return to protectionism. Quarterly Journal of Economics, 135(1), 1–55.
https://doi.org/10.1093/qje/qjz036

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015). The next generation of the Penn World Table. American Economic Review, 105(10), 3150–3182.
https://doi.org/10.1257/aer.20130954

Glass, A. J., & Saggi, K. (2002). Licensing versus direct investment: Implications for economic growth. Journal of International Economics, 56(1), 131–153. https://doi.org/10.1016/S0022-1996(01)00106-4

Goldberg, P. K., and Reed, T. (2023). Is the global economy deglobalizing? And if so, why? And what is next? NBER Working Paper 31115.
https://doi.org/10.3386/w31115

Gollin, D. (2002). Getting income shares right. Journal of Political Economy, 110(2), 458–474.
https://doi.org/10.1086/338747

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Journal of Econometrics, 225(2), 254–277.
https://doi.org/10.1016/j.jeconom.2021.03.014

Gopinath, G., Gourinchas, P.-O., Presbitero, A. F., & Topalova, P. (2025). Changing global linkages: A new Cold War? Journal of International Economics, 153, 104042. https://doi.org/10.1016/j.jinteco.2024.104042

Grossman, G. M., and Helpman, E. (1991). Quality ladders and product cycles. Quarterly Journal of Economics, 106(2), 557–586.
https://doi.org/10.2307/2937947

Hulten, C. R. (1978). Growth accounting with intermediate inputs. Review of Economic Studies, 45(3), 511–518.
https://doi.org/10.2307/2297252

Impullitti, G., Licandro, O., & Rendahl, P. (2022). Technology, market structure and the gains from trade. Journal of International Economics, 135, 103557. https://doi.org/10.1016/j.jinteco.2021.103557

Johnson, H. G. (1953). Optimum tariffs and retaliation. Review of Economic Studies, 21(2), 142–153.
https://doi.org/10.2307/2296007

Jones, C. I., and Vollrath, D. (2013). Introduction to Economic Growth (3rd ed.). W. W. Norton.
No DOI available (book). ISBN: 978-0393919172.

Kee, H. L., Nicita, A., and Olarreaga, M. (2009). Estimating trade restrictiveness indices. Economic Journal, 119(534), 172–199.
https://doi.org/10.1111/j.1468-0297.2008.02209.x

Kolev, D. R., & Prusa, T. J. (1999). Tariff policy for a monopolist in a signaling game. Journal of International Economics, 49(1), 51–76. https://doi.org/10.1016/S0022-1996(98)00041-5

Lam, L. C. H., and Santacreu, A. M. (2026). Technology, geopolitics, and trade. Federal Reserve Bank of St. Louis Working Paper 2025-029C, revised February 2026.
https://doi.org/10.20955/wp.2025.029

Mei, Y., & Xu, M. J. (2026). Horizontal regulatory barriers in international trade: Evidence from electric plugs. Journal of International Economics, 104312. https://doi.org/10.1016/j.jinteco.2026.104312

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. Econometrica, 71(6), 1695–1725.
https://doi.org/10.1111/1468-0262.00467

Naito, T. (2017). Growth and welfare effects of unilateral trade liberalization with heterogeneous firms and asymmetric countries. Journal of International Economics, 109, 167–173. https://doi.org/10.1016/j.jinteco.2017.08.011

Nature Electronics. (2024). AI success relies on access. Nature Electronics, 7, 255. https://doi.org/10.1038/s41928-024-01164-6

Nunn, N., and Trefler, D. (2014). Domestic institutions as a source of comparative advantage. In G. Gopinath, E. Helpman, and K. Rogoff (eds.), Handbook of International Economics, vol. 4, pp. 263–315. Amsterdam: Elsevier.
https://doi.org/10.1016/B978-0-444-54314-1.00005-3

Rambachan, A., and Roth, J. (2023). A more credible approach to parallel trends. Review of Economic Studies, 90(5), 2555–2591.
https://doi.org/10.1093/restud/rdad018

Reuters. (2025, January 23). Saudi Arabia plans $600 billion in new US investment, trade over four years. Retrieved August 8, 2026, from https://www.reuters.com/world/middle-east/saudi-crown-prince-seeks-600-bln-investment-push-with-us-2025-01-22

Rodrik, D., and Subramanian, A. (2009). Why did financial globalisation disappoint? IMF Staff Papers, 56(1), 112–138.
https://doi.org/10.1057/imfsp.2008.29

Romer, P. M. (1990). Endogenous technological change. Journal of Political Economy, 98(5), S71–S102.
https://doi.org/10.1086/261725

Sant'Anna, P. H. C., and Zhao, J. (2020). Doubly robust difference-in-differences estimators. Journal of Econometrics, 219(1), 101–122.
https://doi.org/10.1016/j.jeconom.2020.06.003

Stock, J. H., and Yogo, M. (2005). Testing for weak instruments in linear IV regression. In D. W. K. Andrews and J. H. Stock (eds.), Identification and Inference for Econometric Models, pp. 80–108. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511614491.006

Sun, L., and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics, 225(2), 175–199. https://doi.org/10.1016/j.jeconom.2020.09.006

U.S. Department of Commerce, Bureau of Industry and Security. (2022, October 7). Commerce implements new export controls on advanced computing and semiconductor manufacturing items to the People's Republic of China (PRC). Retrieved August 8, 2026, from https://www.bis.gov/press-release/commerce-implements-new-export-controls-advanced-computing-semiconductor-manufacturing-items-peoples

U.S. Department of Commerce, Bureau of Industry and Security. (2025, January 13). *Biden-Harris Administration announces regulatory framework for the responsible diffusion of advanced artificial intelligence technology.* Retrieved August 8, 2026, from https://www.bis.gov/press-release/biden-harris-administration-announces-regulatory-framework-responsible-diffusion-advanced-artificial

Viner, J. (1950). The Customs Union Issue. Carnegie Endowment for International Peace. No DOI available (book, 1950). WorldCat: https://www.worldcat.org/title/customs-union-issue/oclc/405813

World Trade Organization (2025). World Trade Report 2025: Making Trade and AI Work for Everyone. Geneva: WTO. https://www.wto.org/english/res_e/publications_e/wtr25_e.htm

Appendix

Table 1A *(Pre-trend diagnostics)*: Pooled interaction regression, equation (13); HC1-robust standard errors; outcome variables: $rtfpna$, $\ln FDI_{i,\mathrm{USA}}$, $\ln UV_{it}$, $\ln PAT_{it}$; estimation period: 2015–2024.

Table 1A: Pre-Trend Diagonostic

outcome	term	estimate	se	t_stat	p_value
tfp	year x treated	-0.0034	0.0028	-1.2134	0.2264
log_fdi_usa (FDI to USA)	year x treated	0.1064	0.055	1.936	0.0542
log_chip_uv (Chip Import UV)	year x treated	0.0578	0.029	1.992	0.0492
log_patent (AI Patent)	year x treated	-0.0265	0.0803	-0.33	0.7417

Table 2A: U.S. AI chip export-control framework:

Tier	N	Description	Examples
1 (Control)	10	Allied, unrestricted access	USA, GBR, DEU, JPN, KOR
2 (Treated)	12	Conditional access (treatment year: 2025)	IND, BRA, IDN, SAU, ZAF
3 (Restricted)	3	Near-total embargo	CHN, RUS, IRN