ISET WORKING PAPER SERIES

WP 02-26

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

Lasha Chochua, Irene Iodice

The International School of Economics at Tbilisi State University (ISET) is supported by BP, the Government of Georgia, the Norwegian Ministry of Foreign Affairs, Higher Education Support Program of the Open Society Institute, the Swedish International Development Agency and the World Bank.

International School of Economics at Tbilisi State University
16 Zandukeli Street, Tbilisi 0108, Georgia
www.iset.ge

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

Lasha Chochua † Irene Iodice‡

May 15, 2026

Abstract

Do countries gain more by liberalizing trade together than alone? We study the WTO's 2016 Phase II expansion of the Information Technology Agreement (ITA), which eliminated tariffs on products covering roughly 12% of world goods trade. Using triple-difference structural gravity, importer–product market access rose by 4–6%. Decomposition reveals nearly half reflects general-equilibrium coordination spillovers—exceeding contributions from direct tariff cuts or reduced policy uncertainty. Exploiting variation in coalition size, spillovers turn positive once participants span about two-thirds of world imports—well below the 80% critical-mass benchmark commonly assumed for plurilaterals. Conditional general equilibrium counterfactuals show ITA Phase II reduced members' import price indexes by 1.4 percentage points on average, peaking near 2.0 percentage points by 2019. Joint liberalization yields benefits beyond the sum of individual actions—evidence of the force of many.

*We are grateful to Chad Bown, Maria Pthaskina, Gerald Willmann, Andrew Bernard, Emanuel Ornelas, James Lake, Lars Nilsson, Isabelle Mejean, Fedora Teti, and participants and organizers of the following venues for their valuable feedback: EEA Conference, the Sardinian Empirical Trade Conference, the CESifo Area Conference, the WTO Public Forum, the CITP Community, the Melbourne Empirical Trade Conference, the Kiel-Aarhus Workshop, the ITSG, and the ETSG.

†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: (i) broad multilateral rounds spanning many sectors and countries; (ii) preferential trade agreements among a limited set of countries that grant preferential market access to members; and (iii) issue-specific, opt-in plurilateral deals focused on particular sectors. While the first two have been studied extensively, 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 promising solution to these institutional challenges. Open plurilaterals—agreements among a subset of members whose benefits are extended to all on a most-favored-nation (MFN) basis—avoid the trade diversion associated with preferential treatment.[1] Their key design feature is a critical-mass threshold: commitments take effect when participants account for a sufficiently large share of world trade in covered products. This mechanism preserves non-discrimination while limiting free-riding by non-participants—major markets must join for commitments to take effect—yet still allows a subset of members to move ahead when multilateral consensus proves elusive. The product-specific scope further eases negotiation complexity by reducing cross-issue linkages.

Despite these institutional advantages, a fundamental question remains: do plurilateral agreements generate gains that exceed the sum of their parts? Theory suggests they might, as collective action can create spillover effects that amplify trade beyond what individual countries achieve alone. Yet empirical evidence on these coordination benefits—gains from acting jointly rather than individually—is scarce. Moreover, while practitioners often invoke ”critical mass” benchmarks, there is little theoretical or empirical guidance on where such thresholds should be set. This paper addresses both gaps.

We develop a framework that decomposes the gains from plurilateral agreements into three distinct channels: direct tariff reductions lower trade costs through

[1] Open plurilaterals extend benefits MFN-wide, while closed (exclusive) plurilaterals, such as the Government Procurement Agreement, apply benefits only among signatories.

reductions in applied MFN rates; policy-uncertainty reduction operates by narrowing tariff water—the gap between bound and applied rates; and a coordination channel magnifies trade responses through general-equilibrium feedbacks across global markets. Our main contribution is to isolate and measure this coordination channel. The mechanism is intuitive: when a single country liberalizes on an MFN basis, it improves market access for all exporters to that market; when multiple large economies liberalize simultaneously, exporters gain better access to several major markets at once, creating system-wide competitive pressures that amplify trade responses far beyond individual liberalizations. These spillover effects, channeled through multilateral resistance terms in the gravity equation, grow stronger as coalitions expand—the more countries that coordinate, the more powerful the gains.

We test this theory using the 2016 Phase II expansion of the Information Technology Agreement (ITA), a unique setting for causal analysis. The expansion covered a wide array of high-tech products, including advanced semiconductors, accounting for roughly 12% of world goods trade.[2] Several institutional features make it particularly well-suited for our study. First, the WTO's coordinated consensus on a detailed product list limited scope for unilateral influence. Second, the phased reductions of bound tariffs over four years (2016–2019) followed a predetermined linear schedule, generating temporal variation plausibly exogenous to individual country preferences.

Most importantly, the agreement specified products through two distinct mechanisms that created quasi-experimental variation in coalition size. Attachment A listed established products with clear HS codes, which were liberalized by virtually all 50 ITA participants. By contrast, Attachment B provided textual descriptions for innovative products that lacked standardized international classifications. National customs authorities mapped these descriptions to existing tariff lines through administrative procedures driven by technical coding rules rather than economic or political considerations. This process generated plausibly exogenous, product-level variation in participation: some Attachment B products were liberalized by the full 50-country coalition, while others were adopted by only a subset. We exploit this variation in coalition size to test whether products liberalized by more countries experience disproportionately larger trade increases.

Our empirical strategy leverages this variation in a triple-difference design em-

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

bedded in structural gravity, comparing ITA-liberalized products to non-liberalized controls within the same industries (products identified as candidates by stakeholders but not included). Identification comes from cross-product (ITA vs. non-ITA), cross-country (members vs. non-members), and temporal (pre vs. post) variation, with extensive fixed effects to absorb confounding trends. To address the collinearity of MFN policy with importer–product–time fixed effects in standard gravity, we use a two-stage approach: first estimate the gravity model to recover importer–product–time market-access measures; then decompose these measures to quantify direct tariff, uncertainty, and coordination channels, exploiting the importer–product–time rather than bilateral variation.

The results provide strong evidence for coordination gains. ITA liberalization increased market access by 4–6 percentage points, with coordination effects representing half of this gain. Direct tariff reductions account for 30% of the total effect and uncertainty reduction explains 24%, but the coordination premium—gains emerging from simultaneous action—represents the remaining 46%.

Based on the regression estimates, we calculate the critical-mass threshold and find that coordination benefits turn positive once liberalizing coalitions account for approximately 66% of world imports in covered products—substantially lower than the 80% benchmark commonly assumed by practitioners.

Finally, we quantify ITA effects in terms of welfare (moving beyond market access measures) using a structural-gravity, conditional general equilibrium counterfactual comparing the realized ITA Phase II schedule to a no-reform benchmark. Treating tariff cuts and binding reductions as observed policy shifters, we re-estimate the gravity model with these terms as offsets, holding expenditures and outputs fixed (appropriate for capturing short-run adjustments over our 4-year horizon) so that multilateral resistance adjusts endogenously. Members' import price indexes decline by approximately 1.4 percentage points on average—peaking near 2.0 percentage points by 2019—demonstrating that the gains are general equilibrium effects transmitted through the multilateral-resistance network.

Overall, our findings have important implications for trade policy design. By documenting that coordination yields substantially larger gains than unilateral liberalization—and that these gains grow with coalition size—we provide new quantitative evidence for plurilateral design. Our results suggest that coordination spillovers are positive at participation thresholds lower than those often assumed by policymakers, making coordinated liberalization feasible when broad multilateral progress remains elusive.

1.1 Related Literature

Our paper contributes to three main areas of the literature on international trade.

WTO, MFN Liberalization, and Bindings We quantify the effects of non- discriminatory (MFN) trade policy, building on a large literature that estimates the impact of MFN tariff cuts within structural gravity frameworks. Early studies using simple WTO membership dummies found little evidence of significant trade effects (Rose, 2004), a puzzle later resolved by models that emphasized the non-discriminatory nature of MFN liberalization and the central role of multilateral resistance (e.g., Anderson and Van Wincoop, 2003). More recent reassessments addressing these issues find large positive effects: for instance, Larch et al. (2019) provide empirical evidence of sizable increases in international trade, while Caliendo et al. (2015) show in a structural model with production linkages and firm heterogeneity that most of the global gains stem from multilateral MFN reductions during the Uruguay Round. While Caliendo et al.'s framework reveals important amplification mechanisms, its two-country setup cannot capture the collective action that arises when many countries liberalize simultaneously. We interpret multilateral resistance as a channel for this coordination effect, which we argue is a key mechanism reinforcing the gains from multilateral liberalization. To demonstrate this mechanism empirically, we study a domain-specific MFN setting—the Information Technology Agreement (ITA) Phase II—which provides a unique case to precisely identify the effects of product-specific liberalization, disentangled from the broader cross-issue linkages of WTO negotiating rounds.

Our focus also connects to the literature on how bindings shape policy. Bound rates constrain discretion, shrink tariff "water," and thereby reduce trade policy uncertainty (TPU). Theory models bindings as insurance against beggar-thy-neighbor policy (Bagwell and Staiger, 2011; Beshkar et al., 2015), while empirical work links tighter bindings and credible commitments to higher entry, investment, and trade (Handley and Limao, 2015; Handley and Limão, 2017). We extend this literature by (i) decomposing market-access gains from applied MFN cuts versus uncertainty reduction due to narrowing water, and (ii) isolating a third component—coordination spillovers—arising when multiple countries liberalize simultaneously on an MFN basis.

Plurilateral Agreements, Free Riding, and Critical Mass The theoretical case for open plurilateral agreements is rooted in classic insights about coordination failures. When a subset of countries liberalizes on an MFN basis, non-participants can *free ride* on improved market access without reciprocal concessions. Conversely, preferential agreements create an *exclusion problem* by restricting benefits to coalition members. A broad theoretical literature analyzes these *free riding* and *exclusion* incentives, showing how they can undermine global liberalization (Saggi and Yildiz, 2010; Saggi et al., 2013; Missios et al., 2016; Berens et al., 2021).

The institutional design of open plurilaterals with a "critical mass" threshold directly addresses both problems (Hoekman and Sabel, 2021). Open membership rules solve the exclusion problem, while the critical-mass design tackles free riding by ensuring the coalition internalizes most market-access externalities. However, critical-mass thresholds remain undefined in practice, with legal scholarship often citing an 80% threshold without an economic foundation (Adlung and Mamdouh, 2018). Our paper extends this theoretical literature by providing the first systematic empirical attempt to quantify coordination gains and to calibrate a critical-mass threshold that eliminates coordination failures.

Empirical Evidence on the ITA Empirical work on the ITA has largely focused on Phase I, with early studies like Mann and Liu (2009) and more recent analyses by Gnutzmann-Mkrtchyan and Henn (2018). The latter paper documented a key non-linearity: a complete removal of tariffs leads to greater trade gains than partial reductions, highlighting the value of durable tariff commitments.

We extend this literature by studying the 2016 ITA Phase II expansion, which, has not been analyzed before. Unlike the first phase, which was driven by a few key members, Phase II was the product of a lengthy and more inclusive negotiation. This more representative process may explain why our estimates show significant effects for the average country, while previous work found effects concentrated only among a subset of smaller economies. Furthermore, we exploit the availability of staging matrices to trace the yearly phase-outs of bound tariffs. This provides us with plausibly exogenous variation that allows us to separately identify the effects of both short-term uncertainty reduction and the long-term commitment continuity associated with ITA membership. Finally, we leverage institutional variation between Attachments A and B to obtain quasi-experimental differences in coalition size at the product level. This design enables a direct test of the "force of many" mechanism and delivers the first systematic estimate of coordination gains in an open plurilateral, along with a quantification of the implied critical-mass threshold.

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 is a plurilateral trade agreement administered under the WTO framework. It is "open" in that it extends MFN treatment to all WTO members, regardless of their direct participation. Signatories commit to eliminating bound MFN tariff rates on specified information technology products, thereby establishing zero-tariff regimes for covered goods.

2.1 ITA Phase I: Foundation and Expansion

Launched in 1996 with 29 founding WTO members, the original ITA proved remarkably successful in both membership growth and trade expansion. By 2020, participation had risen to 82 countries, including all major economies in the global electronics industry. Over the same period, global exports of ITA Phase I products nearly quadrupled, reaching $2 trillion in 2020 and accounting for more than 10 percent of world merchandise exports.[3]

2.2 The Path to ITA Phase II: Protracted Negotiations

Building on Phase I's success, negotiators launched expansion talks in 1997, immediately after implementation. These quickly stalled, however, due to disagreements over product coverage, and the initiative remained dormant for more than 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 drafted a preliminary product list covering roughly 357 items at the 6-digit Harmonized System (HS6) level.

Despite initial progress and circulation of a revised list in June 2013, negotiations were suspended the following month due to persistent disagreements over

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

coverage, particularly on LCD panels and machine tools. Compared to the original 29-member agreement, the larger and more diverse membership base had fundamentally altered the negotiating dynamics. After 17 unsuccessful rounds over three years, WTO Director-General Roberto Azevêdo's intervention in 2015 proved decisive.[4] His consensus-building approach focused on products with broad multilateral support, albeit on a narrower scope than initially proposed.

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

2.3 Implementation Framework

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

Standardized staging and absence of safeguards. Unlike many trade agreements that allow broad special and differential treatment, ITA Phase II adopted a uniform staging schedule. All participants committed to eliminating tariffs on covered products in four equal annual cuts—beginning July 2016, followed by reductions each January—culminating in full elimination by January 1, 2019. Limited exceptions applied to about 5% of product–country pairs for three developing countries (Albania, Malaysia, and the Philippines), which were granted seven-year staging periods. Negotiators explicitly rejected proposals for product-specific exceptions or safeguard mechanisms. This "all-or-nothing" approach distinguishes the ITA from other trade agreements that typically permit various flexibility mechanisms.

Product classification and exogenous variation. The final Phase II product list was organized into two attachments. Attachment A contained 191 products with predetermined HS2007 codes, which all 50 signatories agreed to liberalize universally. Attachment B, by contrast, covered 10 technologically advanced products—such as multi-component integrated circuits (MCOs), LED backlight modules, and

[4] A comprehensive discussion of the timeline is available here.

[5] The list of members that joined the expansion is reported in Table 11.

touch-sensitive input devices—that lacked HS2007 codes. This gap reflected the mismatch between five-year HS revision cycles and longer negotiation timelines: these products had entered world markets before receiving dedicated HS subheadings and were therefore absent from the earlier stages of negotiation. Each signatory thus had to map the descriptions to existing national tariff lines, typically at the 8–12-digit level, according to its own customs interpretation.

This mapping exercise generated substantial cross-country heterogeneity in HS6 coverage. For instance, 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), emphasizing principal use; the European Union and China mapped them to HS 5911.90 (Textile products and articles for technical uses), emphasizing material composition; and Canada mapped them to both HS 8486.90 and HS 5911.90, thereby liberalizing both.

2.4 Implications for Empirical Identification

These design features create several advantages for causal identification.

Timing uncertainty: The protracted and interrupted negotiation process (2012–2015) made it difficult for firms or governments to anticipate the final implementation date, reducing concerns about anticipatory behavior.

Product list exogeneity: The Director-General's consensus-driven approach, rather than country-specific bargaining, makes it unlikely that the final list reflects narrow national interests.

Uniform implementation: The absence of special treatment and the adoption of a standardized staging schedule minimize cross-country variation, enabling the identification of common treatment effects.

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

Together, these institutional features provide a quasi-experimental setting for estimating the effects of coordinated liberalization, with Attachment B heterogeneity forming 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 minimize 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. They are evenly distributed between control and treated goods, alleviating identification concerns (see Table 8). 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.

[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.

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 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 summarize these descriptions is reported in Appendix 9. 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 12.5% of world merchandise trade in 2015, totaling almost $2 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.

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

(a) Imports, average inflows (b) Liberalization, average applied tariff

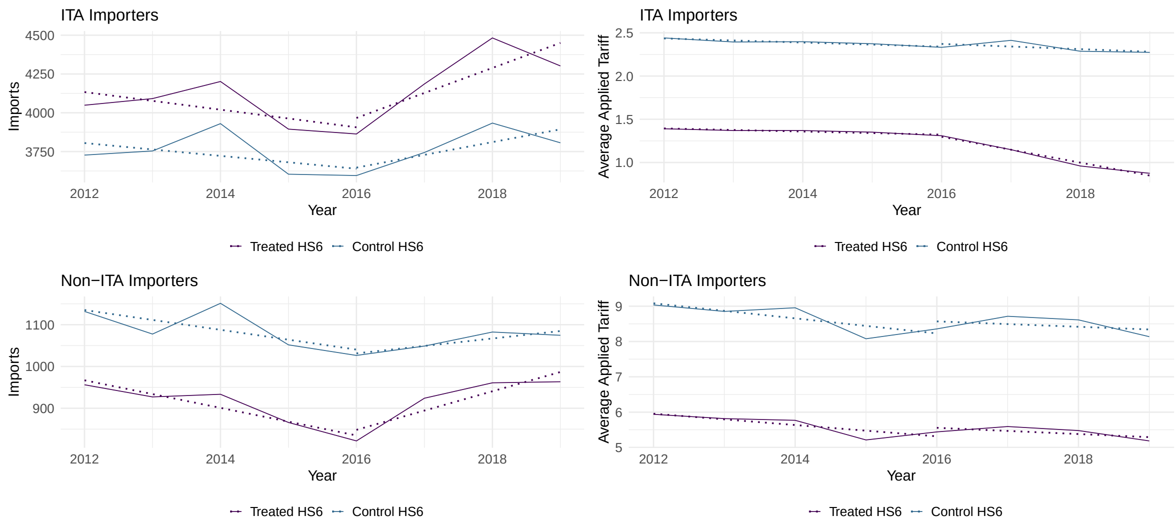

Notes: Imports in thousands of dollars. Applied tariffs are the minimum of applied MFN and preferential rates.

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.

Figure 2: Uncertainty Reduction Over ITA Phase II

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

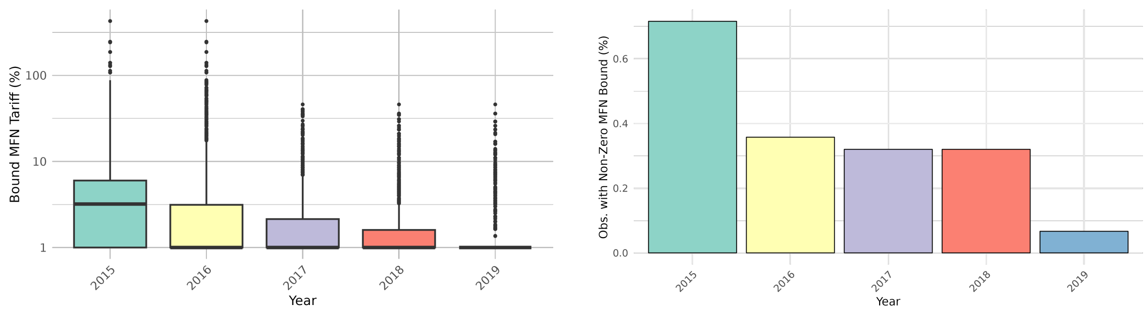

Notes: Includes only HS6 lines without exclusions.

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 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.

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 heterogeneity, 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 formalize how open plurilateral agreements, where countries jointly liberalize on an MFN basis, generate amplification effects through multilateral resistance 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 trade gains through three channels. Coordinated MFN liberalization, as in the case of the ITA, acts on the importer component of trade costs τ_{jkt}, lowering barriers vis-à-vis all trading partners rather than conferring preferential access. Channels 1 and 2 capture these direct effects, while Channel 3 highlights the indirect amplification through multilateral resistance.

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, holding MR terms fixed, 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 Effects via Multilateral Resistance (Unilateral case)

Beyond the direct channels, trade cost reductions generate indirect effects that operate through multilateral resistance (MR) terms. A log-differential of (1) gives

$$d \ln x_{ijkt} \;=\; (1-\sigma)\left[d \ln \tau_{jkt} \;-\; d \ln P_{jkt} \;-\; d \ln \Pi_{ikt}\right]. \tag{9}$$

The first term is the direct trade-creation effect. The latter two are MR adjustments.

Inward MR (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 partly offsets the direct effect by increasing competition in j and making it relatively less attractive. The offset is however *incomplete* because internal trade enters P_{jkt} via τ_{jjkt}, which does not change with an import tariff cut.[8]

Outward MR (Π_{ikt}). If only j liberalizes, the impact on an exporter's outward MR Π_{ikt} is typically small: Π_{ikt} averages access costs over all destinations, so one cheaper market contributes little weight. Hence, under a unilateral cut, the net effect in (9) is positive but attenuated by $d \ln P_{jkt} < 0$.

4.4 Channel 3: Coalition Effects: The Amplification Mechanism

The power of plurilateral agreements lies in coordination. When multiple destinations cut MFN together, exporters' average access improves materially, which in turn shrinks the usual P-offset and amplifies the import response in each liberalizing market.

4.4.1 Intuition

Let $\mathcal{C}_{kt}$ be the set of liberalizing importers for product k at time t. As more $j' \in \mathcal{C}_{kt}$ cut, each exporter faces lower costs in multiple major markets, so Π_{ikt} falls by more. Because P_{jkt} aggregates over suppliers, that system-wide improvement in Π_{ikt} reduces the magnitude of the P-offset.

[8] Internal trade $i = j$ is priced with τ_{jjkt} and is unaffected by MFN changes on imports.

4.4.2 Formal analysis

Starting from (3), log-differentiation with respect to $\{\tau_{j'kt}\}_{j'\in\mathcal{C}_{kt}}$ yields

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

Defining exporter market-access weights

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

gives the compact form

$$d\ln\Pi_{ikt} \;\approx\; \sum_{j'\in\mathcal{C}_{kt}}\omega_{ij'kt}\,d\ln\tau_{j'kt}. \tag{12}$$

Similarly, with supply weights

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

the inward index adjusts according to

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

As the coalition's expenditure share rises, the weighted average in (12) becomes more negative (exporters' global access improves more), which—via (14)—reduces the magnitude of the P-offset.

Interpretation. With home bias, unilateral MFN cuts are partly offset through P_{jkt}. Coordination lowers Π_{ikt} broadly, thereby strengthening the competitive pull from foreign varieties which attenuates that offset and delivers an amplified import response. The amplification grows with the coalition's import size.

What we take to the data. The framework delivers three estimable channels: (i) direct MFN effects, (ii) uncertainty reductions via bindings, and (iii) coordination effects through MR. Because MFN policies vary at the importer–product–time level, identification uses variation across (j, k, t); our empirical strategy recovers $\widehat{\alpha}_{jkt}$ and regresses it on these policy objects and coalition size.

5 Empirical Strategy

Our theoretical framework identifies three channels through which plurilateral agreements generate trade gains: direct tariff reduction, uncertainty reduction through binding commitments, and general equilibrium amplification via multilateral resistance. Testing these channels empirically requires overcoming a fundamental identification 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 estimations cannot identify MFN-based policy effects because of their symmetric nature across trading partners. Recall the bilateral trade cost decomposition in equation (4), where the importer-specific component τ_{jkt} captures MFN policies applied equally to all exporters. In a typical gravity specification,

$$x_{ijkt} = \exp\left(\alpha_{ikt} + \alpha_{jkt} + \alpha_{ijk} + \beta \cdot \ln \tau_{ijkt}\right) \cdot \varepsilon_{ijkt}, \tag{15}$$

the importer–product–time fixed effect α_{jkt} absorbs all variation in τ_{jkt}, including MFN tariff changes. As a result, the coefficient β cannot recover MFN effects for two reasons: (i) for pairs covered by preferential trade agreements, MFN changes are irrelevant because applied rates are already zero; and (ii) for pairs trading under MFN terms, policy changes are identical across all exporters within the same (j, k, t) cell. Including bilateral controls such as PTA membership in $\mathbf{Z}'_{ijt}\boldsymbol{\psi}$ can separate preferential from MFN observations, but once α_{jkt} is included, there is no residual bilateral variation 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 measure of market access:

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

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):

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

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 variable $w_{ijkt}^{\text{appl}} = \ln(1 + t_{ijkt}^{\text{appl}})$ denotes the applied ad valorem tariff rate faced by exporter i in market j for product k at time t. By construction, w_{ijkt}^{appl} equals the MFN rate t_{jkt}^{MFN} when trade occurs under MFN terms, and the preferential rate t_{ijkt}^{pref} under an RTA. The vector $\mathbf{Z}_{ijt}$ collects bilateral controls that vary within the (i, j, t) dimension, such as joint WTO membership or the presence of a preferential trade agreement. Finally, ε_{ijkt} is a multiplicative error term with $\mathbb{E}[\varepsilon_{ijkt}|\cdot] = 1$, ensuring consistency of the PPML estimator.

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 $t > 2015$.[9]

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 measures the effect of applied MFN tariff changes. Lower MFN tariffs reduce trade costs and expand imports, while the gain is partly offset by inward multilateral resistance P_{jkt}. Theory therefore predicts $\beta_1 < 0$.

Channel 2: Uncertainty Reduction via Bindings (β_2). β_2 captures the effect of reductions in tariff overhang when bound rates fall relative to applied rates. Smaller gaps between bound and applied tariffs reduce policy uncertainty and the risk of future reversals. Holding MFN tariffs constant, theory predicts $\beta_2 > 0$.

Channel 3: General Equilibrium Amplification (β_3). β_3 isolates the incremental effect of coordinated liberalization under the ITA, beyond a country's own tariff changes. When many destinations liberalize simultaneously, exporters' outward resistance Π_{ikt} declines, which lowers all importers' inward resistance P_{jkt}. This

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

coordination channel amplifies import responses relative to unilateral MFN cuts. Theory predicts $\beta_3 > 0$.

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

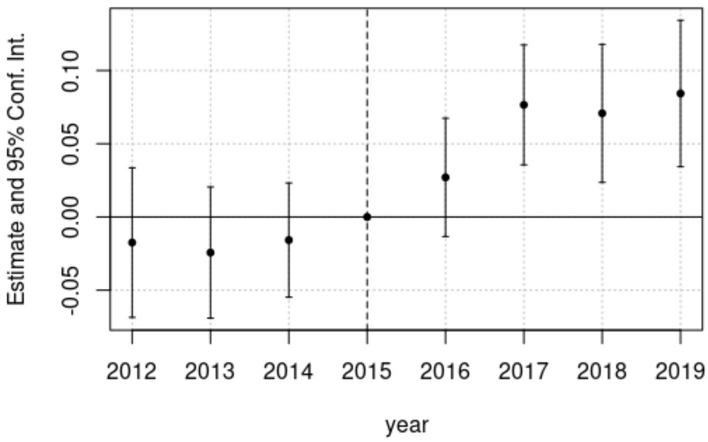

Figure 3 shows the estimated β_n coefficients and 95% confidence intervals. Pre-treatment coefficients are close to zero and statistically insignificant, supporting the parallel-trends assumption. Following the partial implementation in 2016, effects are positive but modest—about 3%—and not statistically 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 Channels

We now turn to the second-stage estimates, which map the recovered importer–product–time effects onto policy variables corresponding to the three channels outlined in Section 4. Table 3 reports how MFN tariff cuts, binding reductions, and coordinated liberalization under the ITA contributed to market-access gains.

Column (1) shows the overall effect of MFN tariff reductions, combining the partial equilibrium trade-creation effect from lowering τ_{jkt} with the general equilibrium adjustment via P_{jkt}. A 10% cut in applied MFN tariffs raises market access by about 3.7%, holding other factors constant.

Column (2) incorporates changes in bound tariffs. The estimate implies that a one–percentage-point cut in the bound rate increases market access by roughly 0.17%, conditional on applied tariffs. For a larger ten–percentage-point cut, the effect is about 1.6%.[10] This supports the prediction that narrowing tariff "water" boosts trade by reducing uncertainty over future policy reversals.[11]

Column (3) adds the ITA coverage indicator. Covered products enjoy nearly 2% higher market access than non-covered products, conditional on tariff changes. This is consistent with the theory: beyond tariff reductions, plurilateral agreements improve market access through the coalition-size amplification mechanism described in Section 4.

Columns (4) and (5) introduce two controls: (i) a dummy for zero MFN tariffs ($\mathbf{1}\{t^{\text{MFN}}_{ikt} = 0\}$) to capture any non-linear effects from zero administrative costs, and (ii) the logarithm of country exports at the product level ($\ln(\text{Export}_{ikt})$) as a proxy for import expenditure.

The first control follows earlier work, though unlike Gnutzmann-Mkrtchyan and Henn (2018) we do not find significant effects. Two factors may explain this difference. First, by not accounting for bound MFN tariffs in staging schedules, their specification may have attributed large short-term uncertainty reduction potential to the zero-tariff dummy. Second, more than two decades after Phase I (1996), the role of reduced administrative costs from less frequent inspections at zero MFN tariffs may have diminished.

The second control addresses potential residual variation in importer–product expenditure not absorbed by fixed effects. This could matter if expenditure dynam-

[10] Computed as $\Delta\widehat{\alpha} = 0.1664 \times \ln(1.10) = 0.01586 \Rightarrow \exp(0.01586) - 1 \approx 0.016$.

[11] β_2 is an elasticity of market access with respect to $\ln(1 + \text{bound cut})$, where the bound cut is expressed in decimals (e.g., 0.10 for a 10 pp cut). Since $\ln(1 + x) \approx x$ for small x, the percentage point change interpretation is a good approximation.

ics were correlated with countries' product-list choices. However, given the inclusive nature of the negotiation process, we do not view this as a serious concern. Consistent with this view, our results remain robust to including country–product exports.

Table 3: Unpacking the Different Channels in the ITA Phase II Expansion

Dependent Variable:	Market Access: $\hat{\alpha}_{jkt}$				
Model:	(1)	(2)	(3)	(4)	(5)
$\ln(1+t_{jkt}^{\text{MFN}})$	-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)
$\mathbb{1}_{\{\text{MFN}=0\}}$					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

Relative importance of channels. To assess the economic significance of each channel, we evaluate the ITA Phase II commitments, which reduced applied MFN and bound tariffs on covered products from their 2015 levels to zero. On average, this corresponds to a 3.8 percentage-point MFN tariff cut and a 7.9 percentage-point binding cut. Applying our Model (3) estimates, the direct tariff channel generates market access gains of about 1.3% and uncertainty reduction contributes roughly 1.1%.[12] The coordination benefit adds a further 2.0% beyond these direct effects. Taken together, the three channels imply a 4.3% improvement in market access for ITA products. Coordination benefits emerge as the dominant driver, accounting for nearly half (46%) of the total predicted effect, followed by direct tariff cuts (29%) and uncertainty reduction (25%).

[12] These values are computed as $0.34 \times \ln(1.038) \approx 0.013$ (1.3%) for direct tariff cuts and $0.14 \times \ln(1.079) \approx 0.011$ (1.1%) for uncertainty reduction.

Table 4: Economic Magnitude of ITA Phase II Liberalization Channels

Channel	Policy Change	Coefficient	Market Access Impact	Share of Total Effect
Direct tariff reduction	3.80 p.p. MFN cut	0.34	+1.27%	29%
Uncertainty reduction	7.85 p.p. binding cut	0.14	+1.06%	25%
Coordination benefits	ITA membership	0.020	+2.00%	46%
Total predicted effect			**+4.33%**	**100%**
Event study estimate			+3% to +8%	

Notes: Baseline year is 2015; endpoints for applied MFN and bound rates are set to zero. Variant B excludes longer transition schedules prior to constructing cuts. Impacts are computed using $\beta_{\text{MFN}} \cdot \ln(1 + \Delta\text{MFN})$, $\beta_{\text{BND}} \cdot \ln(1 + \Delta\text{Bound})$, and the ITA coefficient. Note that given that the changes in decimals are small, the logarithmic terms can be treated as approximately equal to their corresponding percentage-point changes. Coefficients from Table 3, 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

The gravity model highlights why liberalization is more powerful when countries act together rather than alone. What matters for trade flows is not only the trade costs between two countries, but how that cost compares to all other available trading opportunities.

When a single country cuts its MFN tariffs, two forces operate. The direct effect is straightforward: lower trade costs expand imports. But a general equilibrium force pushes the other way: as the country becomes relatively more attractive than still-protected markets, its average price index falls, making it a tougher competitive environment for exporters. This inward multilateral resistance effect partially offsets the initial gains.

Coordinated liberalization weakens this offset. When several large economies reduce tariffs simultaneously, exporters gain access to multiple key markets at once. Their average global access improves—generating a reduction in outward multilateral resistance—so the system-wide adjustment that usually dampens unilateral liberalization becomes much weaker.

Consider semiconductors. If only the United States cuts tariffs, imports rise but the effect is limited because Europe and China remain closed. If all three cut at once, exporters face open doors across the world's largest markets. Profit opportunities expand globally, and competition intensifies across destinations. Each member's market access response is therefore larger in a coalition than it would be in isolation.

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:

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

where $\textit{CoalitionSize}_{kt}$ is measured in two ways: (i) the number of other countries implementing tariff cuts on product k in year t, 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 liberalizing countries varies across products in this group but is identical across importers for Attachment A products. 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 liberalizing countries.

To prevent the coalition–size coefficient β_4 from picking up structural differences between product groups, some specifications include an ITA × 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 5 reports the results. Columns (1) and (3) include the coalition-size interaction, using alternative measures of coalition size. In both cases, the interaction term is positive and statistically significant. With coalition size measured as the number of liberalizing countries, the coefficient is about 0.014 ($p < 0.1$). When measured as the coalition's share of world imports, the coefficient is much larger—about 0.22 ($p < 0.01$). These estimates imply that the coordination benefits of ITA membership increase strongly with the scale of participation.

Columns (2) and (4) add an ITA × AttachmentB control to absorb systematic differences between Attachments A and B. This addition leaves the main results intact: the coalition-size interaction remains positive and significant, with magnitudes similar to the baseline. The Appendix B dummy itself is small and statistically insignificant, suggesting that classification differences between Attachments A and B are not driving the observed effects.

Across all specifications, the standalone ITA coefficient is negative once the interaction is included. This pattern is consistent with the theory: for small coalitions, the offsetting inward multilateral resistance effects can dominate, making the average ITA effect negligible or even negative. As coalition size grows, however, the positive interaction term outweighs this offset, turning the net effect positive. In other words, meaningful gains from plurilateral liberalization emerge only once a critical mass of countries participates.

Table 5: 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)
ITA_{jkt} × 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. Column (4) of Table 5 implies a coalition-size threshold S^* at which the net ITA effect turns positive.[13] The estimate yields $S^* \approx 0.66$, meaning that coordinated liberalization delivers net gains once about two-thirds of world imports in a product are covered. Above this point, the positive interaction effect dominates the negative baseline coefficient. In practice, ITA Phase II coalitions typically exceeded 80 percent of world imports, comfortably above this threshold.

This S^* should be interpreted as an empirical coordination threshold: the coalition share required for multilateral-resistance forces to amplify rather than offset direct tariff cuts. It is not a legal or welfare benchmark, but it provides a quantitative anchor for the long-standing policy notion of "critical mass" in plurilateral agreements.

8 Alternative Channels and Robustness

This section examines alternative explanations for our coordination effects and tests the robustness of our main findings. We first investigate whether our results operate through channels other than multilateral resistance, including heterogeneous trade elasticities and vertical supply chain linkages. We then conduct extensive robustness checks to ensure our findings are not driven by confounding shocks such as the US-China trade war, endogenous trade preferences from active players, or variations in liberalization scope and implementation across products, and standard error clustering.

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 6, 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

[13] Formally, we solve $\hat{\beta}_3 + \hat{\beta}_4 \cdot \ln(1 + S^*) = 0$ using the estimates in Column (4). With $\hat{\beta}_3 = -0.129$ and $\hat{\beta}_4 = 0.254$, this yields $S \approx 0.66$.

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) (Caliendo and Parro, 2015). 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 6, columns 3-4, interacts the ITA indicator with product-level upstreamness. Column 3 indicates that more upstream products display somewhat larger ITA gains, consistent with input–cost spillovers. Yet once coalition heterogeneity is controlled for (column 4), the interaction loses significance, while the core coordination effect is positive and highly significant. This pattern suggests that vertical linkages may play a role but are not the primary driver of the results, likely because ITA products are relatively homogeneous high-tech inputs with similar positions in global supply chains.

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 (Bown, 2020). To ensure our results are not driven by these bilat-

Table 6: 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
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.

eral tensions, we re-estimate our baseline specification while excluding both the US and China from our sample. Table 7, 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 7), 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 7: 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	
$\ln(1+t_{jkt}^{\text{MFN}})$	-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)
$\ln(1+\Delta\lvert t_{jkt}^{\text{BND}}\rvert)\times \text{Treat}_{jkt}$	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. $\ln(1+t_{jkt}^{\text{MFN}})$ is the log MFN ad valorem tariff. $\ln(1+\Delta\lvert t_{jkt}^{\text{BND}}\rvert)\times\text{Treat}_{jkt}$ 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 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 (Table 12 in the Appendix) 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 product–time and various multiway clustering schemes, including product–country and importer–time. Our results (Table 13 in the Appendix) remain consistent and statistically significant across these alternative specifications, confirming the robustness of our main findings.

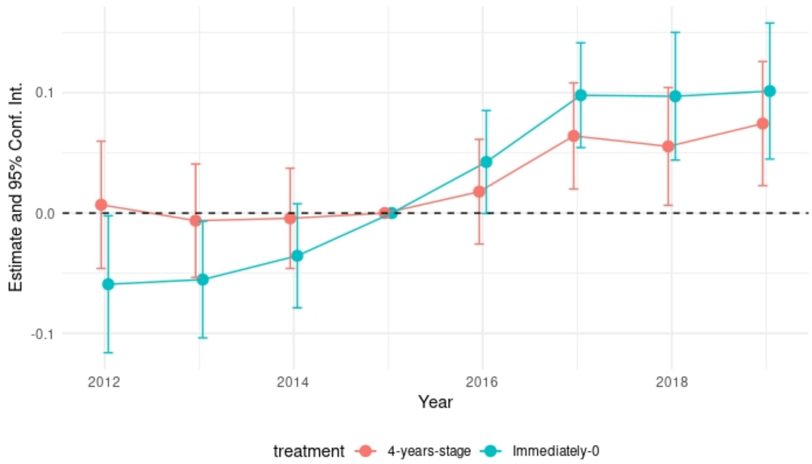

Figure 4: Market Access by Treatment type, 4-years-stage vs Immediately Zero

8.3 Simulating Counterfactual Market Access

We now quantify general-equilibrium gains by simulating market aggregates under a baseline with ITA Phase II schedules versus a counterfactual freezing policies at 2015 levels. The exercise uses conditional GE: we hold expenditures and production fundamentals fixed (appropriate for a 4-year, short-run horizon) and let multilateral resistances adjust endogenously. This delivers model-implied inward MR terms, which serve as CES import price indices to measure the welfare implications of ITA Phase II relative to no agreement.

Objects and assumptions. We work within a structural gravity framework with CES demand (variety-level elasticity $\sigma = 8.74$, taken from the literature). Our counterfactual exercises use *conditional GE*: total expenditure E_{jkt} and output Y_{ikt} are held fixed across scenarios to isolate endogenous multilateral-resistance adjustments from policy changes, appropriate for our 4-year horizon. Bilateral preference wedges (e.g., PTA preferences) remain fixed throughout. The policy reform affects importer–product trade costs through applied MFN tariffs and tariff water (the gap between bound and applied MFN rates). These tariff elements map into iceberg trade costs using a parameter $\gamma = 1/3$, estimated from our reduced-form coefficients $\hat{\beta}_1/\hat{\beta}_2$, ensuring applied and bound MFN changes feed consistently into model-implied trade frictions.

Policy environments. We consider two scenarios, $s \in \{\text{baseline (bsl)}, \text{counterfactual (cnt)}\}$, in which importer–product–time trade costs take the form

$$\ln \tau^{s}_{ijkt} = \ln\big(1 + t^{\text{MFN},s}_{jkt}\big) + \ln\big(1 + \gamma\, w^{s}_{jkt}\big) + \omega^{\text{pref}}_{ijkt}.$$

In the baseline, $(t^{\text{MFN,bsl}}_{jkt}, w^{\text{bsl}}_{jkt})$ follow the ITA Phase II implementation schedule over 2016–2019. In the counterfactual, $(t^{\text{MFN,cnt}}_{jkt}, w^{\text{cnt}}_{jkt})$ are held fixed at their 2015 levels for all $t \geq 2016$. The bilateral preference wedge $\omega^{\text{pref}}_{ijkt}$ is kept constant across scenarios.

Stage-1 gravity with a plug-in policy term. For each scenario s, we estimate the same PPML gravity specification,

$$x_{ijkt} = \exp\Big(\alpha_{ikt} + \alpha_{jkt} + \alpha_{ijk} + (1-\sigma)\ \ln \tau^{s}_{ijkt}\Big) \cdot \varepsilon_{ijkt},$$

where we *plug in* the known policy term $(1-\sigma)\ln \tau^{s}_{ijkt}$ rather than estimating it.[14] The recovered importer–product–time effects $\widehat{\alpha}^{s}_{jkt}$ then absorb the full equilibrium adjustment to the policy shock, combining direct price effects and multilateral-resistance spillovers.

Recovering the MR response. Using $\alpha_{jkt} = \ln E_{jkt} + (1-\sigma)[\ln \tau_{jkt} - \ln P_{jkt}]$ and keeping E_{jkt} fixed across scenarios, the implied change in inward MR is

$$\Delta \ln P_{jkt} \ \equiv\ \ln P^{\text{cnt}}_{jkt} - \ln P^{\text{bsl}}_{jkt} \ =\ -\frac{\widehat{\alpha}^{\text{cnt}}_{jkt} - \widehat{\alpha}^{\text{bsl}}_{jkt}}{1-\sigma}. \tag{21}$$

We demean $\widehat{\alpha}^{s}_{jkt}$ across products within each importer-year before differencing to purge normalization constants.

[14] In GLM/PPML terminology this term is an *offset*—a known component of the linear predictor with coefficient fixed to one.

Aggregating to market indexes and welfare. For importer j in year t, we aggregate product-level MR changes using initial (2015) expenditure shares s°_{jk} to avoid endogeneity from policy-induced expenditure reallocation. Assuming a Cobb–Douglas aggregator across products, the log change in the consumer price index is

$$\Delta \ln P_{jt} \;=\; \sum_{k} s^{\circ}_{jk}\, \Delta \ln P_{jkt}.$$

Under homothetic preferences and fixed nominal expenditure (our conditional GE), the welfare change is $\Delta \ln W_{jt} \;=\; -\,\Delta \ln P_{jt}$.

Results: equilibrium price effects. Comparing the ITA path to the 2015-frozen counterfactual, we find sizable MR declines for members on covered products: the average $\Delta \ln P_{jkt}$ for signatories is -1.43pp, versus essentially zero for non-members (-0.04pp). The effect deepens with staged liberalization: -0.55pp (2016), -1.41pp (2017), -1.70pp (2018), and -2.06pp (2019). These differences are the model-implied equilibrium adjustments consistent with the tariff schedules in the baseline and counterfactual scenarios. They therefore reflect the conditional general-equilibrium response to the ITA reform—including both the direct impact of tariff changes and the induced multilateral-resistance adjustments.

9 Conclusion

This paper has shown that when countries liberalize trade together, the gains are greater than the sum of their individual actions. By examining the Information Technology Agreement, we developed a framework that isolates three distinct channels through which coordinated liberalization works: direct tariff cuts, a reduction in policy uncertainty, and a general-equilibrium amplification effect—what we call the coordination premium.

Our empirical analysis confirms the significance of all three channels. We found that the ITA Phase II expansion increased market access for covered products by 4–6%. The gains were not evenly distributed: direct tariff cuts accounted for about 30%, lower policy uncertainty explained roughly 25%, and the remaining 45% came from the measurable coordination premium. This premium arises because simultaneous MFN liberalization by a group of large importers shifts global market conditions, amplifying trade responses across the entire system. Crucially, this amplification grows with the size of the coalition. For instance, products liberalized by the full group of 50 participants saw roughly twice the trade response compared to those liberalized by only half that many, even when holding tariff changes constant. Our findings also suggest that the conventional "critical mass" rule of thumb—often cited as 80%—is high; our estimate shows that coordination gains become decisively positive once the liberalizing coalition accounts for around two-thirds of world imports in the covered products.

To quantify the economic magnitude of these effects, we conducted a structural counterfactual exercise that isolates the coordination channel's impact on welfare. Treating the observed tariff cuts and binding reductions as known policy changes, we re-estimated the gravity system to allow multilateral resistance terms to adjust endogenously while holding expenditures and production fixed. This exercise reveals that the coordination channel alone lowers members' import price indices by approximately 1.4 percentage points on average, with effects peaking near 2.0 percentage points by 2019.

These findings carry important policy implications. Open plurilaterals are a viable path forward when multilateral consensus is stalled, as they preserve MFN principles while minimizing trade diversion. To be most effective, these agreements should target a sufficiently large coalition to capture the full coordination premium, secure credible bindings to reduce uncertainty, and use a product scope that minimizes complexity. While our study focuses on a high-tech goods setting, future research could explore the application of this framework to services, non-tariff measures, and different product sectors. It would also be valuable to connect these market-access gains to overall welfare, and to investigate how firm-level heterogeneity and supply-chain dynamics interact with these policy changes.

Overall, plurilateral agreements aren't just about adding up individual actions, they amplify them. Accounting for this coordination premium is crucial for designing agreements that sustain MFN-based liberalization and deliver durable gains to the global trading system.

References

ADLUNG, R. AND H. MAMDOUH (2018): "Plurilateral trade agreements: An escape route for the WTO?" *Journal of World Trade*, 52.

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.

BAGWELL, K. AND R. W. STAIGER (2011): "What do trade negotiators negotiate about? Empirical evidence from the World Trade Organization," *American Economic Review*, 101, 1238–1273.

BERENS, S., L. CHOCHUA, AND G. WILLMANN (2021): "The Farsighted Stability of Global Trade Policy Arrangements," CESifo Working Paper 9213, CESifo, Center for Economic Studies and ifo Institute, Munich.

BESHKAR, M., E. W. BOND, AND Y. RHO (2015): "Tariff binding and overhang: theory and evidence," *Journal of international Economics*, 97, 1–13.

BOWN, C. (2020): "How the United States marched the semiconductor industry into its trade war with China," *East Asian Economic Review (EAER)*, 24, 349–388.

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.

CALIENDO, L. AND F. PARRO (2015): "Estimates of the Trade and Welfare Effects of NAFTA," *The Review of Economic Studies*, 82, 1–44.

CONTE, M., P. COTTERLAZ, T. MAYER, ET AL. (2022): *The CEPII gravity database*, CEPII.

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.

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.

HOEKMAN, B. AND C. SABEL (2021): "Plurilateral Cooperation as an Alternative to Trade Agreements: Innovating One Domain at a Time," *Global Policy*, 12, 49–60.

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.

MISSIOS, P., K. SAGGI, AND H. M. YILDIZ (2016): "External trade diversion, exclusion incentives and the nature of preferential trade agreements," *Journal of International Economics*, 99, 105–119.

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

SAGGI, K., A. WOODLAND, AND H. M. YILDIZ (2013): "On the relationship between preferential and multilateral trade liberalization: the case of customs unions," *American Economic Journal: Microeconomics*, 5, 63–99.

Saggi, K. and H. M. Yildiz (2010): "Bilateralism, multilateralism, and the quest for global free trade," *Journal of International Economics*, 81, 26–37.

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

A Appendix

Table 8: Share of Missing Applied Tariffs (%) by ITA Status $\times$ 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 9: 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 10: Share of Imports Covered by ITA (2012–2019)

Year	Share of Imports (%)
2012	82.03
2013	82.21
2014	82.53
2015	82.62
2016	83.00
2017	82.29
2018	82.87
2019	82.22

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		

Table 12: Robustness: Partial Liberalization Adjustments

Dependent Variable:	Market Access: $\hat{\alpha}_{jkt}$	
Model:	(1)	(2)
$\ln(1 + t^{\text{MFN}}_{jkt})$	-0.3390***	-0.3384***
	(0.0974)	(0.1070)
$\ln(1 + \lvert\Delta\bar{t}_{jkt}\rvert)$	0.1374***	0.1284*
	(0.0521)	(0.0696)
ITA_{jkt}	0.0179**	-0.1283***
	(0.0081)	(0.0467)
Partial $\text{Liberalization}_{jkt}$	0.0048	0.0199
	(0.0110)	(0.0135)
$\text{ITA}_{jkt} \times \ln(1 + \text{Share}_{2015})$		0.2525***
		(0.0786)
Fixed effects:		
Importer–Time FE	Yes	Yes
Importer–Product FE	Yes	Yes
Product–Time FE	Yes	Yes
Observations	438,951	438,951
R^2	0.94363	0.94364
Within R^2	6.19×10^{-5}	0.00010

Notes: The dependent variable is the importer–product–time fixed effect from the gravity model (Equation 17). Partial $\text{Liberalization}_{jkt}$ is a dummy for HS6 products where some subheadings were excluded from commitments. ITA_{jkt} indicates ITA Phase II coverage; Share_{2015} is the pre-2015 leave-one-out world import share. Standard errors clustered at the treatment–importer–time level in parentheses. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 13: Robustness: Alternative Standard-Error Clustering

Dependent Variable:	Market Access: $\hat{\alpha}_{jkt}$			
Model:	(1)	(2)	(3)	(4)
Clustering:	*jk* (Importer–Product)		*jk* + *jt* (Importer–Product & Importer–Time)	
$\ln(1 + t^{\text{MFN}}_{jkt})$	-0.3476***	-0.3434***	-0.3363***	-0.3324***
	(0.0986)	(0.0985)	(0.0979)	(0.0979)
$\ln(1 + \lvert\Delta\bar{t}_{jkt}\rvert)$	0.1422***	0.1360**	0.1394***	0.1328**
	(0.0525)	(0.0556)	(0.0531)	(0.0560)
ITA_{jkt}	0.0198***	-0.1034***	0.0241***	-0.0948**
	(0.0074)	(0.0388)	(0.0075)	(0.0409)
$\text{ITA}_{jkt} \times \ln(1 + \text{Share}_{2015})$		0.2192***		0.2116***
		(0.0684)		(0.0722)
Fixed effects:				
Importer–Time FE	Yes	Yes	Yes	Yes
Importer–Product FE	Yes	Yes	Yes	Yes
Product–Time FE	Yes	Yes	Yes	Yes
Observations	433,850	433,850	422,267	422,267
R^2	0.94153	0.94153	0.93856	0.93856
Within R^2	6.43×10^{-5}	9.61×10^{-5}	6.71×10^{-5}	9.48×10^{-5}

Notes: Same sample throughout: longer transition schedules excluded (longer_exp_window $\neq 1$). $\ln(1 + t^{\text{MFN}}_{jkt})$ is the log applied MFN rate; $\ln(1 + \lvert\Delta\bar{t}_{jkt}\rvert)$ captures binding (TPU) reductions. ITA_{jkt} indicates ITA Phase II coverage; Share_{2015} is the pre-implementation (2015) leave-one-out world import share for product *k*. Columns (1)–(2) report one-way clustering at the importer–product level (*jk*); Columns (3)–(4) report two-way clustering at *jk* and importer–time (*jt*). Standard errors in parentheses. Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 14: First-Stage PPML Estimation

Dependent Variable:	Bilateral Trade Flows x_{ijkt}	
Model:	PPML with High-Dimensional FE	
	(1)	(2)
τ_{ijkt}	-0.1051	-7.1932***
	(0.2075)	(0.4635)
FTA Membership (FTA_{ijt})	-0.0123	0.3847***
	(0.0227)	(0.0217)
Both WTO Members (WTO_{ijt})	-0.0784	1.5920***
	(0.1811)	(0.3624)
Distance ($\ln(\text{dist}_{ij})$)		-0.6694***
		(0.0108)
Common Colony (comcol_{ij})		0.2162**
		(0.1035)
Common Language (comlang_{ij})		0.1169***
		(0.0290)
Constant	19.18306***	22.4533***
	(0.1811)	(0.3737)
Fixed Effects:		
Importer–Product–Time (jkt)	Yes	Yes
Exporter–Product–Time (ikt)	Yes	Yes
Importer–Exporter–Product (ijk)	Yes	No
Observations	12,308,023	12,842,494
Clusters (ijk)	2,219,394	2,992,346
Pseudo R^2	0.9859	0.9177

Notes: PPML estimates of the first-stage gravity model. Robust standard errors in parentheses. Column (1): standard errors clustered two-way at the importer–exporter–product (ijk) and importer–time (it) levels. Column (2): standard errors clustered at the importer–exporter–product (ijk) level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.