Tariff Cuts, Policy Uncertainty, and the Force of Many: The Power of Plurilateral Agreements *

Lasha Chochua † Irene Iodice‡

August 26, 2025

Abstract

When countries liberalize trade together, do they gain more than by acting alone? We quantify the coordination benefits of joint liberalization using the Phase II expansion of the Information Technology Agreement (ITA), a plurilateral deal covering more than 12% of global trade, including semiconductors and other advanced electronics. Exploiting the agreement's institutional design, we compare liberalized products to similar ones that were not covered, isolating three channels: applied tariff cuts, reductions in trade policy uncertainty, and coordination spillovers. Liberalized goods gained 4–6% in market access, with coordination spillovers accounting for the largest share. The results reveal a cascade effect: as more countries join the liberalization effort, gains amplify for all—showing that the benefits of acting together exceed those from cutting trade costs alone.

*We are grateful to Chad Bown, Maria Pthaskina, Gerald Willmann, Andrew Bernard, Emanuel Ornelas, James Lake, Lars Nilsson, Fedora Teti, and participants of Cesifo Munich, CITP University of Startchylde, Melbourne Empirical Trade Conference 2024, Kiel-Aarhus 2023 Workshop and ETSG 2023.

†International School of Economics at Tbilisi State University. E-mail: lasha.chochua@iset.ge.

‡Bielefeld University & CESifo Research Network member. E-mail: irene.iodice@uni-bielefeld.de.

1 Introduction

Since the establishment of the General Agreement on Tariffs and Trade (GATT), trade liberalization has unfolded through three distinct pathways: broad multilateral negotiations covering multiple sectors, preferential trade agreements creating discriminatory access, and issue-specific plurilateral agreements targeting particular industries. While the first two forms have been extensively studied, plurilateral agreements remain comparatively underexplored despite their growing importance as the multilateral trading system faces mounting challenges (Bagwell et al., 2016). This gap is consequential. At a time when global negotiations have stalled and protectionist pressures fuel unilateral actions, understanding how coordination affects the gains from liberalization has become critical for the future of the multilateral trade system.

Plurilateral agreements offer a distinctive approach to this challenge. Defined by their non-discriminatory structure, targeted sectoral focus, and—in many cases—open membership, they combine the efficiency gains of narrow scope with the coordination benefits of multilateral action. By extending market access improvements to all trading partners on a most-favored-nation basis, they avoid the trade diversion effects that limit welfare gains from preferential agreements while harnessing general equilibrium spillovers that amplify benefits beyond what individual countries could achieve through unilateral action. Their limited scope reduces cross-issue linkages, making negotiations more manageable while enabling precise evaluation of coordination effects in trade policy.

We develop a structural gravity framework that highlights three channels through which plurilateral liberalization operates: (1) direct reductions in applied MFN tariffs, (2) reductions in trade policy uncertainty through the narrowing of tariff "water"—the gap between bound and applied rates, and (3) a novel coordination mechanism that magnifies trade responses through general equilibrium feedbacks across global markets. The key insight centers on the non-discriminatory nature of MFN liberalization: when members of a plurilateral agreement reduce tariffs, these reductions automatically extend to all trading partners, creating aggregate policy changes that trigger general equilibrium adjustments across global markets. This coordination channel operates through multilateral resistance terms in the gravity equation—when multiple large economies coordinate their liberalization simultaneously, widespread improvements in exporters' global market access intensify competitive pressure in each liberalizing market, amplifying trade responses beyond what individual countries could achieve through sequential action.

The 2016 Phase II expansion of the Information Technology Agreement (ITA) provides a compelling empirical setting for quantifying these effects. Covering roughly 12 percent of world trade,[1] the agreement combined open membership with a "critical mass" rule to ensure broad participation while maintaining non-discriminatory treatment. The expansion covered a wide array of high-tech products—including next-generation semiconductors, advanced telecommunications equipment, and emerging digital technologies—providing a rich context for evaluating targeted liberalization outcomes.

Several institutional features make the ITA Phase II uniquely suited for quasi-experimental analysis. The WTO coordinated consensus among participating countries on a detailed product list, reducing the role of unilateral policy preferences. The liberalization process was implemented through a phased reduction of bound tariffs over four years, with the linear schedule creating plausibly exogenous variation for empirical analysis. Most importantly for our identification strategy, the agreement specified products through two distinct mechanisms: Attachment A listed established products with universally agreed-upon Harmonized System (HS) codes that all participants could readily implement, while Attachment B provided textual descriptions for innovative products lacking standardized international classifications.

This institutional distinction creates quasi-experimental variation in what we term the "force of many"—the number of countries liberalizing each specific product line. For Attachment A products, all 82 ITA participants implemented identical liberalization. However, for cutting-edge Attachment B products, the absence of harmonized codes forced national customs authorities to map product descriptions onto existing tariff schedules through administrative procedures. This mapping process, driven by technical classifications rather than economic considerations, generated exogenous variation: some products were liberalized by the full coalition, others by varying subsets, depending on administrative interpretations.

Our empirical strategy combines a difference-in-differences design with the structural gravity framework to exploit this variation. We compare liberalized ITA products to carefully selected non-liberalized control products within the same industry classifications, implementing several approaches: comparing expansion products to similar non-covered goods, exploiting within-product variation in liberalizing country numbers vs non-members to control for product specific trends.

We find that ITA liberalization increased market access for covered products by

[1] See ITA Symposium: 25 Years of the Information Technology Agreement.

4–6 percent, with coordination effects representing the largest component of this gain. Direct tariff reductions account for approximately 40 percent of the total effect, while uncertainty reduction through binding commitments explains another 20 percent. The remaining 40 percent—the coordination premium—strengthens monotonically with coalition size, with products liberalized by larger country groups experiencing disproportionately larger trade increases. This relationship proves particularly strong for high-tech Attachment B products, where we observe the greatest variation in participation. The results are robust to alternative specifications and unaffected by pre-trend concerns.

Taken together, our findings reveal that coordination yields substantially larger gains than unilateral liberalization, and that these gains grow with coalition size. By documenting how plurilateral agreements amplify the benefits of trade reform, we provide new evidence to inform the design of future initiatives in a global economy where broad-based multilateral progress remains elusive.

Related Literature Our paper contributes to four key areas of literature. First, we quantify the effects of non-discriminatory trade policy changes. While reductions in applied MFN tariffs have been among the most impactful trade policy changes since the mid-20th century, previous studies have often neglected to take into account multilateral resistance when estimating their trade effects (Anderson and Van Wincoop, 2003). There are, however, exceptions, such as Caliendo et al. (2015) and Heid et al. (2021), who estimate the substantial effects of MFN tariff reductions on trade. For instance, Caliendo et al. (2015) found that the MFN tariff reductions from the Uruguay Round explain 90% of subsequent trade gains. Furthermore, Larch et al. (2019) examine the impact of GATT/WTO membership on international trade using structural gravity modeling, overcoming previous methodological gaps by including intra-national trade flows. They find that GATT/WTO membership boosts international trade relative to domestic sales by 72% and increases trade between members by 171%. Additionally, membership facilitates trade with non-members, growing by 88%. Although GATT was more effective in fostering bilateral trade among members, WTO had broader impacts, particularly in promoting trade with non-members. We argue that examining domain-specific liberalization, such as that of the ITA, offers a better case study to identify the effects of tariff cuts that result from product-specific rather than complex liberalization schemes, which can bring about cross-linkage effects that are more challenging to disentangle.

We also contribute to the literature on the effects of WTO membership on trade policy certainty. It has been common in previous empirical studies of the impact of WTO membership on trade to use a simple dummy representing WTO membership, which combines the impacts of reducing/eliminating tariffs with the effects of greater certainty in trade policy. In spite of the difficulties associated with finding significant WTO effects in the early research (Rose (2004)), subsequent studies have discovered such effects for a limited number of members (Subramanian and Wei (2007); Eicher and Henn (2011)). By examining the trade impacts of WTO membership in a more structured manner, we distinguish between the effects of reducing and eliminating MFN tariffs in isolation from those associated with enhanced trade policy certainty, when MFN tariff reductions are the result of a collaborative liberalization effort.

Our paper also contributes to the expanding literature on trade policy uncertainty (TPU) and its impact on international trade. Previous research, such as Handley and Limao (2015), has demonstrated that TPU can significantly hinder market entry and investment, particularly in sectors with high sunk costs, due to the "option value of waiting"—where firms delay investments in anticipation of potential unfavorable changes in trade policy. As further discussed by Handley and Limão (2017), TPU is especially relevant in contexts where tariff bounds are still in place, creating uncertainty about possible future tariff increases. Our study extends this literature by examining how the phased reduction of bound tariffs in the ITA's Phase II expansion helps to mitigate TPU, thereby enhancing market access.[2] Specifically, we quantify the impact of reducing bound tariffs, which decreases uncertainty surrounding future trade policy changes and enables firms to expand their export markets.

Finally, our paper contributes to a limited body of literature focused on the Information Technology Agreement. Previous econometric analyses of trade impacts focused on the first phase of the ITA, with Mann and Liu (2009) being one of the earliest studies available. In a more recent paper, Gnutzmann-Mkrtchyan and Henn (2018) uncovered a non-linearity in the impact of tariff changes: not only is reducing tariffs beneficial, but removing them entirely leads to even greater gains in trade. Essentially, the commitment to durable tariff elimination, achieved through WTO bindings, contributes to more effective import and export facilitation than equivalent unilateral reforms.

We extend the analysis of Gnutzmann-Mkrtchyan and Henn (2018) in two key ways. First, we examine the recent Phase II expansion of the ITA, which, to our knowledge, has not been studied in the literature. Unlike the initial phase led by a few active

[2]For a comprehensive review of the literature on trade policy uncertainty, see Handley and Limão (2022).

WTO members, Phase II featured a longer, more inclusive negotiation process. This inclusivity may explain why Gnutzmann-Mkrtchyan and Henn (2018) found significant effects only for smaller economies, whereas our estimates are significant for the average country.

Second, we leverage the staging matrices signed by participating countries to measure the annual phasing out of bound tariffs, which generally followed a standard four-year linear reduction schedule. This provides us with exogenous variation in bound tariffs, enabling us to identify the effects of reduced trade policy uncertainty and to isolate an additional component we term the "plurilateral effect." We also exploit differences in Attachment A.

The paper is structured as follows: Section 2 explores the historical and institutional context of the ITA and its phase II expansion. Section 3 addresses data collection and presents stylized facts. In Section 4, we outline our identification strategy and its formulation. Section 5 presents our results and Section 6 concludes.

2 Institutional Background

The Information Technology Agreement (ITA) is a plurilateral trade agreement administered under the WTO framework that extends most-favored-nation (MFN) treatment to all WTO members, regardless of their direct participation. Participating countries commit to eliminating bound MFN tariff rates on specified information technology products, effectively establishing zero-tariff regimes for covered goods.

2.1 ITA Phase I: Foundation and Expansion

The original ITA, launched in 1996 with 29 founding WTO members, demonstrated remarkable success in both membership growth and trade expansion. By 2020, participation had increased to 82 countries, encompassing all major economies in the global electronics industry. The economic volume proved substantial: global exports of ITA Phase I products increased nearly fourfold since 1996, reaching $2 trillion in 2020 and representing over 10% of world merchandise exports.[3]

[3] See *ITA Symposium: 25 Years of the Information Technology Agreement.*

2.2 The Path to ITA Phase II: Protracted Negotiations

Building on Phase I's success, negotiators launched expansion talks in 1997, but these proved considerably more challenging. Negotiations commenced in 1997, immediately following Phase I implementation, but stalled by 1998 due to disagreements over product coverage. After remaining dormant for over a decade, momentum resumed in May 2012 when six ITA members—Canada, Japan, South Korea, Singapore, Chinese Taipei, and the United States—jointly submitted a "Concept Paper for Expanding the ITA" to the ITA Committee. By mid-2012, negotiators had developed a preliminary product list encompassing approximately 357 items at the 6-digit Harmonized System (HS6) level.

Despite initial progress and circulation of a revised expansion list in June 2013, negotiations were suspended in July 2013 due to persistent disagreements over product coverage, particularly regarding LCD panels and machine tools. The expanded membership base, now involving substantially more countries with diverse economic priorities, had fundamentally altered the negotiating dynamics compared to the original 29-member agreement. After 17 unsuccessful rounds over three years, WTO General Director Roberto Azevêdo's 2015 intervention proved decisive.[4] His consensus-building approach focused on products with broad multilateral support, albeit with narrower scope than initially proposed.

The final breakthrough came in December 2015, when WTO members formally signed the ITA Phase II expansion. The concluded agreement covered 191 products expressed in HS2007 classification and 10 product descriptions—fewer than initially proposed but representing a feasible consensus among participants. Notably, only 50 of the 80 ITA Phase I members chose to join the expansion, with India being the most significant economy to opt out.[5]

2.3 Implementation Framework

The agreement that eventually emerged established several institutional features crucial for our identification strategy.

Standardized staging approach and no safegards. In contrast to many trade agreements that allow broad special and differential treatment, ITA Phase II adopted a uniform staging schedule. All participants committed to phasing out tariffs on covered

[4]A cohmprensive discussion of the timeline is available at the link here.
[5]The list of members that joined the expansion is reported in Table 11

products in four equal annual cuts—beginning July 2016, followed by reductions at the start of each subsequent year—culminating in full elimination by January 1, 2019. Limited exceptions applied to approximately 5% of product–country pairs for three developing countries (e.g. Albania, Malaysia, and the Philippines), which were granted seven-year staging periods with equal annual reductions. In additions negotiators explicitly rejected proposals for product-specific exceptions or safeguard mechanisms. This "all-or-nothing" approach for covered products distinguishes the ITA from other trade agreements that typically permit various flexibility mechanisms.

Product classification and exogenous variation. The final ITA Phase II product list was organized into two attachments. Attachment A contained 191 products with predetermined HS2007 codes that all 50 signatories agreed to liberalize universally. Attachment B, by contrast, included 10 product descriptions covering technologically advanced goods—such as multi-component integrated circuits (MCOs), LED backlight modules, and touch-sensitive input devices—that lacked specific HS2007 codes. This gap reflects the mismatch between five-year HS revision cycles and longer negotiation timelines. As a result, these innovative products entered world markets before receiving dedicated HS subheadings and were therefore absent from the negotiation, which had started before 2012. As a result, each signatory had to map the descriptions to existing national tariff lines, typically at the 8–12-digit level, based on its own customs interpretation.

This mapping exercise created substantial cross-country heterogeneity in HS6 coverage. For example, MCOs (Item 192) were distributed across more than 60 different HS6 codes, with some liberalized universally and others by only a few countries. Similarly, self-adhesive circular polishing pads for semiconductor wafer manufacturing (Item 197) were classified differently across participants: the United States mapped them to HS 8486.90 (Parts and accessories of machines used in semiconductor manufacturing), focusing on their principal use; the European Union and China mapped them to HS 5911.90 (Textile products and articles for technical uses), emphasizing their textile composition; and Canada mapped them to both HS 8486.90 and HS 5911.90 by linking the description to two distinct national tariff lines, thereby liberalizing both.

2.4 Implications for Empirical Identification

These design features provide several advantages for causal identification.

Timing uncertainty: The extended and interrupted negotiation timeline (2012-

2015) made it difficult for businesses or governments to anticipate the final implementation date, reducing concerns about strategic behavior in anticipation of the agreement.

Product list exogeneity: The Director-General's focus on consensus-building, rather than country-specific preferences, makes it unlikely that the final product list reflects narrow national interests.

Uniform implementation: The absence of special treatment provisions and the adoption of a standardized staging schedule minimize country-specific variation, allowing common treatment effects to be identified.

Administrative variation in Attachment B: Cross-country differences in classification practices—function-oriented, material-based, or dual coverage—were driven by administrative customs procedures rather than strategic trade policy. This generated exogenous variation in both the number of countries liberalizing specific HS6 codes and the share of world imports covered.

These institutional features provide a quasi-experimental setting for estimating the effects of coordinated liberalization, and the Attachment B heterogeneity forms the empirical foundation for our "force of many" identification strategy.

3 Data and Descriptive Analysis

We construct a comprehensive dataset spanning 2012–2019, encompassing product-level trade flows, tariffs, ITA Phase II expansion schedules, and standard gravity model variables. This period covers the formal endorsement of the ITA expansion in December 2015 and its initial implementation beginning July 2016. While expansion schedules mandated full tariff elimination by 2020, we exclude 2020 to avoid COVID-19 pandemic disruptions to international trade.

We begin in 2012 to ensure several pre-treatment years while remaining relevant given IT products' rapid innovation cycles. Earlier periods would risk capturing technologically distinct products not included in Phase II. The 2012 start also coincides with the joint concept paper submitted by six ITA members, ensuring that our pre-treatment period reflects the environment relevant to the agreement. Finally, since there was a HS classification revision in 2012, we minimizing concordance issues across HS revisions.

3.1 Data Sources and Construction

Tariff data. To analyze tariff liberalization and uncertainty reduction, we utilize four tariff categories sourced from the WTO Integrated Database (IDB), Consolidated Tariff Schedule (CTS) and ITA Expansion Schedules: preferential ad valorem duties, MFN ad valorem tariffs, bound MFN tariffs, and yearly bound tariff phase-outs as agreed in the ITA.[6] Missing tariff observations represent around 17% of our sample, but are evenly distributed between control and treated goods, alleviating identification concerns (see Table 9). Our main analysis excludes these missing observations to avoid any potential bias from imputation procedures. As a robustness check, we also estimate our models using imputed values with a carry-forward approach based on the most recent available observation, as countries often fail to report unchanged tariff rates following PTA signature (). This imputation procedure reduces missing observations to approximately 10% of the sample, and our main results remain robust to this alternative treatment of missing data.

ITA expansion schedules. We extract expansion schedules from official WTO documentation, including schedules and staging matrices approved by ITA participants. All liberalized products are defined under the common HS2007 nomenclature. Where exceptions are specified within HS6 categories, the data include an "EX" flag to capture partial liberalization for robustness checks. Constructing the dataset required resolving granularity discrepancies, as staging matrices occasionally specified reductions at a finer level of detail than available trade data.

Trade and gravity variables. Trade data, WTO membership information, and gravity-related variables are sourced from CEPII, specifically using BACI for trade flows and the gravity database described in Conte et al. (2022). We extract BACI directly in HS2007 classification to align with the ITA schedules. We do not impose any country-level sample restrictions: the dataset includes 50 ITA members (a full list is provided in the Appendix) and 176 non-members. Among the members is Taiwan. Because the United Nations does not disseminate trade statistics for Taiwan, COMTRADE—and consequently BACI—do not report Taiwan's trade separately. Following CEPII's recommendation, we use trade flows reported under 'Asia, not elsewhere specified' (ISO code 490) as a proxy for Taiwan. While this category could in principle include trade from other unspecified Asian territories, in practice it almost exclusively

[6] preferential ad valorem duties, MFN ad valorem tariffs, bound MFN tariffs, are calculated using tariff line information with time series identifiers HS_A_0015, HS_A_0070, and HS_A_0025.

captures Taiwan's trade, with only a few minor exceptions for specific reporting countries.

Product coverage definition. We define covered ITA products by consolidating Attachment A (191 HS6 subheadings, 50 partially covered) and Attachment B (10 product descriptions mapped to national tariff lines). For Attachment B products, we track which countries liberalized each HS6 code based on their individual mapping decisions described in Section 2. This creates variation in both the extensive margin (whether a product was liberalized) and the intensive margin (by how many countries), which we exploit in our identification strategy. A table that summarise these descriptions is reported in Appendix 10. Even within Attachment A, some exceptions existed at sub-HS6 levels. We tracked these exceptions to control for partial liberalization, though full liberalization represents over 80% of product lines.

Control group construction. To identify ITA Phase II effects, we require a control group of similar products that could have been included in the expansion. We utilize the list of HS6 products proposed for possible Phase III expansion by the Information Technology and Innovation Foundation in collaboration with industrial groups (Ezell and Long, 2023). This list of 151 HS2007 6-digit product lines provides a natural control group because it identifies products that industry stakeholders and policy experts considered suitable candidates for ITA expansion, suggesting similar technological characteristics and trade patterns to actual Phase II products, but which were not ultimately included in the agreement. After excluding products covered under previous phases and Attachments B to avoid capturing earlier partial liberalization, our control group comprises 151 HS2007 codes. We validate this control group by confirming that treated and control products exhibit parallel pre-trends in trade flows and tariff levels during 2012-2015, supporting our identifying assumption.

Product Coverage and Economic Scope Table 1 summarizes the economic scope of the ITA Phase II expansion across product categories. The expansion covered 191 HS6 codes in Attachment A with universal liberalization by all 50 signatories, and 147 codes in Attachment B with variable country participation. Combined, these attachments represented approximately 15% of world merchandise trade in 2015, totaling $2.4 trillion out of $16 trillion in global trade, demonstrating the substantial economic significance of this plurilateral agreement. Notably, despite containing fewer product codes, Attachment B products account for nearly a quarter ($0.56 trillion) of the total covered trade value, reflecting the high-value nature of these technologically advanced goods and highlighting the economic importance of the cross-country variation in their liberalization that we exploit for identification.

Table 1: ITA Phase II Product Coverage and Economic Scope

Product Category	**HS6 codes**	**Liberalization Coverage**	**2015 Trade Value ($ trillions)**
Attachment A	191	Universal (50 countries)	1.43
Attachment B	147	Variable (1–50 countries)	0.56
Control (proposed Phase III)	151	None	1.24
Total ITA Phase II	**338**	–	**1.99**

Notes: Shows distinct HS2007 6-digit product lines by category and corresponding 2015 world import values. Combined Attachments A and B represented approximately 15% of world merchandise trade ($16 trillion) in 2015.

3.2 Descriptive Patterns

Figure 1 illustrates import trajectories and liberalization dynamics around the ITA Phase II expansion. Panel (a) shows import trajectories for treated and control products among ITA members (upper panel) and non-ITA countries (lower panel) from 2012–2019. For ITA members, imports of both product categories exhibited significant growth after 2016, with treated products displaying notably faster rates. The pre-2016 parallel trends between treated and control products support a plausible causal relationship between the ITA expansion and accelerated import growth for liberalized goods.

Interestingly, non-ITA countries also experienced substantial growth in ITA expansion product imports post-2016, outpacing their control group. This pattern suggests potential indirect benefits from non-discriminatory MFN liberalization, possibly through increased international competitiveness enhancing exporter performances.

Panel (b) confirms liberalization dynamics by illustrating MFN tariff evolution. ITA participants implemented linear tariff decreases on covered products post-2016, while non-ITA members maintained parallel tariff trajectories for both treated and control goods throughout the study period.

3.3 Uncertainty Reduction via Bound Tariff Elimination

ITA Phase II required participating countries to eliminate bound MFN tariffs through staged, time-limited reductions. Signatories committed to lowering bound rates in

Figure 1: Evolution of Import Flows and Tariffs Around ITA Phase II

(a) Imports, average inflows

(b) Liberalization, average applied tariff

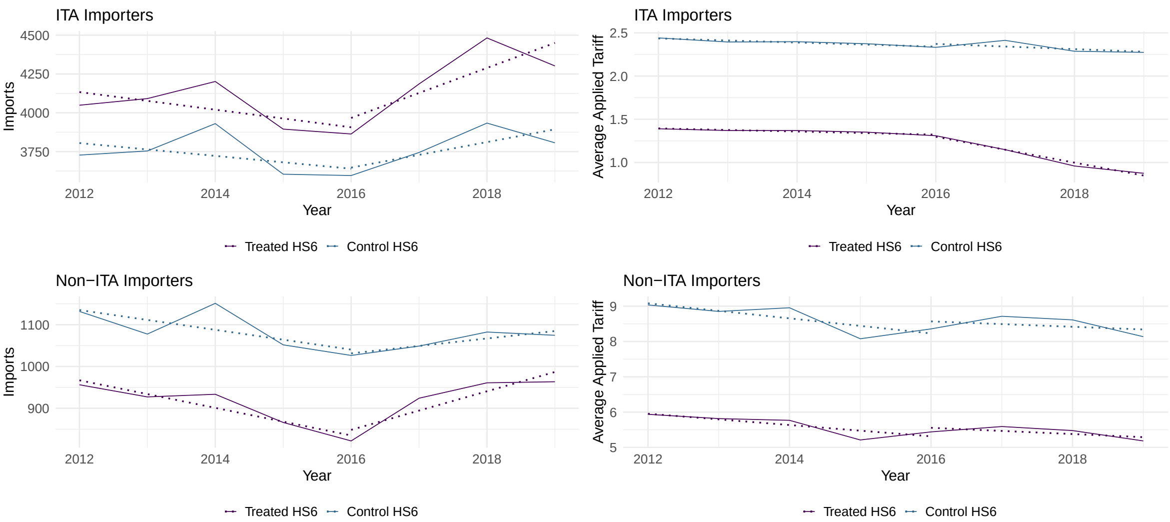

Notes: Imports in thousands of dollars. Applied tariff is the minimum of applied MFN and preferential rates.

four equal annual steps beginning in 2016, with full elimination by 2020. Importantly, reductions in applied MFN tariffs occurred as a direct consequence of these bound commitments, since applied rates had to be adjusted downward in line with the scheduled elimination of bound rates.

Figure 2 illustrates this evolution. Panel (a) shows the distribution of bound MFN tariffs for liberalized products across countries between 2015 and 2019. In 2015, the median bound tariff was around 5%, with an interquartile range of 1–8 percentage points. Both the level and the dispersion declined steadily, with the median and interquartile range approaching zero by 2019.

Panel (b) tracks the share of liberalized product–country pairs still subject to positive bound tariffs. In 2015, roughly 70% had non-zero rates. This share dropped to about 40% in 2016, reflecting the first scheduled reduction, and then fell gradually to around 6% by 2019. In some cases, countries accelerated their schedules—particularly for products with low initial bound rates—by reducing tariffs to zero immediately rather than following the staged path. We account for these accelerations in robustness checks, as they could raise endogeneity concerns if driven by expectations of strong trade effects. Residual positive bound rates after 2019 reflect the seven-year phase-out exceptions granted to developing countries for a limited subset of products, which we exclude from the analysis.

Figure 2: Uncertainty Reduction Over ITA Phase II

(a) Evolution of MFN bindings (b) Share with non-zero bindings

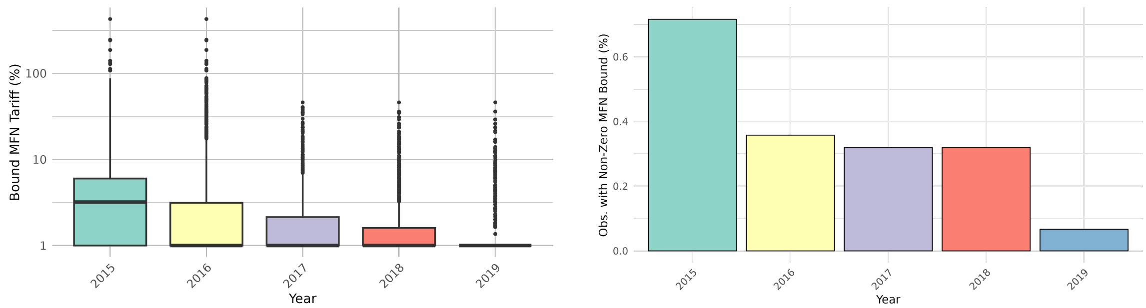

Notes: Includes only HS6 lines without exclusions.

Cross-Country Heterogeneity in Attachment B Products A key element of our empirical strategy exploits the substantial heterogeneity in liberalization patterns across Attachment B products. The classification gap required participants to map product descriptions to existing national tariff lines, generating substantial cross-country variation in HS6 coverage. Table 2 documents this variation across the 147 HS6 lines affected by Attachment B products. The median HS6 code was liberalized by only four countries, with a quarter liberalized by two or fewer countries. At the other extreme, some HS6 lines were liberalized by all 50 participants, creating market share coverage ranging from less than 1% to 100% of world imports. This administrative variation—driven by classification mechanics rather than strategic policy choices—provides the "force of many" variation central to our empirical identification strategy.

Table 2: Panel B: Cross-Product Variation in Attachment B

	Countries liberalizing	**Mkt Share Liberalized**	**Import Value 2015 ($B)**
Minimum	1	0.7%	0.02
25th percentile	2	17.7%	2.1
Median	4	50.5%	4.8
Mean	14.5	50.1%	8.3
75th percentile	32	82.3%	12.7
Maximum	50	100.0%	45.2

Notes: Panel B reports distributional statistics for the number of countries liberalizing each Attachment B HS6 line, their share of world imports (2015), and corresponding 2015 import values.

4 Theoretical Framework

This section develops a theoretical framework analyzing how plurilateral trade agreements generate market access gains through coordinated liberalization. Within the structural gravity framework, we demonstrate that open plurilateral agreements, where countries jointly liberalize on a most-favored-nation basis, create amplification effects that exceed the sum of unilateral liberalizations.

4.1 Gravity Model with Multilateral Resistance

We begin with the canonical structural gravity equation of Anderson and Van Wincoop (2003):

$$x_{ijkt} = \frac{Y_{ikt}E_{jkt}}{Y_{kt}} \left(\frac{\tau_{ijkt}}{P_{jkt}\Pi_{ikt}} \right)^{1-\sigma}, \tag{1}$$

where x_{ijkt} denotes imports of product k by country j from i in year t, Y_{ikt} is exporter i's output, E_{jkt} is importer j's expenditure, Y_{kt} is world output of product k, and $\sigma > 1$ is the elasticity of substitution across source countries and products.

The key insight from structural gravity for analyzing plurilateral agreements lies in the multilateral resistance (MR) terms. The inward MR P_{jkt} is a CES price index capturing the average trade cost that importer j applies to all source countries.. The outward MR Π_{ikt} measures exporter i's average difficulty in accessing markets. These terms are jointly determined by the system:

$$P_{jkt}^{1-\sigma} = \sum_i \left(\frac{\tau_{ijkt}}{\Pi_{ikt}} \right)^{1-\sigma} \frac{Y_{ikt}}{Y_{kt}}, \tag{2}$$

$$\Pi_{ikt}^{1-\sigma} = \sum_j \left(\frac{\tau_{ijkt}}{P_{jkt}} \right)^{1-\sigma} \frac{E_{jkt}}{Y_{kt}}. \tag{3}$$

Because P depends on all exporters' Π's and vice versa, any change in trade costs in one market affects conditions in all others through a network of feedback effects.

4.1.1 Scope and Assumptions

Throughout we work in a structural gravity environment with $\sigma > 1$. We hold $(Y_{ikt}, E_{jkt}, Y_{kt})$ fixed and allow general-equilibrium adjustments only through multilateral resistance terms (P_{jkt}, Π_{ikt}) following the conditional general equilibrium approach of Anderson and Yotov (2016) and the modular trade impact framework of Head and

Mayer (2014). Internal trade $i = j$ enters P_{jkt} with τ_{jjkt}, which is unaffected by MFN tariff changes.

4.2 Trade Cost Structure

To analyze how MFN-based liberalization operates, we decompose bilateral trade costs into components that vary at different levels:

$$\ln \tau_{ijkt} = \ln \tau_{jkt} + \ln \tau_{ijkt}^{\text{bil}} + \kappa_{ijk}. \tag{4}$$

Here κ_{ijk} represents time-invariant country–pair–product factors, such as a common language, colonial ties, or persistent consumption preferences that vary across products but remain stable over time.

4.2.1 Importer Component: MFN Policies

The importer-specific component captures policies applied equally to all trading partners:

$$\ln \tau_{jkt} = \ln(1 + t_{jkt}^{\text{MFN}}) + \ln(1 + \gamma \cdot w_{jkt}), \tag{5}$$

where t_{jkt}^{MFN} is the applied MFN tariff rate, $w_{jkt} = \bar{t}_{jkt} - t_{jkt}^{\text{MFN}}$ represents "tariff water" (the gap between bound and applied rates), and $\gamma > 0$ converts tariff overhang into an ad valorem cost equivalent.[7]

4.2.2 Bilateral Component: Preferential Access

The bilateral component captures country-pair specific frictions, reflecting preferential or discriminatory treatment beyond MFN principles:

$$\ln \tau_{ijkt}^{\text{bil}} = \omega_{ijkt}^{\text{pref}} + \mathbf{Z}'_{ijt}\boldsymbol{\psi}, \tag{6}$$

where $\omega_{ijkt}^{\text{pref}} = \ln \frac{1+t_{ijkt}^{\text{appl}}}{1+t_{jkt}^{\text{MFN}}} \leq 0$ is the *preference wedge* (equal to zero when trade occurs under MFN regime), and $\mathbf{Z}'_{ijt}\boldsymbol{\psi}$ captures other bilateral determinants of trade costs, such as preferential agreements or geographic factors that evolve over time.

[7] We treat both applied tariffs and tariff water in log form for consistency with iceberg trade costs. The applied MFN tariff enters as $\ln(1 + t^{\text{MFN}})$, while tariff water enters as $\ln(1 + \gamma w)$, so that both operate multiplicatively on trade costs. In the empirical analysis, this structure maps directly into using $\ln(1 + t_{jkt}^{\text{MFN}})$ for applied tariffs and $\ln(1 + |\Delta \bar{t}_{jkt}|)$ for binding reductions.

4.3 Three Channels of Plurilateral Liberalization

Having established the trade cost structure, we now examine how plurilateral agreements generate gains through three distinct but interconnected channels that operate simultaneously when countries coordinate their MFN liberalization. First, note that plurilateral agreements like the ITA operate through the importer component of trade costs τ_{jkt}, affecting all bilateral relationships simultaneously rather than creating preferential access. These are discussed in channels 1 and 2, and we then discuss in channel 3 the indirect effects via MR.

4.3.1 Channel 1: Direct Tariff Reduction

The first channel captures the familiar direct effect of cutting applied MFN tariffs. When country j lowers its MFN tariff on product k by $dt_{jkt} < 0$ for all exporters, holding multilateral resistance terms fixed, a log-differentiation gives:

$$d \ln x_{ijkt} = (1 - \sigma) \frac{dt_{jkt}}{1 + t_{jkt}}. \tag{7}$$

Because $\sigma > 1$ and $dt_{jkt} < 0$, this effect is positive: lower tariffs reduce bilateral trade costs and expand imports proportionally to the trade elasticity. The larger the elasticity of substitution σ, the more responsive trade flows are to tariff changes, reflecting greater substitutability between varieties from different sources.

4.3.2 Channel 2: Policy Uncertainty Reduction

Beyond applied tariffs, plurilateral agreements also reduce policy uncertainty by narrowing, on an MFN basis, the gap between bound and applied rates (Handley and Limao, 2015; Handley and Limão, 2017). Let $\bar{t}_{jkt}$ denote the bound rate and $w_{jkt} = \bar{t}_{jkt} - t_{jkt}$ the "tariff water" in the schedule. We model the ad valorem equivalent of policy uncertainty as $\ln(1 + \gamma w_{jkt})$, where $\gamma > 0$ scales the overhang into trade costs.

A reduction in the bound rate ($d\bar{t}_{jkt} < 0$), holding the applied rate fixed, decreases tariff water ($dw_{jkt} < 0$). A log-differentiation gives:

$$d \ln x_{ijkt} = (1 - \sigma) \frac{\gamma \, d\bar{t}_{jkt}}{1 + \gamma w_{jkt}} > 0 \tag{8}$$

This channel operates even without any change in current tariffs: it works by constraining future policy discretion and thereby reducing the risk of trade policy reversals.

4.3.3 Channel 3: Effects via Multilateral Resistance

Beyond the direct channels, trade cost reductions generate indirect effects that operate through multilateral resistance (MR) terms. When a country reduces its import tariffs toward all partners, the full effect can be decomposed as:

$$d \ln x_{ijkt} = (1 - \sigma)\, d \ln \tau_{jkt} + (\sigma - 1) \left(\frac{\partial \ln P_{jkt}}{\partial \ln \tau_{jkt}} + \frac{\partial \ln \Pi_{ikt}}{\partial \ln \tau_{jkt}} \right) d \ln \tau_{jkt}. \tag{9}$$

The first term captures the **direct trade creation effect**: lower bilateral trade costs directly expand imports. The second and third terms represent **general equilibrium adjustments** through multilateral resistance channels.

Inward Multilateral Resistance (P_{jkt}). When country j reduces its MFN tariffs, its inward multilateral resistance P_{jkt} falls as the average cost of sourcing from all suppliers decreases. This adjustment partially offsets the direct trade creation effect by increasing competition in j and making it relatively less attractive. If all sourcing costs into j (including domestic production) fell proportionally, the decline in P_{jkt} would exactly neutralize the direct effect. However, since domestic production typically has a large weight in the price index and its cost does not fall with import tariffs, the offset is incomplete. The net impact on imports remains positive, though smaller than the direct effect alone.

Outward Multilateral Resistance (Π_{ikt}). When only a single country j reduces its MFN tariffs, the impact on any exporter's outward multilateral resistance Π_{ikt} is typically small. This is because Π_{ikt} represents exporter i's average market access cost across all destinations worldwide. Even if country j eliminates all tariffs, this affects only one market among many in the global average, creating minimal change in Π_{ikt}.

However, Π_{ikt} becomes the crucial amplification channel when multiple countries coordinate their liberalization. As we demonstrate in the next section, when a large coalition simultaneously reduces tariffs, the cumulative effect on exporters' global market access costs becomes substantial, generating powerful spillover effects that amplify trade beyond what individual countries could achieve alone.

4.4 Coalition Effects: The Amplification Mechanism

The power of plurilateral agreements lies in **coordination**. When multiple countries simultaneously liberalize, they create spillovers that amplify trade beyond what individual countries achieve alone.

4.4.1 Intuition

Consider what happens when a coalition $\mathcal{C}$ of countries simultaneously cuts MFN tariffs. Each exporter now faces lower costs in multiple markets, reducing their average global market access cost (Π_{ikt} falls more). As exporters' global costs fall, they become more competitive in all markets, increasing competitive global pressure and making each individual country more attractive.

4.4.2 Formal Analysis

Starting from the outward MR condition (3), taking the log-differential with respect to $\tau_{j'kt}$ for all j' in the set of liberalizing countries gives:

$$d\ln \Pi_{ikt} \approx \frac{\sum_{j'} \left(\frac{\tau_{ij'kt}}{P_{j'kt}}\right)^{1-\sigma} \frac{E_{j'kt}}{Y_{kt}} \, d\ln \tau_{j'kt}}{\sum_{j} \left(\frac{\tau_{ijkt}}{P_{jkt}}\right)^{1-\sigma} \frac{E_{jkt}}{Y_{kt}}}. \tag{10}$$

Define the weights representing exporter i's market access share to liberalizing country j' relative to its total global market access:

$$\omega_{ij'kt} \equiv \frac{\left(\frac{\tau_{ij'kt}}{P_{j'kt}}\right)^{1-\sigma} \frac{E_{j'kt}}{Y_{kt}}}{\sum_{j} \left(\frac{\tau_{ijkt}}{P_{jkt}}\right)^{1-\sigma} \frac{E_{jkt}}{Y_{kt}}}. \tag{11}$$

This directly captures how important a liberalizing market is for an exporter relative to all their export opportunities. Substituting this definition gives the compact form:

$$d\ln \Pi_{ikt} \approx \sum_{j'} \omega_{ij'kt} \, d\ln \tau_{j'kt}, \tag{12}$$

showing that the percentage change in Π_{ikt} is a weighted average of the percentage changes in $\tau_{j'kt}$ across all liberalizing importers j'.

Let's define the weights representing importer j's supply access share from exporter i relative to its global suppliers. This captures how important exporter i is as a supplier to importer j relative to all suppliers in the global market.

$$\phi_{ijkt} \equiv \frac{\left(\frac{\tau_{ijkt}}{\Pi_{ikt}}\right)^{1-\sigma} \frac{Y_{ikt}}{Y_{kt}}}{\sum_{i'} \left(\frac{\tau_{i'jkt}}{\Pi_{i'kt}}\right)^{1-\sigma} \frac{Y_{i'kt}}{Y_{kt}}}, \qquad \sum_i \phi_{ijkt} = 1, \tag{13}$$

We can now express the relationship in a compact form, showing how changes in outward multilateral resistances affect the inward MR term.

$$d \ln P_{jkt} \approx -\sum_i \phi_{ijkt} \, d \ln \Pi_{ikt}, \tag{14}$$

This equation shows that the percentage change in P_{jkt} equals the negative weighted average of percentage changes in outward multilateral resistance across all exporters i, where the weights are given by ϕ_{ijkt}.

Interpretation. Under home bias (internal trade has large weight in P_{jkt}), unilateral MFN cuts are partially offset because domestic varieties also enter the price index. Coordination attenuates this offset by improving Π_{ikt} across many exporters at once, thereby strengthening the competitive pull from foreign varieties and deepening the import response in every liberalizing market. The larger the coalition's expenditure share and the broader the product coverage, the stronger these MR feedbacks and the bigger the amplification.

What we take to the data. The framework delivers three estimable channels: (i) direct MFN tariff effects, (ii) uncertainty reductions via bindings, and (iii) coordination effects operating purely through MR. Because MFN policies vary at the importer–product–time level, these channels are naturally identified from variation across (j, k, t) rather than across partners i within a given (j, k, t) dimension. The empirical strategy below builds exactly on this logic.

5 Empirical Strategy

Our theoretical framework identifies three channels through which plurilateral agreements generate market access gains: direct tariff reduction, uncertainty reduction through binding commitments, and general equilibrium amplification via multilateral resistance. Testing these channels empirically requires overcoming a fundamental iden-

Table 3: Channels of MFN Liberalization in an Open Plurilateral (ITA): Members vs. Non-Members

Channel	**Object / Margin**	**Member (liberalizing importer $j \in \mathcal{C}$ and its exporters)**	**Non-member (importer $j \notin \mathcal{C}$ and its exporters)**	**Key refs / sign**
Direct MFN tariff cut	$d\ln\tau_{jkt}$	< 0 for all $i \neq j$; immediate trade creation for x_{ijkt}.	None (no own policy change).	(7): $d\ln x = (1-\sigma)\,d\ln\tau$
Policy uncertainty (bindings)	$d\bar{t}_{jkt}$, dw_{jkt}	$d\bar{t} < 0 \Rightarrow dw < 0$; lowers $\ln(1+\gamma w_{jkt})$ on an MFN basis.	None in own policy; still faces lower uncertainty when buying from members.	(8), (5)
MR (inward)	$d\ln P_{jkt}$	< 0, but *coordination attenuates the offset*: $\lvert d\ln P_{jkt}\rvert$ becomes smaller (can approach 0) as coalition size/coverage grow.	≈ 0 (no own MFN change); only small GE spillovers via exporters' Π.	(??)
MR (outward)	$d\ln \Pi_{ikt}$	< 0 across many i (strong improvement in access to multiple $j \in \mathcal{C}$ simultaneously).	< 0 as well (MFN access to member markets), but *smaller in magnitude* than for member exporters (own country does not liberalize).	(10), (??)
Net effect on x_{ijkt}	Bilateral flow	Large and *amplified*: direct + bindings − attenuated P offset − (smaller) Π offset. Gains rise with coalition expenditure share, size, and product coverage.	Positive spillovers on exports *to* members via $\Pi\downarrow$; no direct/uncertainty effect in own market; overall GE gains smaller than for members.	(??), multiplier $(I-\Omega\Phi)^{-1}$

Notes: $\mathcal{C}$ is the ITA coalition. Under MFN liberalization, $d\ln\tau_{jjkt} = 0$ (internal trade) and $d\ln\tau_{ijkt} = d\ln\tau_{jkt}$ for $i \neq j$. Home bias ($\phi_{jjkt} > 0$) As I remeber, Home bias is an empirical observation in trade literture implying that even after accounting for trade costs domestic trade is much larger than implied by theory. implies unilateral MFN cuts make $d\ln P_{jkt} < 0$, partially offsetting the direct effect; *coordination* reduces this offset because $d\ln\Pi_{ikt}$ becomes more negative across many exporters, entering (**??**) with a minus sign. Both members and non-members enjoy $d\ln\Pi_{ikt} < 0$ from improved access to member markets; for non-members this decline is marginally smaller since their own country does not liberalize. Signs assume $\sigma > 1$.

tification challenge: MFN-based policies affect all trading relationships within each importer–product–time dimension simultaneously, making them perfectly collinear with standard gravity fixed effects.

This section develops a two-stage estimation strategy that exploits the institutional structure of plurilateral agreements to separately identify each channel. We first recover market access effects from a structural gravity equation, then decompose these effects using cross-dimensional variation in policy treatments.

5.1 The Identification Challenge

Standard gravity estimation cannot identify MFN-based policy effects due to their symmetric nature across trading partners. Consider the trade cost decomposition from equation (4), where the importer-specific component τ_{jkt} captures MFN policies applied equally to all partners. In a typical gravity specification,

$$x_{ijkt} = \exp\Big(\alpha_{ikt} + \alpha_{jkt} + \alpha_{ijk} + \beta \cdot \omega_{ijkt}^{\text{pref}} + \text{controls}\Big) \cdot \varepsilon_{ijkt}, \tag{15}$$

the importer–product–time fixed effects α_{jkt} absorb all variation in τ_{jkt}, including MFN tariff changes. The applied tariff coefficient β cannot recover MFN effects because: (i) for country pairs with preferential trade agreements, MFN changes are irrelevant since applied rates are already lower; and (ii) for pairs trading under MFN terms, the policy change is identical across all partners within the same (j, k, t) dimension. Including controls such as PTA membership simply separates preferential from MFN observations, but does not generate any within-cell variation in MFN tariffs. Once α_{jkt} is included, there is no remaining variation at the bilateral level to identify the effect of MFN changes.

5.2 Two-Stage Estimation Framework

5.2.1 Theoretical Foundation

Our solution exploits the fact that importer–product–time fixed effects from structural gravity capture a composite effect:

$$\alpha_{jkt} = \ln\left(\tau_{jkt}^{1-\sigma} \cdot \frac{E_{jkt}}{P_{jkt}^{1-\sigma}}\right). \tag{16}$$

This composite includes both direct trade cost effects (τ_{jkt}) and general equilibrium adjustments through inward multilateral resistance (P_{jkt}). By regressing these recovered fixed effects on policy variables, we can decompose the total effect into constituent channels while preserving the structural interpretation.

5.2.2 Stage 1: Structural Gravity Estimation

We estimate Equation 17 below using Poisson Pseudo Maximum Likelihood (PPML) (Silva and Tenreyro, 2006):

$$x_{ijkt} = \exp\left(\alpha_{ikt} + \alpha_{jkt} + \alpha_{ijk} + \beta \cdot \omega_{ijkt}^{\text{pref}} + \mathbf{Z}'_{ijt}\boldsymbol{\psi}\right) \varepsilon_{ijkt}. \tag{17}$$

Here, α_{ikt} are exporter–product–time fixed effects capturing all supply-side conditions and outward multilateral resistance ($\ln(Y_{ikt}/\Pi_{ikt}^{1-\sigma})$). The α_{jkt} terms are importer–product–time fixed effects that absorb not only demand conditions but also all importer-side policy effects that vary at the (j, k, t) level—including the ITA channels. α_{ijk}, absorb the $(1-\sigma)\kappa_{ijk}$ term from the trade cost structure, capturing time-invariant bilateral trade frictions such as common language, colonial ties, and geographic proximity.

The term $\omega_{ijkt}^{\text{pref}}$ reflects either preferential rates under regional trade agreements (RTAs) or the MFN rate in the absence of preferences. The vector $\mathbf{Z}_{ijt}$ collects bilateral controls that vary within (i, j, t) dimension, such as joint WTO membership or the presence of a preferential trade agreement.

5.2.3 Stage 2: Policy Decomposition

To identify MFN policy effects, we exploit variation across (j, k, t) dimensions by regressing the recovered fixed effects on policy variables:

$$\widehat{\alpha}_{jkt} = \delta_{jk} + \delta_{jt} + \delta_{kt} + \beta_1 \ln(1 + t_{jkt}^{\text{MFN}}) + \beta_2 \ln\big(1 + |\Delta \bar{t}_{jkt}|\big) + \beta_3 \text{ITA}_{jkt} + u_{jkt}. \quad (18)$$

Here, t_{jkt}^{MFN} is the MFN tariff applied by importer j on product k at time t, expressed in decimals. The term $\Delta \bar{t}_{jkt}$ denotes the change in the WTO bound tariff rate, also expressed in decimals. We take the absolute value so that β_2 measures the magnitude of binding reductions; this term is non-zero only for ITA-covered goods affected by Phase II commitments. Finally, ITA_{jkt} is an indicator equal to one if product k is covered under ITA Phase II, importer j is a signatory, and $t > 2015$. [8]

The fixed effects structure is crucial for identification. δ_{jk} (importer–product) fixed effects absorb permanent differences in market access across country–product pairs. δ_{jt} (importer–year) fixed effects control for country-specific macroeconomic trends, exchange rate movements, and other annual shocks. δ_{kt} (product–year) fixed effects capture global product-level shocks such as commodity price changes or technological developments.

Each coefficient in equation (18) maps directly to our theoretical channels:

Channel 1: Direct MFN Tariff Effects (β_1). β_1 captures the total effect of MFN tariff changes, combining direct trade cost reductions with general equilibrium adjustments through inward multilateral resistance. Theory predicts $\beta_1 < 0$ since lower tariffs reduce trade costs and improve market access.

Channel 2: Uncertainty Reduction via Bindings (β_2). $\Delta \bar{t}_{jkt}$ measures changes in bound tariff rates, capturing reductions in policy uncertainty through binding commitments. By controlling for MFN tariffs, $\beta_2 > 0$ reflects that reducing tariff "water" constrains future policy discretion, lowering uncertainty even without immediate tariff cuts.

[8] Following the theoretical specification, we log applied MFN tariffs and measure binding reductions through log changes in the tariff overhang.

Channel 3: General Equilibrium Amplification (β_3). The ITA indicator isolates amplification benefits from simultaneous liberalization. β_3 captures effects operating through the multilateral resistance system: when many countries cut tariffs simultaneously, exporters' outward multilateral resistance Π_{ikt} declines, feeding back into all importers' inward resistance P_{jkt} and amplifying trade beyond unilateral effects. Identification comes from comparing ITA vs. non-ITA products within countries, ITA vs. non-ITA members within products, and pre- vs. post-implementation periods.

Identification Strategy Our identification relies on three sources of variation in the panel. First, *cross-product* variation within importers: each country imports both ITA and non-ITA products, enabling within-country comparisons while controlling for importer–time shocks via δ_{jt}. Second, *cross-country* variation within products: each product is imported by both ITA members and non-members, enabling within-product comparisons that control for product–time shocks via δ_{kt}. Third, *time* variation: the staged implementation of ITA coverage between 2015–2019 generates within-(importer, product) changes in treatment status, helping to separate policy effects from general performance dynamics; time-invariant importer–product heterogeneity is absorbed by δ_{jk}.

The multi-dimensional nature of this variation, combined with our comprehensive fixed-effects structure, ensures that estimated coefficients reflect plausibly exogenous policy changes rather than confounding factors. The approach is particularly powerful because it leverages the institutional features of plurilateral agreements—where participation decisions are made at the agreement level rather than the product level—to generate clean identification of the channels.

6 Results

This section presents our empirical findings. We first document how market access evolved around the ITA Phase II expansion using a dynamic difference-in-differences design. We then turn to the second-stage decomposition, which isolates the three channels from Section 4 and quantifies their relative importance.

6.1 Evolution of Market Access

The first-stage estimates in equation 17 yield importer–product–time fixed effects, $\hat{\alpha}_{jkt}$, which capture all destination–product–specific determinants of trade, including

Figure 3: Market Access Evolution: Liberalized Under ITA vs. Control Goods

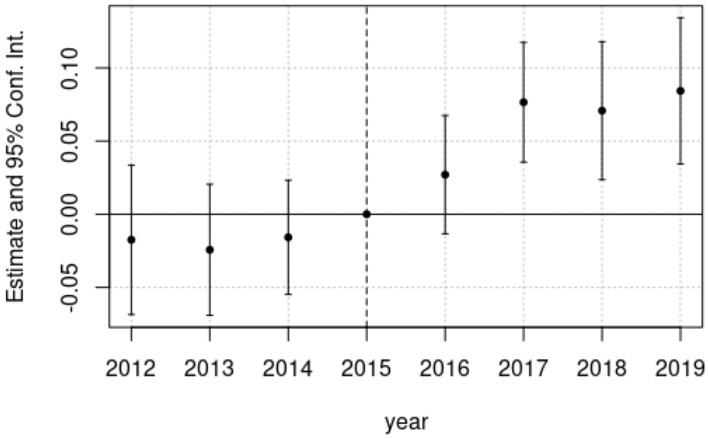

both partial- and general-equilibrium responses to MFN policy changes. While these fixed effects are measured relative to a reference country–product pair and cannot be interpreted as absolute market access levels, their dynamics are informative.

To visualize the aggregate response to ITA Phase II, we compare $\hat{\alpha}_{jkt}$ for fully liberalized goods against a control group of non-liberalized goods. We estimate a dynamic difference-in-differences specification:

$$\hat{\alpha}_{jkt} = \delta_{jk} + \delta_{kt} + \sum_{n=-3}^{4} \beta_n \, \mathbf{1}_{\{\text{event}_t = n\}} \times \mathbf{1}_{\{\text{treat}_{jk}\}} + \mathbf{X}_{jpt} + \epsilon_{jkt}, \tag{19}$$

where n indexes years relative to 2015 (the last pre-treatment year), $\mathbf{1}\{\text{treat}_{jk}\}$ flags ITA Phase II products, and $\mathbf{X}_{jpt}$ controls for importer GDP and product-level exports. Standard errors are clustered at the importer–treatment–year level.

Figure 3 shows the estimated β_n coefficients and 95% confidence intervals. Pre-treatment coefficients are close to zero and statistically insignificant, supporting the parallel-trends assumption. Following the partial implementation in 2016, effects are positive but modest—about 3%—and not statistically significant. Starting in 2017, as the scheduled four-step tariff reductions progressed, effects rise to 6–8% and remain significant through 2019. The shape of the response closely tracks the ITA's implementation schedule, consistent with a causal interpretation.

Overall, the event-study evidence indicates that ITA Phase II liberalization produced economically meaningful and sustained market access gains for covered products, with no sign of anticipatory behavior or confounding pre-trends.

6.2 Decomposing the ITA Effect

To disentangle the contributions of the three policy channels discussed in the theoretical framework, we regress the estimated importer–product–year fixed effects $\hat{\alpha}_{jkt}$ from the first stage on the applied MFN tariff rate, the change in bound tariffs, and an indicator for ITA Phase II product coverage. The resulting coefficients map directly to the theoretical mechanisms outlined in Section 4, providing an empirical breakdown of how tariff liberalization, policy-uncertainty reduction, and coordinated liberalization each shape market-access outcomes, as reported in the second-stage estimates of Table 4.

We now discuss the results for each channel in turn.

Column (1) captures the total market-access effect of MFN tariff reductions, combining the partial-equilibrium trade-creation effect from lowering τ_{jkt} with the general-equilibrium adjustment via P_{jkt}. The coefficient of -0.3474 implies that a 10% decrease in the applied MFN tariff rate is associated with an increase in market access of approximately $0.3474 \times \ln(0.90) \approx 3.7\%$, holding other factors constant.

Adding the change in bound tariffs in Column (2) yields a coefficient of $\hat{\beta}_2 = 0.1664$ on $\ln\big(1 + d\bar{t}_{jkt}\big)$. For small changes, this implies that a one–percentage-point cut in the bound rate (i.e., $d\bar{t} = 0.01$) raises market access by roughly 0.17%, holding the MFN tariff constant. For a larger, ten–percentage-point cut—e.g., from 10% to 0%—the exact effect is $\Delta\widehat{\alpha} = 0.1664 \times \ln(1.10) = 0.01586 \Rightarrow \exp(0.01586) - 1 \approx 0.01598$ (about 1.60%).

This supports the prediction that narrowing tariff "water" boosts trade by reducing uncertainty about future policy reversals.

Column (3) adds the ITA coverage indicator. The estimated coefficient of 0.0192 is a semi-elasticity: switching from non-covered to covered products is associated with an increase in market access of $(\exp(0.0192) - 1) \cdot 100) \approx 1.94\%$, holding applied tariff and bound-rate changes constant. This aligns with the theory: simultaneous MFN liberalization by a large number of countries lowers global outward multilateral resistance and amplifies bilateral trade gains through the coalition-size mechanism described in Section 4.

Robustness. Columns (4) and (5) add two controls.

odel (5) adds two controls to the regression: a dummy variable for zero MFN tariffs ($\mathbf{1}\{t_{ikt}^{\text{MFN}} = 0\}$) and the logarithm of country exports at the product level ($\ln(\text{Export}_{ikt})$) as a proxy of import expenditure.

The inclusion of the first control is supported by previous literature. However, it should be noted that, unlike Gnutzmann-Mkrtchyan and Henn (2018), we do not find any significant results for the zero MFN tariff dummy. There are two possible explanations for these differing findings. First, by not accounting for the value of the MFN-Bound tariff in the staging schedules, Gnutzmann-Mkrtchyan and Henn (2018) may have captured cases of large short-term uncertainty reduction potential in the zero MFN dummy. Second, with respect to the 20 years since Phase I's implementation in 1996, the role of reduced administrative costs associated with less frequent (rigorous) inspections for goods subject to zero MFN tariffs may have diminished.

The second control variable accounts for any variability related to import expenditure at the importer-product level that is not captured by the fixed effects included in the regression, thus ending up in the error term. This might be a concern if import expenditure dynamics were also related to how countries selected the list of products to be treated. As we argued previously, we do not consider this a serious concern given the inclusive nature of the negotiation process. In support of this argument, we observe that our results remain robust to the inclusion of country-product exports.

Table 4: Unpacking the Different Channels in the ITA Phase II Expansion

Dependent Variable:	Market Access: $\hat{\alpha}_{jkt}$				
Model:	(1)	(2)	(3)	(4)	(5)
$\ln(1+t^{\text{MFN}}_{jkt})$	-0.347***	-0.3454***	-0.3377***	-0.3365***	-0.3197***
	(0.0975)	(0.0975)	(0.0973)	(0.0971)	(0.1022)
$\ln(1+\lvert d\bar{t}_{jkt}\rvert)$		0.1664***	0.1381***	0.1384***	0.1378***
		(0.0557)	(0.0522)	(0.0522)	(0.0527)
ITA_{jkt}			0.0192***	0.0191***	0.0188**
			(0.0073)	(0.0073)	(0.0073)
Export_{jkt} (log)				0.0049***	0.0049***
				(0.0005)	(0.0005)
$1_{\{\text{MFN}=0\}}$					0.0042
					(0.0102)
Fixed-effects					
Importer-Time FE	Yes	Yes	Yes	Yes	Yes
Importer-Product FE	Yes	Yes	Yes	Yes	Yes
Product-Time FE	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	438,963	438,951	438,951	438,951	438,951
R^2	0.94361	0.94363	0.94363	0.94367	0.94367
Within R^2	4.45×10^{-5}	5.1×10^{-5}	6.15×10^{-5}	0.00074	0.00074

Notes: The dependent variable is the importer-product-time FE estimated from the gravity Equation 17. Clustered standard errors at the treatment-importer-time level in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

First, a zero-MFN-tariff dummy is insignificant, in contrast to Gnutzmann-Mkrtchyan and Henn (2018); the difference likely reflects our explicit control for binding cuts, which separates uncertainty effects from zero-tariff status. Second, including importer–product log exports leaves the main coefficients virtually unchanged, addressing potential concerns about expenditure-side product selection.

Relative importance of channels. To assess the economic significance of each channel, we calculate the contribution of the ITA Phase II liberalization episode. The average ITA participant reduced MFN tariffs by 2.6 percentage points and bound tariffs by 7.6 percentage points on covered products. Applying our estimates from Model (3), the direct tariff channel generated market access gains of approximately 0.9% (= 0.35 × ln(1.026)), while uncertainty reduction contributed 1.2% (= 0.17 × ln(1.076)). The coordination benefit added 2.0% beyond these direct effects. In total, the three channels combined explain a 4.1% improvement in market access for ITA products. Notably, coordination benefits emerge as the dominant channel, accounting for nearly half (48%) of the total predicted effect, followed by uncertainty reduction (30%) and direct tariff cuts (22%).

Table 5: Economic Magnitude of ITA Phase II Liberalization Channels

Channel	**Policy Change**	**Coefficient**	**Market Access Impact**	**Share of Total Effect**
Direct tariff reduction	2.60 p.p. MFN tariff cut	−0.34	+0.90%	22%
Uncertainty reduction	7.58 p.p. binding cut	0.14	+1.24%	30%
Coordination benefits	ITA membership	0.0200	+2.00%	48%
Total predicted effect			**+4.14%**	**100%**

Notes: Headline policy changes use Variant A (base > 0 and actual cut by 2019, equal-country averaging). Impacts computed as $|\beta_{\text{MFN}}| \cdot \ln(1 + \Delta\text{MFN})$, $\beta_{\text{BND}} \cdot \ln(1 + \Delta\text{Bound})$, and the ITA coefficient. Coefficients from Table 4, Model (3).

7 The Economics of Coordinated Liberalization

Our empirical results point to substantial coordination benefits from ITA Phase II, market access gains that exceed the sum of what countries would achieve by acting alone. This section develops the economic intuition for this finding and links it directly to the multilateral resistance structure of the gravity model.

7.1 Intuition: Why Coordination Matters

In the gravity framework, trade depends not only on the bilateral cost between two countries, but on that cost *relative* to the importer's average cost of sourcing from all potential suppliers. Unilateral liberalization improves a country's attractiveness, but exporters' average access to world markets may remain constrained if other major markets stay closed. Coordinated liberalization changes this relative landscape everywhere, amplifying the gains.

Consider semiconductor tariffs. In one scenario, the United States and the European Union cut their tariffs from 5% to zero. In another, China joins the coalition and makes the same cut. In both cases, lower bilateral costs generate a *direct trade creation effect*: importers in liberalizing countries face lower prices and expand purchases. But the gravity model also embeds a general equilibrium channel: when multiple countries liberalize together, exporters' *average* market access improves, lowering their *outward multilateral resistance.* This improvement feeds into all importers' CES price indices, shifting their *inward multilateral resistance* - the measure of how accessible they are relative to the rest of the world.

The difference between the two scenarios is the magnitude of this shift. With only the U.S. and EU liberalizing, outward resistance falls modestly. Adding China produces a much larger drop, reshaping the competitive landscape. Conditional on its own policy, the U.S. becomes relatively more "distant" in multilateral terms as other large markets open. Exporters facing tougher competition in the EU and China may redirect some shipments toward the U.S., where relative profitability improves. This *general equilibrium trade diversion* arises entirely from the global reshuffling of market access conditions.

In our empirical framework, these shifts are captured by the importer–product–year fixed effects. Coordinated MFN liberalization by a large number of countries moves these fixed effects more than unilateral or small-coalition liberalization, exactly as predicted by the theory.

7.2 Testing the Coalition Size Mechanism

We test whether the effect of ITA membership depends on the size of the implementing coalition by extending equation (18) to include an interaction between the ITA indicator and a measure of coalition size:

$$\widehat{\alpha}_{jkt} = \delta_{jk} + \delta_{jt} + \delta_{kt} + \beta_1 \ln\big(1 + t^{\mathrm{MFN}}_{jkt}\big) + \beta_2 \ln\big(1 + |d\bar{t}_{jkt}|\big) \\ + \beta_3\, \mathrm{ITA}_{jkt} + \beta_4\, \mathrm{ITA}_{jkt} \times \ln\big(1 + \mathrm{CoalitionSize}_{kt}\big) + \varepsilon_{jkt}, \quad (20)$$

where *CoalitionSize* is measured in two ways: (i) the number of other countries implementing tariff cuts on product k in year t (leave–one–out count), and (ii) the share of world imports of k (excluding the importer) accounted for by the liberalizing coalition, using pre–implementation (2015) import shares.

The source of variation differs by measure and product group. For the *count* measure, cross–country variation exists *only* for Attachment B products: in 2016, members either added a product through Attachment B or did not, so the number of liberalising countries varies across products in this group but is identical across importers for Attachment A products (apart from the common staging schedule over time). For the *market–share* measure, there is product–level variation even within Attachment A, because the 2015 distribution of world imports differs across products, but this variation is larger for Attachment B and is directly linked to differences in the number of liberalising countries.

To prevent the coalition–size coefficient β_4 from picking up structural differences between product groups, some specifications include an ITA $\times$ AttachmentB interaction. This absorbs any systematic difference in the average ITA effect between Attachments A and B (e.g., from technology, trade intensity, or classification rules), ensuring that β_4 is identified from within–group variation in coalition size—by design, driven entirely by Attachment B products for the count measure and by both groups for the market–share measure.

7.3 Results

Table 6 reports the results. Column (1) shows the baseline specification without coalition-size interactions, in which MFN tariff cuts and bound tariff reductions both have the expected signs and are statistically significant. The ITA coefficient is small and statistically insignificant, suggesting that, on average, the direct ITA effect (conditional on other channels) is close to zero when coalition size is accounted for.

Columns (2) and (4) introduce the coalition-size interaction term for the two alternative measures. The interaction is positive and statistically significant for both definitions: $\hat{\beta}_4 = 0.0219$ ($p < 0.05$) when size is measured as the number of liberalizing countries, and $\hat{\beta}_4 = 0.2542$ ($p < 0.01$) when size is measured as the coalition's share of global imports. In both cases, the standalone ITA coefficient becomes more negative, consistent with the interpretation that the average ITA effect depends strongly on coalition size: for small coalitions, net coordination effects may be negative, but as the coalition grows, the interaction term dominates, turning the net effect positive.

The Attachment B dummy is small and statistically insignificant in both specifications where it is included (Columns 2 and 4), suggesting that differences between Attachments A and B products are not a major driver of the observed coordination effects.

Table 6: Coalition Size and Coordination Benefits

Dependent Variable:	Market Access: $\hat{\alpha}_{jkt}$			
Model:	(1)	(2)	(3)	(4)
	Coalition Size: # of Liberalizing Countries		*Coalition Size: Mkt Share of Liberalizing Coalition*	
$\ln(1+t^{\text{MFN}}_{jkt})$	-0.3356***	-0.3379***	-0.3336***	-0.3376***
	(0.0973)	(0.0974)	(0.0972)	(0.0973)
$\ln(1+\lvert d\bar{t}_{jkt}\rvert)$	0.1369***	0.1403***	0.1321**	0.1365**
	(0.0527)	(0.0535)	(0.0553)	(0.0563)
ITA_{jkt}	-0.0329	-0.0670	-0.1051***	-0.1293***
	(0.0317)	(0.0417)	(0.0387)	(0.0424)
$\text{ITA}_{jkt} \times \ln(1+\text{Coalition Size})$	0.0139*	0.0219**	0.2211***	0.2542***
	(0.0083)	(0.0102)	(0.0682)	(0.0722)
$\text{ITA}_{jkt} \times$ Appendix B Dummy		0.0137		0.0178
		(0.0136)		(0.0116)
Fixed effects:				
Importer–Time FE	Yes	Yes	Yes	Yes
Importer–Product FE	Yes	Yes	Yes	Yes
Product–Time FE	Yes	Yes	Yes	Yes
Observations	438,936	438,936	438,936	438,936
R^2	0.94363	0.94363	0.94363	0.94363
Within R^2	6.52×10^{-5}	6.72×10^{-5}	9.40×10^{-5}	9.88×10^{-5}

Notes: The dependent variable is the importer–product–time FE estimated from the gravity Equation 17. Coalition Size is measured either as the number of other countries that have implemented (cols. 1–2) or the pre-2015 import market share of the liberalizing coalition (cols. 3–4). Clustered standard errors at the treatment–importer–time level in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Quantifying the critical mass. For the global market share measure, we can use the specification in Column (4) of Table 6 to compute the threshold at which the net coordination effect of ITA participation turns positive. We solve for S^* in:

$$\hat{\beta}_{\text{ITA}} + \hat{\beta}_{\text{ITA}\times\text{Share}} \cdot \ln(1+S^*) = 0,$$

where S^* is the coalition's share of global imports for the product in question, fixed at 2015 pre-ITA levels.

Using the Column (4) coefficients:

$$-0.129284 + 0.254224 \cdot \ln(1+S^*) = 0,$$

$$\ln(1+S^*) = \frac{0.129284}{0.254224} \approx 0.5087,$$

$$1 + S^* \approx e^{0.5087} \approx 1.663 \quad \Rightarrow \quad S^* \approx 0.663.$$

Thus, the critical mass is approximately 66% of world imports for the relevant product. When the liberalizing coalition accounts for more than this share, the interaction term outweighs the negative baseline ITA effect, and the net coordination benefit turns positive. The actual ITA Phase II coalition, typically covering over 80%

of global IT trade for included products, comfortably exceeded this threshold.

8 Alternative Channels and Robustness

In this section we first show robustness checks to enhance the validity of our results.

8.1 Alternative Channels

Heterogeneous trade elasticity. We have assumed that elasticities are common across products as they all belong to IT group. We here relax this assumption and test whether our results are driven by heterogeneity in trade elasticity. Using HS6-specific trade elasticities from Fontagné et al. (2022), we interact the ITA indicator with product-level elasticity measures. Table 7, column 1 and 2, reveals that ITA benefits are indeed heterogeneous across elasticity levels: products with higher trade elasticities exhibit smaller coordination effects, consistent with the theoretical prediction that policy interventions have diminished impact in markets where competition is already intense due to high substitutability between varieties. Specifically, the interaction coefficient of -0.006 indicates that for each unit increase in trade elasticity, the ITA's positive effect decreases by approximately 0.6 percentage points. This heterogeneity aligns with our multilateral resistance framework, where coordination benefits should be largest in differentiated product markets with lower baseline competition intensity. Importantly, controlling for this elasticity heterogeneity does not eliminate our main coordination effect, which remains positive and significant, confirming that our results are not simply driven by unobserved variation in demand-side substitutability across product categories.

Supply Chain Position and Vertical Linkages. A key alternative explanation for our coordination effects operates through supply chain mechanisms rather than multilateral resistance. In vertically integrated industries, tariff reductions create spillovers both downstream (lower input costs) and upstream (increased derived demand) (**?**). If ITA products liberalized by larger coalitions systematically occupy specific supply chain positions, our coordination benefits might reflect vertical linkage amplification rather than general equilibrium trade creation through multilateral resistance. To test this channel, we employ upstreamness measures from Antràs et al. (2012) that quantify each product's position in global supply chains. Table 7, columns 3-4, interacts the ITA indicator with product-level upstreamness. Column 3 shows upstream products (above median upstreamness) exhibit larger ITA benefits (0.031 vs. 0.018), consistent with input cost spillovers. However, column 4 demonstrates that controlling for supply chain heterogeneity does not impact substantially on our baseline coordination effect, which remains positive and highly significant. This suggests vertical linkages contribute to ITA benefits but do not account for the coordination patterns we document, supporting our multilateral resistance interpretation.

Table 7: Coordination Benefits with Alternative Channels Controlled

Dependent Variable:	Market Access: $\hat{\alpha}_{jkt}$			
Model:	(1)	(2)	(3)	(4)
Variables				
$\ln(1+t^{\text{MFN}}_{jkt})$	-0.3331***	-0.3279***	-0.3328***	-0.3297***
	(0.0970)	(0.0969)	(0.0971)	(0.0970)
$\ln(1+\Delta\bar{t}_{jkt})$	0.1280**	0.1199**	0.1259**	0.1214**
	(0.0518)	(0.0550)	(0.0529)	(0.0552)
ITA_{jkt}	0.0445***	-0.0861**	0.0582***	-0.0595
	(0.0122)	(0.0390)	(0.0217)	(0.0463)
$\text{ITA}_{jkt} \times \sigma_k$	-0.0029**	-0.0034***		
	(0.0012)	(0.0012)		
$\text{ITA}_{jkt} \times \ln(1+\text{Share}_{k,2015})$		0.2391***		0.1957***
		(0.0696)		(0.0689)
$\text{ITA}_{jkt} \times \text{Upstream}_k$			-0.0213**	-0.0172
			(0.0107)	(0.0107)
Fixed effects:				
Importer–Time FE	Yes	Yes	Yes	Yes
Importer–Product FE	Yes	Yes	Yes	Yes
Product–Time FE	Yes	Yes	Yes	Yes
Fit statistics:				
Observations	438,951	438,951	438,444	438,444
R^2	0.94363	0.94364	0.94366	0.94366
Within R^2	7.73×10^{-5}	1.10×10^{-4}	6.93×10^{-5}	9.45×10^{-5}

Notes: The dependent variable is the importer–product–time fixed effect from the first-stage gravity estimation. $\Delta\bar{t}_{jkt}$ denotes the change in the bound MFN rate (in fractions). σ_k is the product-level elasticity proxy; $\text{Share}_{k,2015}$ is the pre-2015 import market share; Upstream_k measures upstreamness. Standard errors clustered at the treatment–importer–time level are in parentheses. Signif. codes: ***: 0.01, **: 0.05, *: 0.1.

8.2 Robustness to Confounding Shocks and Sample Restrictions

The US-China Trade War. A key concern is that our results are confounded by the US-China trade war, which began in 2018 and directly affected many IT-related goods. For example, the US imposed Section 301 tariffs on Chinese products, including semiconductors, and implemented export restrictions on Huawei that had global spillovers. To ensure our results are not driven by these bilateral tensions, we re-estimate our baseline specification while excluding both the US and China from our sample. Table 8, column 1, shows that our key coefficients—the direct tariff effect, the uncertainty reduction effect, and the plurilateral effect—remain statistically significant and of similar magnitude to our main findings. This demonstrates that the ITA Phase II's benefits are a genuine multilateral phenomenon, robust to the most significant trade policy shocks of the period.

Exclusion of Major Players. We also consider the possibility that our results are driven by the strategic behavior of key countries. The ITA's negotiation process was led by a group of countries (Canada, Japan, South Korea, Singapore, Chinese Taipei, and the United States) that signed the 2012 "Concept Paper," while a major player, India, opted out. To test for potential selection bias, we re-estimate our model after applying two sample restrictions: first, we exclude all of the 2012 signatory countries from the ITA participants, and second, we exclude India from the control group. In both cases, the coefficients of interest remain significant and stable (Table 8), reinforcing our conclusion that the plurilateral effect is a broad-based phenomenon rather than a result of the strategic choices of a few key players.

Table 8: MFN Tariffs, Uncertainty, and ITA Coordination

Dependent Variable:	Market Access: $\hat{\alpha}_{jkt}$					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Sample Exclusion:	USA & CHN		All Active Members		INDIA	
$\ln(1+t^{\text{MFN}}_{jkt})$	-0.3476***	-0.3434***	-0.3363***	-0.3324***	-0.3295***	-0.3252***
	(0.0986)	(0.0985)	(0.0979)	(0.0979)	(0.0978)	(0.0977)
ITA_{jkt}	0.0198***	-0.1034***	0.0241***	-0.0948**	0.0201***	-0.1072***
	(0.0074)	(0.0388)	(0.0075)	(0.0409)	(0.0074)	(0.0390)
$\ln(1+\Delta\lvert t^{\text{BND}}_{jkt}\rvert)\times\text{Treat}_{jkt}$	0.1422***	0.1360**	0.1394***	0.1328**	0.1388***	0.1326**
	(0.0525)	(0.0556)	(0.0531)	(0.0560)	(0.0520)	(0.0551)
$\text{ITA}_{jkt}\times\ln(1+\text{Share}_{2015})$		0.2192***		0.2116***		0.2263***
		(0.0684)		(0.0722)		(0.0687)
Fixed effects:						
Importer–Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Importer–Product FE	Yes	Yes	Yes	Yes	Yes	Yes
Product–Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	433,850	433,850	422,267	422,267	435,150	435,150
R^2	0.94153	0.94153	0.93856	0.93856	0.94305	0.94305
Within R^2	6.43×10^{-5}	9.61×10^{-5}	6.71×10^{-5}	9.48×10^{-5}	6.04×10^{-5}	9.45×10^{-5}

Notes: The dependent variable is the importer–product–time fixed effect estimated from the gravity model (Equation 17). Columns 1–2 exclude USA and China; Columns 3–4 exclude all active members; Columns 5–6 exclude India. $\ln(1+t^{\text{MFN}}_{jkt})$ is the log MFN ad valorem tariff. $\ln(1+\Delta\lvert t^{\text{BND}}_{jkt}\rvert)\times\text{Treat}_{jkt}$ measures the effect of changes in bindings (policy uncertainty) interacted with treatment. ITA_{jkt} is an indicator for products covered by the ITA, and Share_{2015} is the pre-2015 import share. Clustered standard errors at the treatment–importer–time level in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1.

Phased versus Immediate Liberalization. The majority of ITA Phase II products were subject to a linear, four-year tariff phase-out schedule. However, for a small

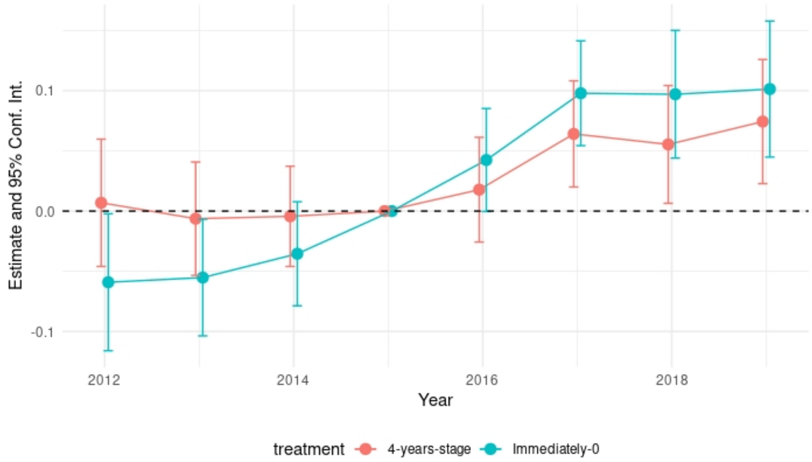

Figure 4: Market Access by Treatment type, 4-years-stage vs Immediately Zero

number of products, tariffs were eliminated immediately upon implementation. The decision to zero-out tariffs immediately could be endogenous to the expected market access gains for those products. We test whether our results are driven by this potentially endogenous subgroup by re-estimating our baseline model while excluding all products that received immediate tariff elimination. Figure 4 confirms that both groups followed similar pre-treatment trends.

Partial Liberalization. Some products were liberalized with exclusions, allowing countries to exclude parts of coarser tariff lines. Since our trade data is at the HS6 level, we cannot determine the exact importance of these finer tariff lines. To account for this, we add a control dummy for all HS6 codes with exclusions to see if these products exhibit different coordination effects. The results, reported in Table **??**, show that our main findings are robust to this control, suggesting that these partial liberalizations do not substantially alter our conclusions.

Standard Error Clustering. Finally, we verify the stability of our results by re-estimating all regressions using different standard error clustering structures to account for potential serial correlation in market access. We experiment with various multi-way clustering schemes, including product-country and importer-time. Our results remain consistent and statistically significant across these alternative specifications, confirming the robustness of our main findings.

9 Conclusion

This paper has examined how plurilateral agreements generate trade gains that go beyond the effects of unilateral or bilateral liberalization. Using the Information Technology Agreement as a motivating case, we developed a structural gravity framework that decomposes the effects of coordinated MFN liberalization into three channels: direct tariff reductions, reductions in policy uncertainty, and a general-equilibrium amplification mechanism that arises only when countries liberalize together.

Our empirical analysis, which exploits variation in ITA commitments across countries, products, and time, confirms the presence of all three channels. Direct tariff cuts raise trade flows substantially, while bound-rate reductions generate additional gains through reduced policy uncertainty. Most importantly, we show that coordinated liberalization produces a measurable coordination premium: coalition-wide tariff cuts amplify trade responses well beyond the sum of isolated national reforms. This amplification operates through multilateral resistance, highlighting the systemic nature of plurilateral liberalization.

These results contribute to the literature in two ways. Theoretically, we extend the structural gravity framework to show how coordination changes the role of multilateral resistance, creating non-additive gains from joint action. Empirically, we provide new evidence on the trade effects of plurilateral agreements, an increasingly important instrument in the world trading system.

From a policy perspective, our findings underscore the value of plurilateral initiatives as a pragmatic vehicle for advancing liberalization in a multilateral setting under deadlock. By combining MFN treatment with coalition-based coordination, plurilaterals can deliver trade expansion on a scale that individual reforms cannot achieve. For future research, our framework can be extended to quantify distributional effects across firms and sectors, and to analyze how plurilateral rules interact with the growing landscape of preferential trade agreements.

In sum, plurilateral agreements do not merely aggregate national reforms; they amplify them. Recognizing and harnessing this coordination premium is key to sustaining liberalization in the global trading system.

References

Anderson, J. E. and E. Van Wincoop (2003): "Gravity with gravitas: A solution to the border puzzle," *American economic review*, 93, 170–192.

Anderson, J. E. and Y. V. Yotov (2016): "Terms of trade and global efficiency effects of free trade agreements, 1990–2002," *Journal of International Economics*, 99, 279–298.

Antràs, P., D. Chor, T. Fally, and R. Hillberry (2012): "Measuring the upstreamness of production and trade flows," *American Economic Review*, 102, 412–416.

Bagwell, K., C. P. Bown, and R. W. Staiger (2016): "Is the WTO passé?" *Journal of economic literature*, 54, 1125–1231.

Caliendo, L., R. C. Feenstra, J. Romalis, and A. M. Taylor (2015): "Tariff Reductions, Entry, and Welfare: Theory and Evidence for the Last Two Decades," Working Paper 21768, National Bureau of Economic Research.

Conte, M., P. Cotterlaz, T. Mayer, et al. (2022): *The CEPII gravity database*, CEPII.

Eicher, T. S. and C. Henn (2011): "In search of WTO trade effects: Preferential trade agreements promote trade strongly, but unevenly," *Journal of International Economics*, 83, 137–153.

Ezell, S. and T. Long (2023): "How Expanding the Information Technology Agreement to an "ITA-3" Would Bolster Nations' Economic Growth," Tech. rep., Information Technology and Innovation Foundation.

Fontagné, L., H. Guimbard, and G. Orefice (2022): "Tariff-based product-level trade elasticities," *Journal of International Economics*, 137, 103593.

Gnutzmann-Mkrtchyan, A. and C. Henn (2018): "Peeling away the layers: Impacts of durable tariff elimination," *Journal of International Economics*, 115, 259–276.

Handley, K. and N. Limao (2015): "Trade and investment under policy uncertainty: theory and firm evidence," *American Economic Journal: Economic Policy*, 7, 189–222.

Handley, K. and N. Limão (2017): "Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States," *American Economic Review*, 107, 2731–2783.

Handley, K. and N. Limão (2022): "Trade Policy Uncertainty," *Annual Review of Economics*, 14, 363–395.

Head, K. and T. Mayer (2014): "Chapter 3 - Gravity Equations: Workhorse,Toolkit, and Cookbook," in *Handbook of International Economics*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Elsevier, vol. 4 of *Handbook of International Economics*, 131–195.

Heid, B., M. Larch, and Y. V. Yotov (2021): "Estimating the effects of non-discriminatory trade policies within structural gravity models," *Canadian Journal of Economics/Revue canadienne d'économique*, 54, 376–409.

Larch, M., J.-A. Monteiro, R. Piermartini, and Y. V. Yotov (2019): "On the Effects of GATT/WTO Membership on Trade: They Are Positive and Large After All," Working Paper.

Mann, C. and X. Liu (2009): "The information technology agreement: sui generis or model stepping stone," *Multilateralizing regionalism: Challenges for the global trading system*, 182–216.

Rose, A. K. (2004): "Do We Really Know That the WTO Increases Trade?" *American Economic Review*, 94, 98–114.

Silva, J. M. C. S. and S. Tenreyro (2006): "The Log of Gravity," *The Review of Economics and Statistics*, 88, 641–658.

Subramanian, A. and S.-J. Wei (2007): "The WTO promotes trade, strongly but unevenly," *Journal of International Economics*, 72, 151–175.

A Appendix

Table 9: Share of Missing Applied Tariffs (%) by ITA Status × Treatment: Original vs. After Fill

Group	**Applied tariff (orig)**	**Applied tariff2 (after fill)**
ITA member, Treated (1)	2.03%	0.67%
ITA member, Control (2)	3.05%	0.92%
Non-ITA, Treated (3)	30.21%	18.55%
Non-ITA, Control (4)	31.46%	19.13%
Overall	17.95%	10.69%

Table 10: Selected Products Covered in Attachment B of the ITA Phase II Expansion

Item	Description
192	**Multi-component integrated circuits (MCOs)**: A combination of one or more monolithic, hybrid, or multi-chip integrated circuits with at least one of the following components: silicon-based sensors, actuators, oscillators, resonators or combinations thereof, or components performing the functions of articles classifiable under headings 8532, 8533, 8541, or inductors classifiable under heading 8504. These are integrated into a single body for assembly onto a printed circuit board. Participants defined components, sensors, actuators, resonators, and oscillators in terms of their microelectronic or mechanical structure and their physical functions (e.g., converting signals, generating oscillations).
193	**Light-Emitting Diode (LED) backlight modules**: Lighting sources with one or more LEDs, connectors, and other passive components, used as backlight illumination for LCDs.
194	**Touch-Sensitive Data Input Devices (touch screens)**: Input devices without display capability that detect touch location via resistive, capacitive, acoustic, infrared, or other technology.
195	**Ink cartridges**: Includes cartridges with or without integrated print heads, for insertion into apparatus under HS 844331, 844332, or 844339. Includes toner and solid ink shapes.
196	**Printed matter granting software/data access**: Includes digital access rights to games, apps, services, or online content. Tariff elimination applies only to the physical printed matter, not the content or service regulation.
197	**Self-adhesive circular polishing pads**: Used in the manufacture of semiconductor wafers.
198	**Boxes, cases, crates, and similar articles**: Of plastic, specially shaped or fitted for the packing of semiconductor wafers, masks, or reticles (HS 392310 or 848690).
199	**Vacuum pumps**: Used principally in the manufacture of semiconductors or flat panel displays.
200	**Plasma cleaner machines**: Remove organic contaminants from electron microscopy specimens and holders.
201	**Portable interactive electronic education devices**: Designed primarily for children.

Table 11: List of ITA II Members

ALB	AUS	AUT	BEL
BGR	CAN	CHE	CHN
COL	CRI	CYP	CZE
DEU	DNK	ESP	EST
FIN	FRA	GRC	GTM
HKG	HRV	HUN	IRL
ISL	ISR	ITA	JPN
KOR	LTU	LUX	LVA
MLT	MNE	MRT	MYS
NLD	NOR	NZL	PHL
POL	PRT	ROU	SGP
SVK	SVN	SWE	THA
TWN	USA		