

Iodice, Irene

Working Paper

Transparency as trade policy: Value of advance notifications in technical regulations

CESifo Working Paper, No. 12015

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Iodice, Irene (2025) : Transparency as trade policy: Value of advance notifications in technical regulations, CESifo Working Paper, No. 12015, Munich Society for the Promotion of Economic Research - CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/338315>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CES ifo

**12015
2025**

Original Version: July 2025
This Version: Oktober 2025

Working Papers

Transparency as Trade Policy: The Value of Advance Notifications in Technical Regulations

Irene Iodice

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
<https://www.cesifo.org>

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge

- from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
- from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
- from the RePEc website: <https://ideas.repec.org/s/ces/ceswps.html>

Transparency as Trade Policy: The Value of Advance Notifications in Technical Regulations

Irene Iodice

Bielefeld University, CESifo*

This version - October 29, 2025[†]

Abstract

Lack of advance information about regulatory changes in destination markets increases uncertainty over compliance costs. This paper quantifies the value of advance notification using newly compiled WTO Technical Barriers to Trade (TBT) timelines matched to detailed French firm–product–destination export data for 1995–2007. When regulations are notified before enforcement, the export-participation loss from TBTs falls by more than half (from 7 to 3 percentage points), driven by fewer temporary exits and greater entry among smaller firms. Exploiting cases of delayed notification, I show that exports rebound immediately once information becomes available, indicating that uncertainty—rather than adjustment time—is the key mechanism. A heterogeneous-firm real-options model explains this pattern: by narrowing the dispersion of possible regulatory outcomes, advance notice reduces the option value of waiting. Calibrated tariff equivalents imply that timely notification can be worth up to 28 percentage points. (JEL F13, F14, D84)

*Contact information: irene.iodice@uni-bielefeld.de.

[†]I thank Lionel Fontagné, Lisandra Flach, Mathieu Parenti, Angelo Secchi, Ariell Reshef, Gerald Willmann, Luca Macedoni, Anne-Célia Disdier, Maria Bas, Paola Conconi, Gianluca Orefice, Thierry Mayer, Alessia Lo Turco, Marcus Biermann, and Peter Egger for their valuable feedback. I also thank participants at various conferences for their insightful comments. This research was supported by a public grant overseen by the French National Research Agency (ANR) as part of the “Investissements d’avenir” program (reference: ANR-10-EQPX-17, Centre d’accès sécurisé aux données).

1 Introduction

Technical regulations—covering product standards, testing, and labeling—represent a significant source of friction in international trade. Firms frequently cite these measures as substantial barriers to market access, alongside traditional tariffs (European Commission, 2019; International Trade Centre, 2016). Exporters must tailor products and production processes to meet divergent requirements across destination markets. Interestingly, firms' complaints focus less on the substance of regulations than on the procedures surrounding them: exporters primarily report difficulties obtaining timely, reliable information about new or changing requirements rather than objections to the technical content itself. While an extensive literature documents how and why countries deploy technical measures for protectionist purposes (Lee and Swagel, 1997; Beverelli et al., 2014; Orefice, 2017; Herghelegiu, 2018; Grossman et al., 2021; Mei, 2024), far less is known about whether greater procedural transparency can mitigate these frictions. This gap is particularly striking because, unlike publicly accessible tariff schedules, information on non-tariff measures often remains scattered and non-systematic (UNCTAD, 2015), creating substantial policy uncertainty for foreign firms navigating an opaque regulatory landscape.

This paper addresses this gap by estimating the economic value of timely information about new product rules. I study the transparency procedures in the World Trade Organization's Agreement on Technical Barriers to Trade (TBTs). Under the TBT Agreement, governments are expected to notify proposed technical regulations before enforcement, publishing official documents that detail requirements with sufficient lead time. When this process functions as intended, firms learn the "rulebook" before regulations take effect, allowing them to assess compliance requirements and make informed entry decisions.

I provide the first causal estimates of the economic value of timely notification. I combine all TBTs raised as Specific Trade Concerns at the WTO from 1995 to 2011

with highly disaggregated French firm–product–destination customs data covering 1995–2007. Specific Trade Concerns (STCs) are formal complaints filed when WTO members believe a trading partner’s technical regulation creates unnecessary obstacles to trade—they represent contentious and potentially trade-restrictive TBTs. To obtain precise identification, I employ textual analysis on WTO documents to reconstruct for each product–country pair, whether a measure was notified and, if so, the exact dates it was notified and entered into force. This new dataset reveals substantial variation in compliance: about 45% of challenged TBTs were implemented without any prior notification.

Motivated by this pattern, I develop a heterogeneous-firm real-options framework in the spirit of [Dixit et al. \(1994\)](#) and [Handley and Limão \(2017\)](#). Absent advance notice, firms face uncertainty about the regulatory environment—modelled as a mean-preserving spread in the compliance-cost multiplier—which raises the option value of waiting to invest in market access. Notification reveals the requirement, narrowing dispersion and thereby lowering the real-option value of waiting. The model predicts a productivity cutoff that is weakly increasing in the dispersion of regulatory outcomes: unchanged when entry would be profitable under any plausible compliance cost, and strictly increasing when uncertainty makes profitability contingent on drawing a less restrictive requirement. Four testable implications follow: (i) effects concentrate on the extensive margin; (ii) notification reduces temporary exits and accelerates re-entry; (iii) gains are largest for small and medium sized, less-productive firms; and (iv) participation rebounds when information arrives, even if notice follows enforcement—ruling out a pure adjustment-time story.

I test these predictions using a difference-in-differences design at the firm-product-destination-time level, exploiting variation in TBT implementation across countries and time. Identification relies on two comparisons: (i) regulated products versus unregulated products within the same HS 2-digit sector and destination-time cell; and (ii) within regulated products, differences between pre-notified and non-pre-notified

TBTs via interaction terms. The specification includes firm-product-destination fixed effects and HS2×destination×time fixed effects, absorbing time-invariant heterogeneity at the firm-product-destination level and common shocks at the HS2-destination-time level. The identifying assumption is parallel trends: absent the TBT, treated and comparison products within an HS2-destination cell would have evolved similarly in firms' trade margins.

The central finding is that timely notification carries substantial economic value. When firms have access to official TBT documentation before enforcement, the disruption to export activity is significantly reduced. Specifically, prior notification mitigates the decline in export participation by about 4 percentage points—more than halving the baseline participation loss. The effect is especially pronounced for smaller firms and is driven largely by a reduction in temporary exits (6 percentage points) and an increase in entry (5 percentage points).

To isolate uncertainty from adjustment time, I exploit cases where TBTs were notified only after enforcement. If adjustment time were the constraint, export participation would remain depressed even after delayed notification. Instead, participation recovers soon after firms receive information, supporting the view that notification primarily reduces uncertainty rather than merely providing lead time.

I confirm the robustness of these findings by verifying the parallel pre-trends assumption and testing against alternative channels and estimation strategies, including the substitutability between tariffs and non-tariff barriers and the use of Poisson pseudo-maximum likelihood (PPML).

To quantify these effects, I exploit the fact that the compliance-cost multiplier enters multiplicatively like an iceberg trade cost, so reductions in dispersion admit a tariff-equivalent interpretation. The implied magnitude is large: receiving timely notification is economically equivalent to avoiding a tariff increase of up to 28 percentage points, based on calibration of the model. This highlights the critical importance of the

notification process—not only as a procedural obligation but as a meaningful policy tool for reducing trade-related uncertainty.

Related Literature. This study contributes to several strands of research.

First, it advances the literature on how non-tariff measures—particularly Technical Barriers to Trade (TBTs)—shape firm-level export behavior. Following [Orefice \(2017\)](#), [Herghelegiu \(2018\)](#), [Fontagné and Orefice \(2018\)](#), and [Singh and Chanda \(2021\)](#), I use WTO Specific Trade Concerns (STCs) to identify trade-restrictive TBTs. The closest study is [Fontagné and Orefice \(2018\)](#), who find that contested TBTs reduced French export participation by 5 percentage points. I extend their work by examining whether the trade impact of a TBT depends on advance notification, rather than STC alone. I make three contributions: methodologically, I separate notified from unnotified TBTs to causally estimate the effect of timely notification; conceptually, I develop a real-options framework highlighting policy uncertainty as the mechanism; and quantitatively, I calculate transparency’s economic value in tariff-equivalent terms. Building on literature documenting how TBTs raise entry costs ([Bao and Qiu, 2012](#); [Fontagné and Orefice, 2018](#); [Beghin et al., 2015](#)), stem from regulatory heterogeneity ([Grossman et al., 2021](#); [Iodice and Reverdy, 2025](#); [Mei and Xu, 2024](#)), or enhance trade via harmonization ([Schmidt and Steingress, 2022](#)), I show that implementation procedures are critical yet overlooked determinants of these outcomes.

Second, this paper contributes to the broader literature on transparency in trade policy. While earlier work often relied on qualitative indices or treaty provisions ([Francois, 2001](#); [Helble et al., 2009](#); [Lejárraga and Shepherd, 2013](#)), and more recent studies have proxied transparency using notification propensity ([Ing et al., 2018](#)), this paper’s methodology offers a more direct measure. By exploiting the exact timing of each notification relative to its enforcement, I provide a granular estimate of the value of predictability itself—an aspect emphasized by scholars ([Helble et al., 2009](#)) but rarely measured for non-tariff measures.

Third, this paper extends the literature on trade policy uncertainty (TPU) to non-tariff measures. Seminal real-options models (Bernanke, 1983; Dixit et al., 1994) and recent applications show that uncertainty in trade policy discourages investment when market entry entails irreversible costs (Handley and Limão, 2015, 2017; Carballo et al., 2018; Alessandria et al., 2019; Liu and Ma, 2020; Coelli, 2022). I adapt this framework to NTMs, which impose significant, destination-specific sunk compliance costs (Grossman et al., 2021; Macedoni and Weinberger, 2022). The key innovation is to treat delayed or missing notification as a hidden source of TPU. In the model, this manifests as a mean-preserving spread in the distribution of compliance costs. Conceptually, opacity increases perceived policy variance and raises the option value of waiting—deterring both entry and persistence in export markets.

The remainder of the paper is structured as follows. Section 2 provides institutional background on WTO transparency provisions. Section 3 describes the dataset construction. Section 4 develops a real-options framework for exporting under regulatory uncertainty. Section 5 outlines the empirical strategy and reports the baseline results and robustness. Sections 6 and 7 investigate the underlying mechanisms and heterogeneous effects across firms. Section 8 quantifies the economic value of notification in tariff-equivalent terms, and Section 9 concludes.

2 Institutional Framework

The WTO Agreement on TBTs allows members to adopt technical regulations to pursue legitimate policy goals—such as protecting health, safety, the environment, or national security—even when these regulations depart from international standards. However, this right is subject to disciplines aimed at preventing protectionism. Measures must avoid arbitrary or unjustifiable discrimination and should not be more trade-restrictive than necessary to achieve their objectives.¹

¹The TBT Agreement states: “Members shall ensure that technical regulations are not prepared, adopted or applied with a view to or with the effect of creating unnecessary obstacles to international

A key discipline is procedural transparency. To ensure predictability, the TBT Agreement requires members to notify proposed technical regulations in advance. This is done through a formal “Notification”—a document outlining the proposed measure—ideally allowing a 60-day comment period for trade partners to review and prepare.²

If a member believes these principles—including transparency—have been breached, it may raise a Specific Trade Concern (STC) in the WTO TBT Committee. Analysis of STCs reveals a striking pattern: while the stated objectives of challenged TBTs are often legitimate (e.g., health, safety, environment; see Figure 1a), concerns are frequently raised over procedural issues, particularly transparency and access to information (Figure 1b).³ This mirrors exporters’ concerns documented in UNCTAD (2019) and EU Commission (2019).⁴

For example, Mexico’s 1998 ban on chlorofluorocarbons (CFCs) in refrigerators and air conditioners aligned with international environmental commitments,⁵ but was enacted as an emergency measure on September 22nd and notified only post-implementation, on October 12th.⁶ The US raised an STC shortly after, questioning the emergency rationale and citing exporters’ uncertainty over compliance procedures, such as testing and certification costs.⁷

This case illustrates how deviations from notification procedures can generate trade frictions. Motivated by the prevalence of transparency-related STCs and such examples, this paper investigates the economic value of timely notification.

trade.” (Art. 2.2)

²A notification is “a transparency obligation requiring member governments to report trade measures to the relevant WTO body if the measures might have an effect on other Members” (World Trade Organization, 2017).

³Following the ITC survey (2016), I group “missing information” and “unreasonable time” as procedural implementation flaws.

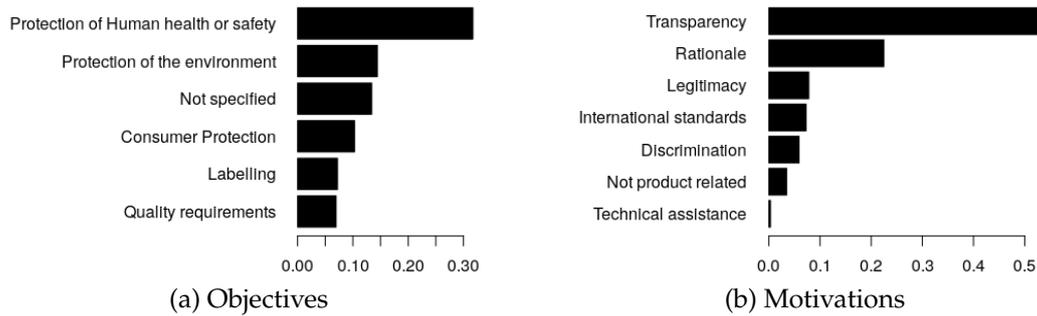
⁴Interviews with trade representatives in Holzer (2019) similarly distinguish between concerns over a measure’s implementation process and its substantive content.

⁵The Montreal Protocol (1989, strengthened in the 1990s) aimed for a global CFC phase-out.

⁶Notification identifier: G/TBT/Notif.98.485.

⁷The US delegate stated: “She questioned the nature of the emergency... US exporters were uncertain about how to comply with the regulation.” (G/TBT/M/14, par. 35)

Figure 1: Objectives and Motivations Behind WTO TBT-Specific Trade Concerns



Notes: The STC database records 598 TBT-objective and 478 TBT-motivation occurrences. I grouped "unreasonable time" and "missing information" under Transparency, consolidated "unnecessary barriers to trade" under Rationale, and broadened Discrimination to include "differential treatment."

3 Data and Stylized Facts

This section presents the data and the construction of a new dataset that recovers the implementation timelines of contested TBTs at the WTO. I first outline the limitations of the WTO's Specific Trade Concerns (STC) database in identifying notification and enforcement dates. I then describe how I supplement this with regulatory documents to track the timing of TBTs and link these data to French firm-level exports. I conclude with stylized facts on the timing of TBT notifications.

3.1 Specific Trade Concerns Database

Identifying which technical regulations act as trade barriers is inherently challenging. Some measures may promote trade, as in the case of harmonized standards (Schmidt and Steingress, 2022). Following prior studies (Fontagné et al., 2015; Fontagné and Orefice, 2018; Herghelegiu, 2018), I focus on potentially restrictive TBTs using the WTO's database of Specific Trade Concerns (STCs) raised from 1995 to 2011.⁸ The

⁸Available at https://www.wto.org/english/res_e/publications_e/wtr12_dataset_e.htm.

database lists 318 STCs covering 403 distinct TBTs, with information on the implementing country, affected products (HS codes), policy objectives, and the concerns raised by WTO members. Importantly, STCs also target measures that were never formally notified.

However, the STC database only records the timing of WTO Committee discussions, not the dates when measures were notified or enforced.⁹ This limits the ability to distinguish between trade effects driven by the regulation's content and those arising from the lack of advance notice.

3.2 New Database on the Timelines of TBTs

To fill this gap, I reconstruct the full regulatory timeline for each measure, identifying both the official notification date and the entry-into-force date. These dates are often missing from standard STC and non-tariff measure databases.¹⁰ The reconstruction draws on multiple sources.

The first source is the WTO's Technical Barriers to Trade Information Management System (TBT IMS), which contains official notification documents. Using the unique "Document Symbol" identifier associated with many STCs, I scrape and parse these documents to extract notification and enforcement dates.¹¹ Where dates were missing or revised, I supplement this information by manually reviewing Revisions and Addenda.

The second source is the Minutes of the WTO TBT Committee, which summarize member discussions of each STC. These Minutes often clarify whether and when the measure was implemented, whether it had been notified, or if its entry was postponed

⁹Previous studies used the timing of STC discussions as a proxy for regulatory timing.

¹⁰Most NTM databases provide static regulatory snapshots and lack unique identifiers needed to track timelines.

¹¹See Appendix A.2.5 for extraction details. The TBT IMS is accessible at <http://tbtime.wto.org/en/Notifications/Search>.

or accelerated. I parsed and manually coded these discussions to extract relevant timeline information.¹²

Combining these sources, I recover both notification and implementation dates for 301 of the 403 contested TBTs—about 75% of the sample. Relying solely on notification documents would bias the sample toward more transparent cases, but by incorporating Committee Minutes, which cover all measures, I mitigate this concern. Recovery rates are comparable across notified measures (77%) and unnotified ones (67%), suggesting balanced coverage of transparent and opaque cases.

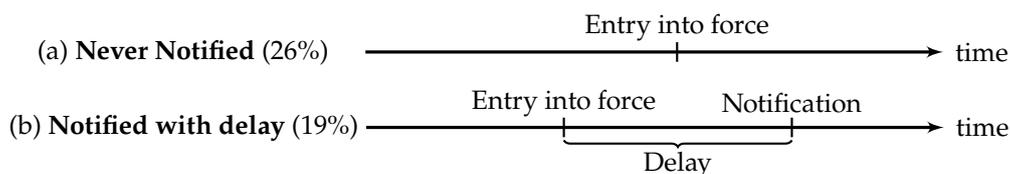
3.3 Stylized Facts about TBTs

A substantial share of TBTs are introduced without prior notice: 45% of the measures in my sample with known implementation timelines were enforced before any WTO notification. Table 1 shows that these measures are spread across multiple sectors, with chemicals, textiles, food products, and machinery most frequently affected. The sectoral composition of timely notified TBTs closely resembles that of all TBTs, suggesting that late or missing notifications are not disproportionately concentrated in particular industries.

The leading countries introducing new TBTs include the EU, China, Mexico, Korea, the US, and India (see Appendix, Table 9). China stands out due to its relatively recent WTO membership and active use of technical measures. Mexico and India, in particular, show a higher incidence of contested TBTs introduced without timely notification. An analysis of domestic industrial growth—proxied by destination-level export values from the BACI database (Appendix, Table 9)—indicates that TBTs without prior notification tend to target sectors experiencing faster export growth. However, the differences in sectoral growth rates between notified and unnotified cases remain modest.

¹²See Appendix A.2 for the data extraction process.

Figure 2: Types of TBT Implementation Without Timely Notification



I distinguish two types of TBTs without timely notice, as illustrated in Figure 2: (a) measures that are *Never Notified*, and (b) measures *Notified with Delay*. Never-notified cases include regulations, such as Mexico’s 1997 spirit labeling rule, which has never been submitted to the WTO notification system.¹³ Delayed notifications refer to measures implemented before being officially submitted, typically with a lag of about three months. For example, the Mexican ban on CFCs was enforced weeks before its notification. These cases reflect gaps in the timely implementation of the WTO’s transparency obligations.

Table 1: Product Coverage of TBTs by Notification Status

Sector	# TBT	% Timely Notified
Chemicals	381	50
Textiles	329	51
Vegetables	231	53
Machinery/Electrical	176	55
Foodstuffs	167	53
Animal Products	119	57
Miscellaneous	101	65
Metals	73	64
Transportation	58	57
Raw Hides, Skins, Leathers	41	56
Stone/Glass	37	32
Footwear/Headgear	36	50
Plastic/Rubbers	34	68
Wood Products	32	78
Mineral Products	17	71

Notes: A TBT is defined by a regulatory identifier paired with an HS4 code. The "% Notified" indicates the share of TBTs where official notification was provided before enforcement.

¹³This measure was contested by the EU, as documented in the WTO STC Database (item 20).

3.4 French Firm Level Data

The firm level data comes from two different sources: (i) the French customs, which reports exports for each firm by destination, product and month for the period between 1995 and 2007 and (ii) BRN (*Régime du bénéfice réel normal*), the French firm level administrative database which provides information on firms' balance-sheets, over the same period.¹⁴ The link to BRN, while it reduces the sample to relatively larger firms, allows to identify their principal activity. In this way, only manufacturing industries are selected, which are the ones directly interested by changes in technical requirements of production. Note that this dataset has been used in several trade related papers dealing with French data, such as Eaton et al. (2011) and Mayer et al. (2014).

I aggregate the monthly trade data at the semester level so that the final panel includes 26 periods.¹⁵ The choice of using a semestral panel comes from a trade off between the possibility to look at the specific dates (months) in which measures are enforced and the fact that there is both seasonality and lumpiness in the export behavior. For example, the median exporter ships a certain HS4 product in a destination twice a semester.

3.5 Descriptive Statistics about the Estimating Sample

Given the availability of French firm exports in the period 1995-2007, I use a sample of 123 TBTs that have been object of a STC by the European Union and that have been introduced in the same period.¹⁶ These 123 contested TBTs have been raised against a total of 31 different countries, with China being the largest contested country. Table 2

¹⁴Firms are obliged to comply with BRN status if they earn annual revenues larger than 763K €. The dataset is accessed through facilities provided by the INSEE (the French Statistical Institute) and were made available for analysis after careful screening to avoid disclosure of individual information.

¹⁵Semestral data are aggregated over two distinct six-month periods within a calendar year. The first period spans from January to June, and the second from July to December.

¹⁶Ideally, one would want to identify those TBTs that are trade restrictive specifically to French exporters, however, European countries participate as a single entity within the TBT Committee. On the other hand, technical regulation are homogeneous across EU countries and TBT measures are applied in a non-discriminatory way to all trading partners.

Table 2: Summary Statistics of TBT Coverage Across Countries and Products

	Mean	Median	Min	Max
# TBTs by Country	4	5.2	1	27
# HS4 categories by TBT	32	3	1	314

Notes: TBT is identified by a regulation identifier ("symbol"). Country is the one that implements the TBT ("maintaining country").

shows that there is a large variability in terms of HS4 categories covered by each TBT. The one that covers the largest number of products is the Mexican “Mandatory standard on Labelling of Industrial Products”, which interests several products across a large number of different HS4, from textile to food.

I limit the sample of firms to those that export to destinations outside the EU-27, as the STCs raised by the EU are directed towards non-EU countries. As for the products, since the WTO STC database typically records them at the 4-digit HS level, I group the export data at this level, which is coarser than the original 8-digit Combined Nomenclature.¹⁷

I then merge this data with tariff information from TRAINS at the HS4 level.¹⁸ The final estimation sample comprises product-destination pairs that satisfy two conditions: (i) the destination is economically significant (above the 10th percentile of French extra-EU exports), and (ii) at least one French firm exports the HS4 product to that destination for more than 4 semesters during the sample period. This second restriction ensures I focus on stable trade relationships where firms’ participation decisions are meaningful, excluding sporadic or experimental shipments.¹⁹

¹⁷The STC database contains 6% of HS2 products, 62% HS4s and the remaining 32% are HS6 goods. I keep the level of the analysis at 4-digit HS, aggregating the HS6 at the HS4 level.

¹⁸HS 4-digit applied tariff data is a simple average tariff within HS-4 headings of the HS-6 applied tariff level data; this aggregation is directly provided by TRAINS. Unfortunately, the database has many empty entries. In the literature, various algorithms have been proposed to increase the number of observations. I apply the interpolation procedure suggested in [Beverelli et al. \(2014\)](#).

¹⁹A similar cleaning procedure is applied in [Fontagné and Orefice \(2018\)](#). Destinations in the bottom 10th percentile of total French exports account for less than 0.001% of total extra-EU French export value. The number of countries in the sample is reduced by around 10%: from 223 non-EU countries

Table 3: Summary Statistics: French Export Markets and TBT Coverage

Panel A: Full Sample (All Products)		
	All Possible Markets	Active Markets
#(HS4, Country) pairs	44,563	29,460
#(Unique Exporters)	23,950	13,756
Avg. Semester Exports (EUR bn)	—	30.3
Panel B: TBT-Affected Markets Only		
	All Possible Markets	Active Markets
#(HS4, Country) pairs	705	483
#(Unique Exporters)	2,593	1,268
Avg. Semester Exports (EUR bn)	—	1.27
<i>TBT-Affected as Share of Full Sample:</i>		
Market pairs (%)	1.6	1.6
Exporters (%)	10.8	9.2
Export value (%)	—	4.2

Notes: All figures report semester averages over 26 semesters (1995–2007). “All Possible Markets” refers to all HS4–country pairs where at least one French firm exports for more than 4 semesters, representing stable market opportunities. “Active Markets” counts, in an average semester, only those pairs with positive exports. For “All Possible Markets,” “Unique Exporters” counts distinct firms that ever exported to each HS4–country pair for more than 4 semesters during 1995–2007. For “Active Markets,” it counts only firms with positive exports in an average semester. Total exports are reported only for Active Markets, as non-active markets contribute zero by definition.

Table 3 summarizes the estimation sample, with all statistics representing semester averages across the 26-semester period. The distinction between ‘All Possible Markets’ and ‘Active Markets’ is central to my empirical approach: the former (44,563 HS4-country pairs) defines the relevant choice set—stable market opportunities where French firms could potentially export—while the latter (29,460 pairs on average) captures realized export outcomes in a typical semester. “Unique Exporters” counts firms that met the presence criterion for each HS4–country pair at any point during 1995–2007 (All Possible Markets) versus only firms with positive exports in an average semester (Active Markets).

The gap between these columns reveals substantial churning in participation: approximately 66% of potential markets are active in any given semester, and active exporters represent 57% of all eligible firms. This variation in firms’ entry and exit decisions is

to 201. Excluded countries include Aruba, Samoa, Tonga, Cook Islands, etc.

the focus of the analysis.

Although TBT-affected markets represent less than 2% of both possible product-destination pairs (705 out of 44,563) and active export relationships (483 out of 29,460), they account for around 1.27 billion euros in average semester exports—roughly 4.2% of total extra-EU exports—highlighting that despite their rarity, TBTs target trade flows of considerable economic importance.

4 Theoretical Framework

This section develops a real-options framework that explains how timely TBT notifications reduce regulatory uncertainty and foster export participation. The analysis builds on partial-equilibrium models of trade under uncertainty (Dixit et al., 1994; Handley and Limão, 2017) and treat *official notification* as a reduction in the dispersion of regulatory outcomes.

4.1 Exporting Under Uncertain Standards

Consider a firm with productivity φ operating in a monopolistically competitive market. Per-period operating profit from exporting is

$$\pi(\varphi, d) = a d^{1-\sigma} \varphi^{\sigma-1}, \quad \sigma > 1, \quad (1)$$

where σ is the elasticity of substitution, a is a destination-specific demand shifter, and $d \geq 1$ is an iceberg compliance-cost multiplier capturing TBT restrictiveness: economically, the per-unit cost uplift of producing to the destination's standard (e.g., certified inputs, different testing/labeling costs, or process changes).²⁰

²⁰See, e.g., Parenti and Vannoorenberghe (2025), where the unit cost of complying with a foreign standard increases with the standard's stringency (e.g., a tighter emissions cap).

To serve the market, the firm pays also a sunk adaptation cost K_s ; once paid, profits are discounted by $\beta \in (0, 1)$.²¹

Timing and information. At $t = 0$ the incumbent requirement is $d_0 \geq 1$. A new TBT takes effect at $t = 1$. At $t = 0$ the firm chooses whether to (i) *enter now* and pay K_s , or (ii) *wait* until $t = 1$ when the realized requirement is observed and then decide. Notification can arrive *before* enforcement (timely), *at or after* enforcement (delayed), or *never*. Notifications are modelled as reducing the dispersion of d ; under timely notification, dispersion is (approximately) zero at $t = 0$.²²

4.2 Two-Period Option–Timing Illustration

To isolate the timing value of transparency, consider a two-period model. At $t = 1$, the TBT is enforced, and the compliance multiplier becomes stochastic. For tractability, adopt a two-point prior with fixed mean $d_0 \geq 1$ and dispersion parameter $\delta \in [0, 1)$, defining high- and low-restrictiveness scenarios:

$$d \in \{d_H, d_L\}, \quad d_H = d_0(1 + \delta), \quad d_L = d_0(1 - \delta), \quad \Pr(d_H) = \Pr(d_L) = \frac{1}{2}, \quad (2)$$

with δ chosen so $d_L \geq 1$.²³ Given these two possible realizations, one express the productivity thresholds that make export entry profitable under each outcome:

$$\varphi_H \equiv \left(\frac{K_s}{a d_H^{1-\sigma}} \right)^{1/(\sigma-1)}, \quad \varphi_L \equiv \left(\frac{K_s}{a d_L^{1-\sigma}} \right)^{1/(\sigma-1)}, \quad \varphi_L \leq \varphi_H.$$

²¹Allowing $K_s(d)$ or ex-ante over-compliance $K^R \geq K_s(d_{\max})$ preserves the option-value logic. In addition, an exogenous survival probability is absorbed into β ; results are unaffected.

²²An observationally equivalent formulation loads uncertainty on demand: with $\pi(\varphi) = m \varphi^{\sigma-1}$ and $m \equiv a d^{1-\sigma}$, notification reduces the dispersion of m —either by clarifying the requirement (d) or by revealing market acceptance/willingness to pay for the compliant product (a). All comparative statics go through unchanged.

²³Results extend to any mean-preserving contraction of d .

There are three regions: (i) $\varphi \geq \varphi_H$ where both realizations are profitable; (ii) $\varphi \in [\varphi_L, \varphi_H)$ where only the low-cost realization is profitable; (iii) $\varphi < \varphi_L$ where neither are profitable.

Invest now vs. wait. At $t = 0$ the firm compares the expected discounted value from investing immediately (V_0) with the value of waiting (V_1):

$$V_0(\varphi) = -K_s + \pi(\varphi, d_0) + \beta \mathbb{E}[\pi(\varphi, d)], \quad (3)$$

$$V_1(\varphi) = \beta \mathbb{E}[\max\{\pi(\varphi, d) - K_s, 0\}], \quad (4)$$

and chooses $F_0(\varphi) = \max\{V_0(\varphi), V_1(\varphi)\}$. If entry cannot be postponed, the value is $\Omega_0(\varphi) = \max\{V_0(\varphi), 0\}$ (outside option is 0). The real-option value of waiting is the wedge $F_0(\varphi) - \Omega_0(\varphi) \geq 0$, which increases with dispersion δ (Appendix A.1).

The entry cutoff is

$$\varphi^*(\delta) = \inf\{\varphi : V_0(\varphi) \geq V_1(\varphi)\}.$$

Because $d^{1-\sigma}$ is strictly convex for $\sigma > 1$, reducing dispersion via timely notification ($\delta \rightarrow 0$) lowers the option value of waiting and weakly lowers the entry cutoff, an effect that operates in region (ii). This intuition is formalized by comparing the firm's value from immediate entry with the value of waiting for uncertainty to resolve.

Proposition 1 (Cutoff monotonicity). *Holding $\mathbb{E}[d] = d_0$ fixed, the entry cutoff $\varphi^*(\delta)$ is weakly increasing in dispersion δ . It is unchanged when both standard realizations induce entry at $t = 1$, and strictly increasing when only the favorable realization would induce entry upon resolution.*

Proof sketch. For $\sigma > 1$, the function $d^{1-\sigma}$ is decreasing and strictly convex. Hence $\pi(\varphi, d)$, and therefore $\max\{\pi(\varphi, d) - K_s, 0\}$, are convex in d . By Jensen's inequality, a mean-preserving spread in d increases $\mathbb{E}[\max\{\cdot\}]$, raising V_1 relative to V_0 and thus

(weakly) increasing the cutoff $\varphi^*(\delta)$. When both realizations satisfy $\pi(\varphi, d) \geq K_s$, waiting merely forgoes the period-0 profit $\pi(\varphi, d_0)$ while leaving the period-1 decision unchanged, so the cutoff remains locally invariant. See Appendix A.1.

Temporary exits. The framework extends naturally to temporary exits. After paying the sunk adaptation cost K_s , suppose the firm faces a small per-period fixed continuation cost $f > 0$. The post-adaptation rule is to export in period t iff $\pi(\varphi, d_t) \geq f$. Under uncertainty about d_t , intermediate-productivity firms optimally suspend exporting while the perceived compliance burden is high and resume when information arrives or the requirement is favourable. Formally, temporary exit occurs for $\varphi \in [\varphi_f(d_L), \varphi_f(d_H))$, where $\varphi_f(d) \equiv \left(\frac{f}{a d^{1-\sigma}}\right)^{1/(\sigma-1)}$. A mean-preserving contraction of d (timely notification) shrinks this interval, reducing the incidence of temporary exits and accelerating re-entry. The same logic applies if a one-off recertification cost K_r is required under the new realised requirement.

4.3 Implications

The model treats timely notification as lowering δ *before* the investment decision, delayed notification as reducing δ only *at or after* enforcement, and never-notified measures as leaving δ high. By Proposition 1, notification reduces the option value of waiting and weakly lowers the entry cutoff, encouraging earlier adaptation and sustained participation. Because d enters multiplicatively like an iceberg trade cost, the *transparency premium*—a reduction in dispersion—admits a tariff-equivalent interpretation.²⁴

This framework generates four testable predictions. *First*, notification should operate primarily on the extensive margin—lowering the option value of waiting increases entry likelihood while leaving conditional export values largely unchanged. *Second*,

²⁴Section 8 implements this using a PPML estimate and the equivalence between $d^{1-\sigma}$ and $(1 + \tau)^{1-\sigma}$ in CES revenue.

among firms that have already entered, notification should reduce short-lived exits and accelerate entry, with little effect on permanent exits driven by fundamentals. *Third*, the benefits of notification should be greatest for low- and mid-productivity firms near the entry cutoff. *Fourth*, export participation should rebound when information arrives—when δ collapses upon notification—even if notice follows enforcement, ruling out a pure adjustment-time explanation.

5 Research Design

This section develops an empirical strategy to test the theoretical predictions. I exploit variation in TBT notification timing to estimate how regulatory uncertainty affects firm export decisions. The analysis proceeds in three steps: variable definitions, identification strategy, and validation.

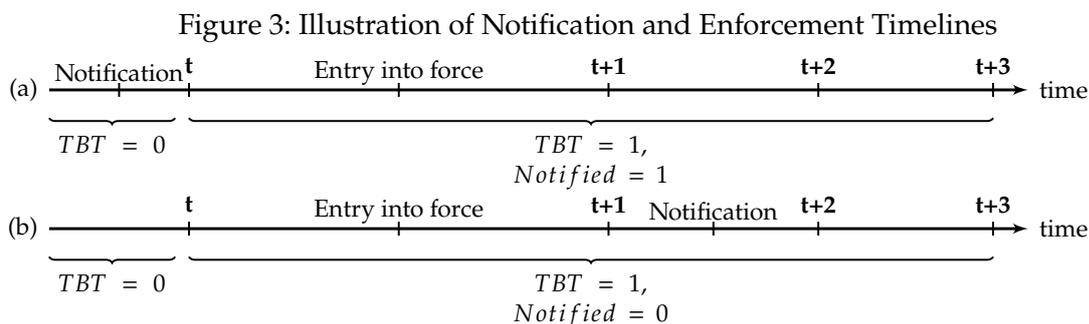
5.1 Definition of Variables

The enforcement date of TBT regulations is denoted by E_{pd} , where pd represents the combination of a 4-digit HS product category and destination country. Since the analysis uses semester-level data, I_{pd} denotes the semester during which TBT enforcement begins. The main treatment variable is $TBT_{pdt} = \mathbb{1}\{t \geq I_{pd}\}$, which equals one from the semester of TBT introduction onwards. In markets where no restrictive TBT has been introduced, this indicator remains zero throughout.

The notification date N_{pd} is used to define whether notification was timely, $Notified_{pd} = \mathbb{1}\{N_{pd} \neq \text{NA and } N_{pd} \leq E_{pd}\}$. This indicator equals one if the TBT was notified by the time of enforcement and zero otherwise—encompassing both measures that were never notified and those notified only after enforcement began. In the baseline analysis, I focus on this binary distinction: whether firms had access to official notification before the TBT took effect. I return to the timing of delayed notifications in Section 6, where

I exploit variation in when information became available to disentangle uncertainty from adjustment costs.

Figure 3 illustrates these indicators for a given country-product pair. In both panels, the TBT is enforced in semester t . Panel (a) shows timely notification (Notified = 1), while Panel (b) shows a case without timely notification (Notified = 0)—here specifically showing late notification, though this category also includes measures never notified at all.



Notes: Illustration of two ideal timelines. Each timeline is product-country specific, but the index pd is dropped.

5.2 Empirical Strategy

To quantify the effect of transparency, I estimate the following specification:

$$y_{ipdt} = \alpha \text{TBT}_{pdt} + \beta (\text{TBT}_{pdt} \times \text{Notified}_{pd}) + \mu_{ipd} + \mu_{HS2,d,t} + \varepsilon_{ipdt}, \quad (5)$$

where y_{ipdt} denotes the export outcome of firm i , HS-4 product p , destination d , and semester t . The specification covers both extensive and intensive margins: the extensive margin is a binary indicator for any positive exports; the intensive margin is the log of export value (EUR).

The coefficient α captures the average effect of a TBT when it is implemented without prior notification, while β measures the incremental mitigating effect of timely noti-

cation. A negative and statistically significant β indicates that advance notice reduces the trade-disrupting impact of the regulation, consistent with the interpretation that transparency lowers uncertainty or facilitates adjustment.

The regression includes HS-2–destination–semester fixed effects ($\mu_{HS2,dt}$) to account for common demand shocks, changes in destination-specific price indices, and sector-level policies. It also includes firm–product–destination fixed effects (μ_{ipd}) to control for time-invariant characteristics such as firm productivity, product quality, or the strategic importance of a market–product combination.

Identification relies on comparing changes in firm-level trade margins before and after the enforcement of a TBT, contrasting treated product–destination pairs with untreated ones within the same (narrowly defined) sector–market context. The assumption is that in the absence of the policy shock, treated and untreated products within the same HS-2–destination cell would have followed parallel trends. Under this assumption, α and β can be interpreted as causal effects. Standard errors are clustered at the product–destination–semester level to allow for arbitrary correlation within treatment cells (Abadie et al., 2017).

5.3 Identification

A key identification concern is reverse causality: regulators might introduce TBTs in response to rising imports, or the EU might challenge measures following poor export performance. Such endogeneity would bias estimates toward reflecting prior trade dynamics rather than causal policy effects.

Two features of the research design mitigate this concern. First, STCs are raised at the EU level rather than by individual member states, making it unlikely that challenges respond specifically to French export patterns. Second, the specification includes saturation-level fixed effects that absorb most plausible confounds.

The firm-product-destination fixed effects μ_{ipd} control for all time-invariant determinants of the trade relationship, while the HS2-destination-time fixed effects $\mu_{HS2,dt}$ absorb time-varying shocks common to all products within a sector-destination pair. Identification stems from within-firm variation: comparing the same firm's exports of treated versus untreated HS-4 products to the same destination, relative to the sectoral average. For bias to remain, unobserved shocks would need to (i) affect specific HS-4 products differentially within HS-2 sectors, (ii) coincide precisely with TBT enforcement dates, and (iii) correlate systematically with WTO notification status.

I validate this strategy testing for pre-trends, estimating the following specification:

$$y_{ipdt} = \alpha \text{TBT}_{pdt} + \beta (\text{TBT}_{pdt} \times \text{Notified}_{pd}) + \sum_{n=-3}^{-1} \beta_n \mathbb{1} \{ \text{TBTtime}_{pdt} = n \} + \mu_{ipd} + \mu_{HS2,dt} + \varepsilon_{ipdt}, \quad (6)$$

where $n = 0$ denotes the semester of enforcement, $n < 0$ indexes pre-enforcement periods, and the reference category includes all periods at least 4 semesters before enforcement ($n \leq -4$). This cutoff reflects the longest advance-notification window in the data (14 months, approximately 3 semesters), allowing the specification to distinguish anticipation effects from spurious pre-trends.

Under the identifying assumption, pre-enforcement coefficients β_n should be zero for non-notified TBTs, where anticipation effects are absent, but may differ from zero for notified TBTs due to advance preparation.

5.4 Baseline Results

Table 4 presents estimates of Equation 5 for both the extensive margin (export participation) and intensive margin (export value). The results reveal two key findings.

First, TBTs impose substantial trade costs. The introduction of a restrictive TBT reduces the probability that a firm exports a given product to the affected destination by 3.9

Table 4: TBTs and timely notification effect - Baseline Model

DEPENDENT VARIABLE	Participation		Export Value	
	(1)	(2)	(3)	(4)
TBT	-0.039*** (0.008)	-0.072*** (0.019)	0.012 (0.035)	-0.002 (0.090)
TBT × Notified		0.042** (0.021)		0.018 (0.096)
Observations	8,155,082	8,155,082	4,492,290	4,492,290
R2	0.247	0.247	0.748	0.748
Firm-HS4-Country FE	Yes	Yes	Yes	Yes
HS2-Country-Time FE	Yes	Yes	Yes	Yes

Notes: Participation is a dummy variable for firm-product export participation (the extensive margin of exports), meaning positive trade flows into a certain product-destination market combination. The observations in cols 1 and 2 are larger than in 3 and 4 since participation accounts for trade zeros. Significance levels: * < 0.1, ** < 0.05, *** < 0.01.

percentage points (column 1), while leaving the intensive margin—conditional export values—statistically unchanged (column 3). This pattern aligns with [Fontagné and Orefice \(2018\)](#).

Second, timely notification significantly mitigates these disruptions. When TBTs are notified before enforcement, the negative effect on participation is reduced by 4.2 percentage points (column 2)—offsetting more than half of the baseline decline. Firms receiving advance notice face a net participation loss of only 3 percentage points ($0.072 + 0.042$), compared to 7.2 percentage points when notification is absent. The intensive margin remains unaffected (column 4), confirming that transparency operates primarily through participation decisions rather than through adjustments in shipment size.

These findings align with the theoretical predictions: transparency reduces uncertainty, lowering the option value of waiting and supporting continued market participation. As expected, the impact concentrates on the extensive margin—firms choose to serve the market rather than delay entry or exit temporarily.

5.5 Robustness Checks

This subsection presents four robustness checks that validate the identification strategy and address potential confounding factors.

Testing for Pre-Existing Trends To validate the parallel trends assumption, I estimate the event study specification from Equation 6. The results, reported in Appendix Table 10, show that coefficients for all pre-TBT periods are statistically insignificant, confirming the absence of pre-trends and supporting a causal interpretation of the baseline results. Importantly, firms do not exit markets before TBTs are actually enforced, ruling out anticipatory selection as a confound.

Accounting for Sectoral Heterogeneity. A potential concern is that the notification effect could reflect sectoral composition: if timely notifications are concentrated in sectors with lower trade elasticities, the estimated mitigation might spuriously capture differences in demand sensitivity rather than reduced uncertainty. To address this, I augment the baseline specification by interacting the TBT treatment with product-level elasticity estimates from Broda and Weinstein (2006) and Fontagné et al. (2022). The interaction term is negative—consistent with higher-elasticity products being more sensitive to trade frictions—but small and statistically insignificant. Critically, the main coefficients on *TBT* and *TBT* × *Notified* remain stable and significant, indicating that the mitigating effect of notification is not driven by sectoral differences in demand elasticities. The limited role of elasticity heterogeneity likely reflects the saturation of HS2–country–time fixed effects, which absorb much of the sector-level variation. Full results are reported in Appendix Table 11.

Controlling for Concurrent Trade Policies. A concern is that TBTs might be introduced alongside changes in tariffs, potentially confounding the estimates. Prior work suggests that TBTs sometimes substitute for tariff reductions (Beverelli et al., 2014;

Orefice, 2017), meaning that countries may tighten technical regulations as they lower tariffs. This raises the possibility that the observed effects reflect the combined impact of both policy instruments rather than the informational value of notification itself. To address this, I augment the baseline specification with product-level applied tariff rates at the destination level. As shown in Table 12, controlling for tariffs leaves the magnitude and significance of the notification interaction unchanged. This indicates that the mitigating effect of timely notification operates independently of concurrent tariff changes and reflects the value of transparency per se.

Alternative Estimator: Poisson Pseudo Maximum Likelihood. Given the prevalence of zero trade flows in firm-level data, I re-estimate the baseline model using a Poisson Pseudo Maximum Likelihood (PPML) estimator, which accounts for heteroskedasticity and retains zero-valued export observations (Santos Silva and Tenreyro, 2006; Correia et al., 2019). Results in Table 13 confirm that restrictive TBTs significantly reduce exports and that timely notification mitigates this decline, consistent with the OLS estimates. The consistency of results across estimators strengthens the conclusion that regulatory transparency reduces the trade-dampening effects of TBTs.

6 Channels: Investment Decisions and Export Dynamics

The baseline results demonstrate that timely notification mitigates TBT-induced trade disruptions on the extensive margin. This section investigates the underlying mechanism by decomposing participation responses into permanent exits, temporary exits, and entry. The theoretical framework predicts that advance notification lowers the option value of waiting, so firms commit to participation rather than suspending exports or delaying entry until uncertainty resolves.

6.1 Decomposition of the Extensive Margin

I classify export dynamics for each firm–product–destination triplet into three categories. *Permanent exits* occur when firms cease serving a market and never return. *Temporary exits* capture firms that suspend exports but subsequently re-enter. *Entry* identifies firms absent from a market for at least two semesters that start exporting. This decomposition separates decisions driven by fundamentals (permanent exits) from those reflecting real-options behavior under uncertainty (temporary exits and delayed entry). The model predicts that notification should primarily affect the latter margins, where waiting has value, while leaving permanent exits largely unchanged.

Table 5 reports the estimates. For temporary exits, a TBT introduced without notice raises the exit probability by 5.4 percentage points. The notification interaction of -6.3 percentage points implies that advance notice eliminates this effect entirely ($0.054 - 0.063 \approx 0$). Similarly, entry falls by 7.2 percentage points without notification but only 2 percentage points when firms receive advance notice ($-0.072 + 0.053 \approx -0.02$), offsetting roughly three-quarters of the decline. Together, these results indicate that firms respond to unannounced measures by suspending exports and delaying entry—behavior that vanishes when compliance requirements are clarified ex ante.

Permanent exits show little response to notification. While TBT introduction raises permanent exits by 6.2 pp when pooling all measures, separating notified from unnotified cases yields small and imprecise coefficients. This is consistent with permanent exits reflecting fundamental incompatibilities that transparency cannot resolve.

Overall, the decomposition shows that transparency operates through timing: advance notice curbs temporary exits and sustains entry, while leaving exits driven by fundamentals mostly unaffected—evidence that notification reduces the option value of waiting by clarifying requirements before enforcement.

Table 5: Decomposition of Export Extensive Margins

DEPENDENT VARIABLE	Permanent Exit:		Short-term Exit:		Entry:	
	(1)	(2)	(1)	(2)	(1)	(2)
$\mathbb{1}\{\text{TBTtime} = 0\}$	0.062*** (0.019)	0.027 (0.039)	0.001 (0.014)	0.054** (0.021)	-0.025*** (0.008)	-0.072*** (0.021)
$\mathbb{1}\{\text{TBTtime} > 0\}$	0.003 (0.014)	-0.006 (0.019)	0.019 (0.013)	0.042* (0.022)	-0.001 (0.007)	0.002 (0.019)
$\mathbb{1}\{\text{TBTtime} = 0\} \times \text{Notified}$		0.040 (0.039)		-0.063*** (0.023)		0.053** (0.021)
$\mathbb{1}\{\text{TBTtime} > 0\} \times \text{Notified}$		0.008 (0.023)		-0.024 (0.026)		-0.006 (0.019)
Observations	4,584,520	4,584,520	3,866,822	3,866,822	4,241,645	4,241,645
R ²	0.507	0.507	0.422	0.422	0.279	0.279
Firm-HS4-Country FE	Yes	Yes	Yes	Yes	Yes	Yes
HS2-Country-Time FE	Yes	Yes	Yes	Yes	Yes	Yes

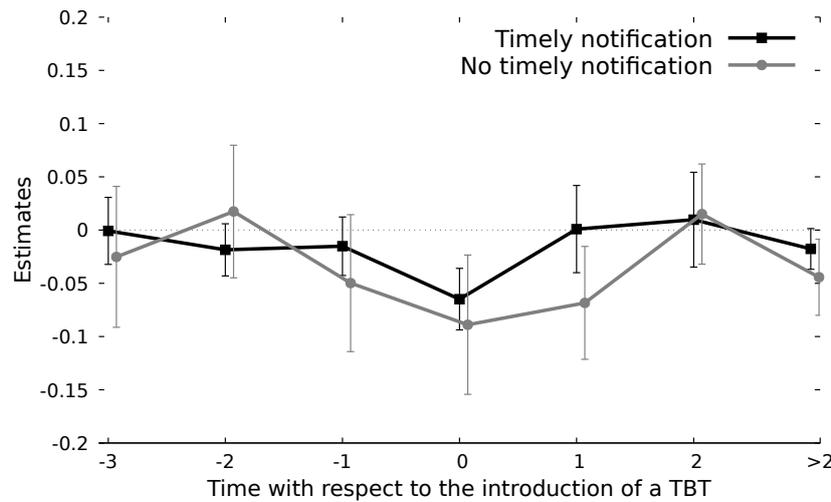
Notes: Permanent Exit is a binary indicator set to 1 if the firm was present at least once in the last two semesters but will not ever serve the product-destination market again. Short-term exit takes value 1 if the firm was serving the market in the current semester but will not serve the market for the next two semesters, but will resume serving it afterwards. The dummy entry takes value 1 if the firm was not present in the product-destination at least for the last two semesters but it operates in the current one and will do so in the future. Standard errors in parentheses are clustered at (product, country, time). Significance levels: * < 0.1, ** < 0.05, *** < 0.01.

6.2 Duration of the Notification Effect

To examine persistence, I estimate an event-study with leads and lags around enforcement. Two patterns emerge from Figure 4. First, pre-trends are flat and statistically indistinguishable across groups. Second, the participation gap opens at enforcement—unnotified measures exhibit a larger immediate decline—and then narrows, closing within two semesters.

This convergence raises a key question: does it reflect adjustment time or information arrival? If firms mainly need time to adapt physically, delayed notifications issued after enforcement should have little immediate effect. If uncertainty is the barrier, participation should rebound when notification arrives, regardless of timing. The next subsection tests this distinction by asking whether delayed notifications trigger an immediate export recovery.

Figure 4: Export Participation Before and After TBT Enforcement by Notification Status



Notes: Coefficients from Equation 15, showing changes in export participation before and after TBT enforcement. Trade disruptions are short-lived when advance notification is provided. Full estimates are reported in Table 14 (Appendix). The difference in the first two post-enforcements semester is statistically significant at the 5% level ($F=4.25, p=0.0391$).

6.3 Mechanisms: Waiting vs. Adjustment Time

Does timely notification help primarily by providing adjustment time or by resolving uncertainty? If adjustment time is critical, delayed notifications issued after enforcement should have minimal immediate impact—firms would still need time to comply. If uncertainty is the key barrier, participation should rebound immediately when information arrives, even post-enforcement.

To test this, I track export participation for up to two semesters after enforcement, comparing TBTs that were notified on time, notified with delay, or never notified. The specification extends the baseline model by interacting post-enforcement time dummies with these three notification statuses.

The results in Table 6 reveal three patterns. First, the initial decline in participation at enforcement is markedly smaller for measures introduced with timely notification

(5-6 percentage points) compared to those lacking prior notice (8 points). Second, for markets where no notification is ever issued, participation remains persistently depressed in the following semester (9-10 points), consistent with unresolved uncertainty. Third, and critically, when a notification eventually arrives—typically one semester after enforcement—participation rebounds substantially, closing the gap relative to timely-notified measures.

Importantly, at the time of enforcement ($t=0$), there is no significant difference in export dynamics between measures that are eventually notified and those that are never notified.²⁵ This suggests that both groups of firms confront comparable initial uncertainty. The divergence in trajectory occurs only after notification is issued, supporting the argument that it is the provision of information—not the passage of time—that enables firms to resume exports.

In the long run (2-3 semesters), participation stabilizes at a 2–3 percentage-point decline, consistent with Section 6.2. Overall, the evidence points to uncertainty—rather than adjustment time—as the key mechanism: timely notification reduces the option value of waiting by clarifying compliance requirements; when notification is delayed, exports rebound only once information arrives.

7 Firm-Heterogeneous Responses

The theoretical framework predicts heterogeneous effects across firms. Smaller and less productive firms face higher real-option values of waiting: they are closer to the threshold productivity cutoff where uncertainty makes delaying market entry optimal. Reducing uncertainty through advance notification lowers this threshold, enabling a wider range of firms to profitably engage in trade.

²⁵Estimating a saturated model with separate indicators for the two groups yields no significant difference at $t=0$ ($F(1, 974, 934) = 1.3, p = 0.3$).

Table 6: Export Participation After TBT Enforcement, by Timing of Notification Arrival

	Dependent variable: participation	
	(1)	(2)
<i>Semester 0 (enforcement)</i>		
Timely notification	-0.054*** (0.015)	-0.063*** (0.015)
No notification yet	-0.077** (0.033)	-0.077** (0.033)
<i>Semester +1 (relative to semester 0)</i>		
Still no notification	-0.092*** (0.032)	-0.100*** (0.031)
Already notified	0.014 (0.022)	0.005 (0.021)
Notification arrives now	0.060 (0.038)	0.040 (0.034)
<i>Semester +2 (relative to semester 0)</i>		
Still no notification		0.002 (0.028)
Already notified		0.018 (0.023)
Notification arrives now		0.012 (0.032)
<i>Long-run (semester $\geq +2$ in col. 1, $\geq +3$ in col. 2)</i>		
	-0.023*** (0.008)	-0.025*** (0.009)
Observations	8,155,082	8,155,082
R ²	0.247	0.247
Firm-HS4-Country FE	Yes	Yes
HS2-Country-Time FE	Yes	Yes

Notes: The dependent variable equals 1 if firm i exports product p to destination d in semester t , 0 otherwise. "Timely notification" denotes measures whose notice preceded or coincided with enforcement. "Notification arrives now" captures measures first notified in the indicated semester. "Still no notification" refers to measures for which no notice has yet been issued. Standard errors (in parentheses) are clustered by product–destination–semester. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7 tests this prediction by interacting the notification indicator with quartiles of firm size—measured by domestic sales (col. 1) and head-count employment (col. 2), used as proxies for productivity.

The interaction terms reveal that notification benefits are concentrated among small and mid-sized firms. For the smallest quartile (Bin 1), advance notice raises participation by 5-6 percentage points. Firms in Bin 2 show similarly large effects (4-6 percentage points), while Bin 3 exhibits comparable point estimates (4-5 percentage points) that are marginally significant with sales but imprecise with employment. In contrast, the largest firms (Bin 4) show small (1-2 percentage points) and statistically insignificant

effects, suggesting that these exporters—with greater internal resources and larger profit margins—are less sensitive to policy uncertainty.

These patterns, consistent across both firm-size measures, confirm the model’s prediction: the value of transparency is largest where the real-option value of waiting is highest. Timely notification particularly benefits small and mid-sized exporters, for whom regulatory uncertainty poses the greatest barrier to trade participation.

Table 7: Notification Effects by Firm Size (Bins Based on Domestic Sales or Employment)

	Dependent Variable: Participation	
	(1) Sales Bins	(2) Employment Bins
TBT	-0.066*** (0.022)	-0.063*** (0.022)
TBT × Notified × Bin 1	0.049** (0.025)	0.062** (0.025)
TBT × Notified × Bin 2	0.041* (0.024)	0.060** (0.024)
TBT × Notified × Bin 3	0.046* (0.024)	0.038 (0.024)
TBT × Notified × Bin 4	0.015 (0.024)	0.008 (0.024)
Bin 2	0.131*** (0.001)	0.125*** (0.001)
Bin 3	0.250*** (0.002)	0.231*** (0.001)
Bin 4	0.356*** (0.002)	0.307*** (0.002)
Observations	6,780,420	6,973,438
R ²	0.302	0.298

Notes: Coefficients from OLS regressions of export participation on TBT measures and their interactions with firm size bins. Bin 1 is the reference category. Bin numbers increase with firm size: column (1) uses domestic sales bins, column (2) uses employment bins. Standard errors clustered at the product–destination level are reported in parentheses.

Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

8 Tariff-Equivalent Magnitude

To express the transparency effect in tariff units, note that under CES demand with constant markups, both ad valorem tariffs and iceberg-type regulatory costs enter multiplicatively in revenues: $Y \propto d^{1-\sigma}(1 + \tau)^{1-\sigma}$. Hence the PPML interaction coefficient $\hat{\beta}$ on TBT × Notified from the baseline specification can be mapped to an equivalent ad valorem tariff change.

PPML estimates expected revenues in levels and retains zero flows, so coefficients reflect adjustments on both the intensive (value) and extensive (participation) margins, while accommodating firm–product–destination triplets that switch in and out of exporting.

Equating revenues under a notified TBT to those obtained via a compensating change in tariffs yields the tariff-equivalent magnitude (in percentage points):

$$\Delta t = (100 + t_0) \left(\exp\left(\frac{\hat{\beta}}{\sigma-1}\right) - 1 \right),$$

where t_0 is the baseline tariff (in percentage points) and $\sigma > 1$ is the elasticity of substitution.²⁶ needed to replicate that gain. The algebraic derivation is in Appendix A.4.

Applying this formula with $\hat{\beta} = 0.450$ (Table 13) and conventional values $\sigma \in [3, 7]$, timely notification is equivalent to a tariff reduction of roughly 8–28 percentage points, depending on σ and t_0 (Table 8). For example, with $\sigma = 3$ and $t_0 = 5.5\%$ (the median extra-EU tariff), the implied reduction is 26.6 percentage points; with $t_0 = 10.6\%$ (the median in TBT-affected markets), it is 27.8 percentage points.

These magnitudes depend on CES demand with constant markups and vary with the assumed σ and baseline tariff. The full mapping and sensitivity checks are in Appendix A.4.

Table 8: Tariff Equivalent of Notification under Alternative σ Values

Elasticity σ	Δt if $t_0 = 5.5\%$	Δt if $t_0 = 10.6\%$
3	26.6 pp	27.8 pp
4	17.2 pp	17.9 pp
5	12.7 pp	13.2 pp
7	8.3 pp	8.7 pp

Notes: Δt expresses the tariff equivalent of timely notification in percentage points, assuming a baseline tariff of either 5.5% (the extra-EU median) or 10.6% (the median in markets with a TBT). Calculations use $\hat{\beta} = 0.450$ from column (2) of Table 13. “pp” denotes percentage points.

²⁶We report $\Delta t > 0$ as the magnitude of the tariff change. It can be read as the compensating tariff increase that would undo the revenue gain from notification—or, equivalently, the same-magnitude cut

9 Conclusion

This paper provides causal evidence on the economic value of transparency in administering Technical Barriers to Trade. Issuing a WTO notification before a regulation enters into force halves its negative impact on export participation—particularly for smaller exporters—and yields a benefit equivalent to avoiding a tariff increase of up to 28 percentage points.

I extend the trade-policy uncertainty literature to technical regulations by identifying *advance notification* as a concrete channel that lowers uncertainty. The timing evidence—exports rebound at information arrival rather than gradually—supports a real-options mechanism in which firms delay investment until compliance costs are clarified. In short, *how* a regulation is introduced matters as much as *what* it requires.

For policymakers, enforcing the 60-day advance-notice guideline is a low-cost, high-impact lever, delivering trade-facilitation benefits comparable to sizable tariff reductions. With multilateral rule-making stalled and dispute settlement impaired, the WTO's credibility increasingly hinges on functions it can still perform well; ensuring timely, complete, and searchable notifications is one such function.

An open question remains: why do governments sometimes choose opacity? Since most contested TBTs are resolved without litigation (Ghodsi and Michałek, 2016), delayed notification may grant domestic industries a temporary protective window with limited retaliation risk. Understanding this political-economy rationale is a promising direction for future research.

References

- ABADIE, A., S. ATHEY, G. W. IMBENS, AND J. WOOLDRIDGE (2017): "When should you adjust standard errors for clustering?" Tech. rep., National Bureau of Economic Research.
- ALESSANDRIA, G. A., S. Y. KHAN, AND A. KHEDERLARIAN (2019): "Taking stock of trade policy uncertainty: Evidence from china's pre-wto accession," Tech. rep., National Bureau of Economic Research.
- BAO, X. AND L. D. QIU (2012): "How do technical barriers to trade influence trade?" *Review of International Economics*, 20, 691–706.
- BEGHIN, J. C., A.-C. DISDIER, AND S. MARETTE (2015): "Trade restrictiveness indices in the presence of externalities: An application to non-tariff measures," *Canadian Journal of Economics/Revue canadienne d'économique*, 48, 1513–1536.
- BERNANKE, B. S. (1983): "Irreversibility, uncertainty, and cyclical investment," *The quarterly journal of economics*, 98, 85–106.
- BEVERELLI, C., M. BOFFA, AND A. KECK (2014): "Trade policy substitution: Theory and evidence from Specific Trade Concerns," Tech. rep., WTO Staff Working Paper.
- BRODA, C. AND D. E. WEINSTEIN (2006): "Globalization and the Gains from Variety," *The Quarterly journal of economics*, 121, 541–585.
- CARBALLO, J., K. HANDLEY, AND N. LIMÃO (2018): "Economic and Policy Uncertainty: Export Dynamics and the Value of Agreements," NBER Working Papers 24368, National Bureau of Economic Research, Inc.
- COELLI, F. (2022): "Trade policy uncertainty and innovation: Evidence from China," *Available at SSRN 4169514*.
- CORREIA, S., P. GUIMARAES, AND T. ZYLKIN (2019): "PPMLHDFE: Stata module for Poisson pseudo-likelihood regression with multiple levels of fixed effects," .
- DIXIT, A. K., R. K. DIXIT, AND R. S. PINDYCK (1994): *Investment under uncertainty*, Princeton university press.
- EATON, J., S. KORTUM, AND F. KRAMARZ (2011): "An anatomy of international trade: Evidence from French firms," *Econometrica*, 79, 1453–1498.
- FONTAGNÉ, L., H. GUIMBARD, AND G. OREFICE (2022): "Tariff-based product-level trade elasticities," *Journal of International Economics*, 137, 103593.
- FONTAGNÉ, L. AND G. OREFICE (2018): "Let's try next door: Technical Barriers to Trade and multi-destination firms," *European Economic Review*, 101, 643–663.
- FONTAGNÉ, L., G. OREFICE, R. PIERMARTINI, AND N. ROCHA (2015): "Product standards and margins of trade: Firm-level evidence," *Journal of international economics*, 97, 29–44.
- FRANCOIS, J. F. (2001): "Trade policy transparency and investor confidence: some implications for an effective trade policy review mechanism," *Review of International Economics*, 9, 303–316.

- GHODSI, M. G. AND J. J. MICHAŁEK (2016): "Technical barriers to trade notifications and dispute settlement within the WTO," *Equilibrium. Quarterly Journal of Economics and Economic Policy*, 11, 219–249.
- GROSSMAN, G. M., P. MCCALMAN, AND R. W. STAIGER (2021): "The "New" Economics of Trade Agreements: From Trade Liberalization to Regulatory Convergence?" *Econometrica*, 89, 215–249.
- HANDLEY, K. AND N. LIMÃO (2015): "Trade and investment under policy uncertainty: theory and firm evidence," *American Economic Journal: Economic Policy*, 7, 189–222.
- (2017): "Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States," *American Economic Review*, 107, 2731–83.
- HELBLE, M., B. SHEPHERD, AND J. S. WILSON (2009): "Transparency and regional integration in the Asia Pacific," *World Economy*, 32, 479–508.
- HERGHELEGIU, C. (2018): "The political economy of non-tariff measures," *The World Economy*, 41, 262–286.
- HOLZER, K. (2019): "Addressing tensions and avoiding disputes: Specific trade concerns in the TBT Committee," *Global Trade and Customs Journal*, 14, 102–116.
- ING, L. Y., O. CADOT, AND J. WALZ (2018): "Transparency in non-tariff measures: An international comparison," *The World Economy*, 41, 884–912.
- IODICE, I. AND C. REVERDY (2025): "Protection or Protectionism: The Effect of Technical Regulations on Input Sourcing," .
- LEE, J.-W. AND P. SWAGEL (1997): "Trade barriers and trade flows across countries and industries," *Review of economics and statistics*, 79, 372–382.
- LEJÁRRAGA, I. AND B. SHEPHERD (2013): "Quantitative evidence on transparency in regional trade agreements," *OECD Publishing*.
- LIU, Q. AND H. MA (2020): "Trade policy uncertainty and innovation: Firm level evidence from China's WTO accession," *Journal of International Economics*, 127, 103387.
- MACEDONI, L. AND A. WEINBERGER (2022): "Quality heterogeneity and misallocation: The welfare benefits of raising your standards," *Journal of International Economics*, 134, 103544.
- MAYER, T., M. J. MELITZ, AND G. I. OTTAVIANO (2014): "Market size, competition, and the product mix of exporters," *The American Economic Review*, 104, 495–536.
- MEI, Y. (2024): "Regulatory Protection and the Role of International Cooperation," *International Economic Review*, 65, 817–850.
- MEI, Y. AND M. XU (2024): "Horizontal regulatory barriers in international trade: Evidence from electric plugs," *Available at SSRN 4281420*.
- OREFICE, G. (2017): "Non-Tariff Measures, Specific Trade Concerns and Tariff Reduction," *The World Economy*, 40, 1807–1835.

- PARENTI, M. AND G. VANNOORENBERGHE (2025): "A simple theory of deep trade integration," *Journal of the European Economic Association*, 23, 891–940.
- SANTOS SILVA, J. AND S. TENREYRO (2006): "The log of gravity," *The Review of Economics and statistics*, 88, 641–658.
- SCHMIDT, J. AND W. STEINGRESS (2022): "No double standards: quantifying the impact of standard harmonization on trade," *Journal of International Economics*, 103619.
- SINGH, R. AND R. CHANDA (2021): "Technical regulations, intermediate inputs, and performance of firms: Evidence from India," *Journal of International Economics*, 128, 103412.

Online Appendices

Online Appendices

A.1 Derivations for the Investment Problem

This appendix provides the algebra behind Section 4. We solve the infinite-horizon problem with survival/discount factor $\beta \in (0, 1)$ and then show how the two-period “snapshot” nests inside it.

Setup and Per-Period Profits

Consumers allocate expenditure between a numéraire and a CES aggregate. Demand for variety i is

$$q_i = E P^{\sigma-1} p_i^{-\sigma}, \quad \sigma > 1,$$

with E total expenditure and P the CES price index. Firms draw productivity φ_i and face monopolistic competition. A technical standard raises marginal cost by an ad valorem factor $d \geq 1$. With constant marginal costs, the optimal markup implies

$$p_i = \frac{\sigma}{\sigma - 1} \frac{d}{\varphi_i}.$$

Per-period operating profit is therefore

$$\pi(\varphi_i, d) = a d^{1-\sigma} \varphi_i^{\sigma-1}, \quad a \equiv E \sigma^{-\sigma} (\sigma - 1)^{\sigma-1} P^{\sigma-1}. \quad (7)$$

For compactness, define $\theta \equiv \sigma - 1 > 0$ and $c(d) \equiv a d^{1-\sigma}$, so $\pi(\varphi, d) = c(d) \varphi^\theta$.

Exporting requires a one-time sunk adaptation cost K_s paid at the time of entry. Conditional on survival/discount β , the present value of a *profit stream only* under constant d is

$$V(\varphi, d) \equiv \frac{\pi(\varphi, d)}{1 - \beta} = \frac{c(d)}{1 - \beta} \varphi^\theta. \quad (8)$$

(Note: I keep K_s *outside* of $V(\cdot)$ to avoid double counting.)

Timeline and Uncertainty

At $t = 0$ the standard is d_0 . From $t \geq 1$ it becomes either

$$d_h = d_0(1 + \delta) \quad \text{with prob. } \frac{1}{2}, \quad d_l = d_0(1 - \delta) \quad \text{with prob. } \frac{1}{2},$$

with $\delta \in [0, 1)$ measuring dispersion (lack of notification \Rightarrow larger δ). This is a mean-preserving spread in d . Because $x \mapsto x^{1-\sigma}$ is convex for $\sigma > 1$, Jensen's inequality implies $\mathbb{E}[d^{1-\sigma}]$ rises with dispersion.

Define

$$c_0 \equiv c(d_0), \quad c_h \equiv c(d_h) = c_0(1 + \delta)^{1-\sigma}, \quad c_l \equiv c(d_l) = c_0(1 - \delta)^{1-\sigma}, \quad \bar{c} \equiv \frac{1}{2}(c_h + c_l).$$

Note $c_l \geq c_h$ since $1 - \sigma < 0$.

Values: Invest Now vs. Wait One Period

In this appendix we define both values as present values at $t = 0$.

Invest immediately ($t = 0$). Pay K_s once at $t = 0$, earn period-0 profit under d_0 , and from $t \geq 1$ profit under the realised $d_j \in \{d_h, d_l\}$ forever:

$$\mathbb{E}_0 \Pi_0(\varphi) = c_0 \varphi^\theta + \frac{\beta}{1-\beta} \bar{c} \varphi^\theta - K_s = \left(c_0 + \frac{\beta}{1-\beta} \bar{c} \right) \varphi^\theta - K_s. \quad (9)$$

Wait until $t = 1$. At $t = 0$ earn nothing. At $t = 1$, after observing $d_j \in \{d_h, d_l\}$, invest iff $V(\varphi, d_j) - K_s \geq 0$. Discounting (survival back) to $t = 0$ yields

$$\mathbb{E}_0 \Pi_1(\varphi) = \beta \left[\frac{1}{2} \max \left\{ \frac{c_h}{1-\beta} \varphi^\theta - K_s, 0 \right\} + \frac{1}{2} \max \left\{ \frac{c_l}{1-\beta} \varphi^\theta - K_s, 0 \right\} \right]. \quad (10)$$

State-Specific Profitability Cutoffs if Waiting

Define the (wait-to- $t = 1$) state cutoffs by the zero-profit conditions inside the max operators:

$$\varphi_j^W \equiv \left(\frac{K_s(1-\beta)}{c_j} \right)^{1/\theta}, \quad j \in \{h, l\}, \quad \varphi_l^W \leq \varphi_h^W \text{ since } c_l \geq c_h. \quad (11)$$

Entry Cutoff φ^*

Let $F_0(\varphi) = \max\{\mathbb{E}_0 \Pi_0(\varphi), \mathbb{E}_0 \Pi_1(\varphi)\}$ and $\Omega_0(\varphi) = \max\{\mathbb{E}_0 \Pi_0(\varphi), 0\}$. The entry cutoff φ^* is the smallest φ such that $F_0(\varphi) = \Omega_0(\varphi)$; equivalently, the smallest solution to $\mathbb{E}_0 \Pi_0(\varphi) = \mathbb{E}_0 \Pi_1(\varphi)$, accounting for the max operators in (10). This yields three regions:

(H) High-productivity region: $\varphi \geq \varphi_h^W$ (both states would be profitable if you waited). Then $\mathbb{E}_0 \Pi_1(\varphi) = \beta \left(\frac{\bar{c}}{1-\beta} \varphi^\theta - K_s \right)$ and

$$\varphi^{*\theta} = \frac{(1-\beta)K_s}{c_0} \iff \varphi^* = \left(\frac{(1-\beta)K_s}{a d_0^{1-\sigma}} \right)^{1/\theta}. \quad (12)$$

In this region φ^* is *independent* of δ .

(M) Intermediate region: $\varphi_l^W \leq \varphi < \varphi_h^W$ (only the good/low-cost state would be profitable if you waited). Then

$$\mathbb{E}_0 \Pi_1(\varphi) = \beta \cdot \frac{1}{2} \left(\frac{c_l}{1-\beta} \varphi^\theta - K_s \right)$$

and

$$\varphi^{*\theta} = \frac{K_s (1 - \frac{1}{2}\beta)}{c_0 + \frac{\beta}{1-\beta} \cdot \frac{1}{2} c_h}. \quad (13)$$

Since $c_h = c_0(1 + \delta)^{1-\sigma}$ falls with δ (as $1 - \sigma < 0$), the denominator declines and φ^* increases with δ .

(L) Low-productivity region: $\varphi < \varphi_l^W$ (even the low-cost state would not justify entry if you waited). Then $\mathbb{E}_0 \Pi_1(\varphi) = 0$ and

$$\varphi^{*\theta} = \frac{K_s}{c_0 + \frac{\beta}{1-\beta} \bar{c}} = \frac{(1-\beta)K_s}{(1-\beta)c_0 + \beta \bar{c}}, \quad (14)$$

with $\bar{c} = \frac{1}{2}(c_h + c_l)$. Because \bar{c} rises with dispersion in d (convexity of $d^{1-\sigma}$), the denominator increases in δ , so φ^* falls with δ in this corner case.

Region validity. Each closed-form expression (12)–(14) is valid only if the resulting φ^* lies in the corresponding region (check against (11)). Economically relevant calibrations need mass on regions (H) and (M), which deliver the “uncertainty raises the option value of waiting” comparative static.

Option Value of Waiting and Comparative Statics

The real option at $t = 0$ is $F_0(\varphi) - \Omega_0(\varphi) = \max\{\mathbb{E}_0 \Pi_1(\varphi) - \mathbb{E}_0 \Pi_0(\varphi), 0\}$. Because $\mathbb{E}_0 \Pi_1$ contains max operators, the option value is inherently piecewise. In region (M), for example,

$$\mathbb{E}_0 \Pi_1(\varphi) - \mathbb{E}_0 \Pi_0(\varphi) = \beta \cdot \frac{1}{2} \left(\frac{c_l}{1-\beta} \varphi^\theta - K_s \right) - \left(c_0 + \frac{\beta}{1-\beta} \bar{c} \right) \varphi^\theta + K_s,$$

which increases with δ because c_h falls and \bar{c} rises as dispersion in d increases. At the cutoff φ^* (region by region), $\partial\varphi^*/\partial\delta > 0$ in (M) and $= 0$ in (H), matching the narrative in Section 4.

Proposition (Transparency lowers the entry cutoff). Holding the mean of d fixed, a mean-preserving increase in its dispersion δ makes the entry cutoff $\varphi^*(\delta)$ *weakly* increasing: it is unchanged in region (H) and *strictly* increasing in region (M). In region (L), φ^* decreases with δ (a corner case that is typically empirically negligible).

Two-Period ‘‘Snapshot’’ as a special case

The snapshot used in Section 4 truncates the horizon at $t = 1$. Obtain it from the infinite-horizon expressions by replacing the post- $t = 1$ value of a profit stream, $\frac{c_j}{1-\beta}\varphi^\theta$, with a single one-period payoff $c_j\varphi^\theta$. Equivalently, inside every continuation term ‘‘from $t \geq 1$ onward,’’ replace

$$\frac{c_j}{1-\beta} \longrightarrow c_j, \quad \varphi_j^W = \left(\frac{K_s(1-\beta)}{c_j}\right)^{1/\theta} \longrightarrow \hat{\varphi}_j^W = \left(\frac{K_s}{c_j}\right)^{1/\theta}, \quad j \in \{h, l\}.$$

Snapshot payoffs

$$\begin{aligned} \mathbb{E}_0 \Pi_0^{\text{snap}}(\varphi) &= c_0\varphi^\theta + \beta\bar{c}\varphi^\theta - K_s, & \bar{c} &= \frac{1}{2}(c_h + c_l), \\ \mathbb{E}_0 \Pi_1^{\text{snap}}(\varphi) &= \beta \left[\frac{1}{2} \max\{c_h\varphi^\theta - K_s, 0\} + \frac{1}{2} \max\{c_l\varphi^\theta - K_s, 0\} \right]. \end{aligned}$$

Snapshot state cutoffs (if the firm waits)

$$\hat{\varphi}_j^W = \left(\frac{K_s}{c_j}\right)^{1/\theta}, \quad j \in \{h, l\}, \quad \hat{\varphi}_l^W \leq \hat{\varphi}_h^W \text{ since } c_l \geq c_h.$$

Snapshot entry cutoff φ^* (piecewise). Solving $\mathbb{E}_0 \Pi_0^{\text{snap}}(\varphi) = \mathbb{E}_0 \Pi_1^{\text{snap}}(\varphi)$ yields:

(H) High productivity: $\varphi \geq \hat{\varphi}_h^W$.

$$\mathbb{E}_0 \Pi_1^{\text{snap}}(\varphi) = \beta(\bar{c}\varphi^\theta - K_s) \Rightarrow \varphi^{*\theta} = \frac{(1-\beta)K_s}{c_0} \iff \varphi^* = \left(\frac{(1-\beta)K_s}{a d_0^{1-\sigma}}\right)^{1/\theta}.$$

Here $\partial\varphi^*/\partial\delta = 0$.

(M) Intermediate: $\hat{\varphi}_l^W \leq \varphi < \hat{\varphi}_h^W$.

$$\mathbb{E}_0 \Pi_1^{\text{snap}}(\varphi) = \beta \cdot \frac{1}{2}(c_l\varphi^\theta - K_s) \Rightarrow \varphi^{*\theta} = \frac{K_s(1 - \frac{1}{2}\beta)}{c_0 + \frac{1}{2}\beta c_h}.$$

Since $c_h = c_0(1 + \delta)^{1-\sigma}$ decreases in δ , $\partial\varphi^*/\partial\delta > 0$.

(L) Low productivity: $\varphi < \hat{\varphi}_l^W$.

$$\mathbb{E}_0 \Pi_1^{\text{snap}}(\varphi) = 0 \Rightarrow \varphi^{*\theta} = \frac{K_s}{c_0 + \beta\bar{c}}.$$

Because \bar{c} increases with dispersion in d , $\partial\varphi^*/\partial\delta < 0$ in this corner case.

A.2 Construction of the TBT Timeline Dataset

A.2.1 Sources and Matching Process

The starting point is the WTO's Technical Barriers to Trade Information Management System (TBT IMS), which contains all official notifications of technical regulations submitted by WTO members. Each notification is identified by a unique Document Symbol, typically formatted as G/TBT/N/[CountryISO]/[Number] (e.g., G/TBT/N/TUR/142). Using this identifier, I matched each Specific Trade Concern (STC) to its corresponding regulatory documents.

I automated the scraping of notification documents from the WTO TBT IMS website using a combination of Python libraries, Requests to handle the HTTP communication and BeautifulSoup.

A.2.2 Date Extraction Procedures

Most notifications follow a standardized format, which allowed me to extract notification and proposed enforcement dates automatically using regular expressions. In cases where this information was missing or later revised, I parsed Revision documents. Addenda and Corrigenda, which typically covered translations or minor updates, were reviewed as needed but rarely altered the timeline.

Overall, about 80% of dates were extracted automatically, with the remainder verified manually due to missing fields, non-standard phrasing, or inconsistencies.

A.2.3 Committee Minutes Parsing and Timeline Classification

For cases where notifications were absent or incomplete, I turned to the WTO TBT Committee Minutes (“MIN” documents), which summarize member discussions. I parsed these texts using an algorithm that extracted all potential timestamps, which were then manually reviewed to identify relevant dates (e.g., “entered into force on...”, “will apply from...”).

If no specific dates were mentioned, I classified the measures into three categories based on the discussion content:

- **Y**: the measure was already in force.
- **NY**: the measure was still being drafted.
- **NA**: the information was inconclusive.

For measures classified as in force (**Y**), I assigned the entry-into-force date as the semester when the STC was first raised—reflecting the likely timing of the measure’s adoption.

When conflicting dates appeared across sources, I prioritized the Committee Minutes, followed by any revisions, and lastly the initially proposed dates in the Notification. This reflects the fact that implementation timelines often change without formal documentation, but are discussed in WTO meetings.

A.2.4 Recovery Rate and Sample Balance

This procedure allowed me to recover the notification and implementation timelines for 301 of the 403 contested TBTs—roughly 75% of the sample. Relying solely on notification documents would bias the sample toward transparent measures, but by incorporating information from the Committee Minutes, I mitigate this concern. Recovery rates are comparable for notified TBTs (77%) and unnotified ones (67%), suggesting balanced coverage of transparent and opaque cases.

A.2.5 Illustrative Example of a Notification Document

Figure 5 shows an example of a WTO Notification, with the notification and proposed enforcement dates highlighted.

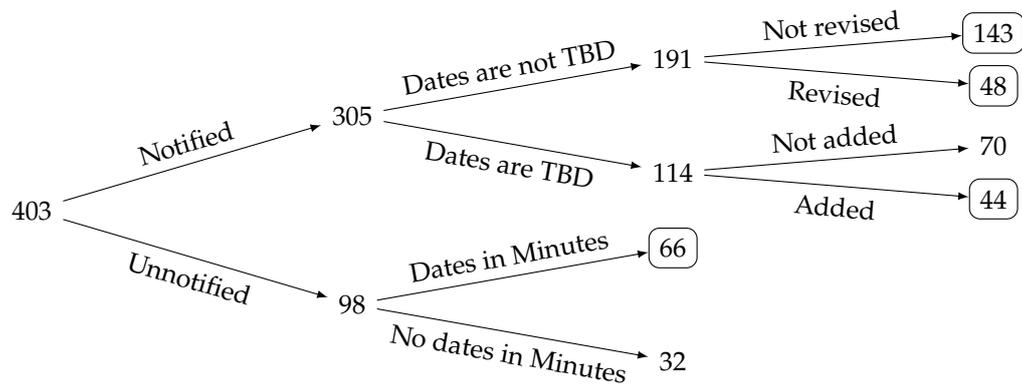
Figure 5: Example of a Notification Document

WORLD TRADE ORGANIZATION		G/TBT/N/ARG/101 23 May 2003 (03-2765)
Committee on Technical Barriers to Trade		Original: Spanish
NOTIFICATION		
The following notification is being circulated in accordance with Article 10.6.		
1.	Member to Agreement notifying: ARGENTINA If applicable, name of local government involved (Articles 3.2 and 7.2):	
2.	Agency responsible: National Institute of Vitiviniculture Name and address (including telephone and fax numbers and E-mail and Web site addresses, if available) of agency or authority designated to handle comments regarding the notification shall be indicated if different from above: <i>Idem</i> National Enquiry Point	
3.	Notified under Article 2.9.2 [X], 2.10.1 [], 5.6.2 [], 5.7.1 [], other:	
4.	Products covered (HS or CCCN where applicable, otherwise national tariff heading. ICS numbers may be provided in addition, where applicable): Wine	
5.	Title, number of pages and language(s) of the notified document: Wine – Sulphate Content (2 pages, in Spanish)	
6.	Description of content: Establishes the maximum limits for sulphate content, expressed as potassium sulphate, both in wine that is in circulation and in wineries.	
7.	Objective and rationale, including the nature of urgent problems where applicable: The need to establish, as an exporting country, the appropriate limits for these products through essential production and conservation techniques, as laid down by the International Organization of Vine and Wine (OIV).	
8.	Relevant documents: INV Resolution No. 14/2003	
9.	Proposed date of adoption: 30 April 2003 (Official Journal) Proposed date of entry into force: 8 May 2003	
10.	Final date for comments: -	
11.	Texts available from: National enquiry point [X], or address, telephone and fax numbers and E-mail and Web site addresses, if available, of other body: Punto Focal de la República Argentina Dirección Nacional de Comercio Interior (DNCI) Avda. J. A. Roca 651, Piso 4º, Sector 22 (1322) Buenos Aires Fax: 54 11 4349 4072 Tel.: 54 11 4349 4067 E-mail: focalote@meccon.gov.ar Web site: http://www.puntofocal.gov.ar	

A.2.6 Data Sources Summary: Tree Diagram

Figure 6 displays a tree summarizing the sources used to recover timelines for each measure.

Figure 6: Data sources for the database on TBT timelines



Notes: The frequency of TBTs is displayed in a tree format, organized by the source of information used to retrieve the timelines of TBTs. The edges of the tree indicate the attributes that identify the information source. Final nodes with rectangular frames indicate the cases where the timeline could be identified.

A.3 Additional Tables

A.3.1 Descriptive Statistics

Table 9: Country Coverage of TBTs and Export Growth Rates

Country	# TBT	% w/o Timely Notif.	Avg. Annual Export Growth (%)		
			All HS4	With TBT	w/o Timely Notif.
EU	1184	50.7	2.10	2.30	2.20
CN	848	46.7	6.60	5.80	5.20
MX	489	65	2.60	1.40	1.50
KR	473	20.3	2.40	3.10	2.90
US	440	46.1	1.50	1.70	1.70
IN	432	47.2	6.20	6.40	6.60
ID	299	94.3	4.50	2.40	4.10
BR	265	64.2	3.20	5.50	5.50
TW	214	78.5	N/A	N/A	N/A
JP	191	1.6	1.20	1.30	2.50
ZA	177	2.3	4	2.80	6.40
CA	174	1.7	N/A	N/A	N/A
AR	172	5.2	1.90	3.10	7
EG	171	100	5.70	1.80	1.80
NZ	167	100	2.70	2.80	2.80
MY	31	41.9	3.50	3.10	N/A
CO	28	7.1	3.50	2.70	N/A
TR	17	100	N/A	N/A	N/A
EC	10	60	3.60	2.90	N/A
IL	8	100	3.10	1.60	1.60
PE	7	14.3	5.10	6.60	6.70
KE	6	100	N/A	N/A	N/A
MD	2	100	N/A	N/A	N/A
TH	1	100	4.40	5.20	5.20
UY	1	100	-0.40	3.50	3.50
Average	228	57.3	3.40	3.30	4

Note: A product refers to an HS4 code. A TBT is a regulatory measure affecting an HS4 code. “% Without Timely Notification” refers to the share of TBTs that were either enforced before notification or never notified. Missing growth rates reflect discontinuous trade series for some country-product pairs.

Table 10: Testing for pre-trends

DEPENDENT VARIABLE	Participation:		
	(1)	(2)	(3)
TBT	-0.074*** (0.020)	-0.075*** (0.020)	-0.076*** (0.020)
TBT × Notified	0.042** (0.021)	0.043** (0.021)	0.043** (0.021)
1{TBTtime = -1}	-0.020 (0.013)	-0.021 (0.013)	-0.021 (0.013)
1{TBTtime = -2}		-0.011 (0.012)	-0.012 (0.012)
1{TBTtime = -3}			-0.007 (0.015)
Observations	8 155 082	8 155 082	8 155 082
R ²	0.247	0.247	0.247
HS2-Country-Time FE	Yes	Yes	Yes
Firm-HS4-Country FE	Yes	Yes	Yes

Notes: Participation equals 1 if a firm exports a given HS-4 product to a destination in a semester. Standard errors clustered at (HS4, country, time). Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Effects of TBTs, controlling for product elasticity

DEPENDENT VARIABLE	Participation			
	(1)	(2)	(3)	(4)
TBT	-0.027** (0.012)	-0.060*** (0.021)	-0.029*** (0.010)	-0.067*** (0.021)
TBT × Elasticity (BW)	-0.003 (0.002)	-0.003 (0.002)		
TBT × Notified		0.043** (0.021)		0.049** (0.021)
TBT × Elasticity (FGO)			-0.002 (0.001)	-0.002 (0.001)
Observations	8,117,694	8,117,694	8,092,448	8,092,448
R ²	0.247	0.247	0.247	0.247
HS2-Country-Time FE	Yes	Yes	Yes	Yes
Firm-HS4-Country FE	Yes	Yes	Yes	Yes

Notes: The dependent variable is a binary indicator equal to one if a firm exports a positive amount of a product to a destination in a given year. Elasticity (BW) and Elasticity (FGO) refer to product-level demand elasticity estimates from [Broda and Weinstein \(2006\)](#) and [Fontagné et al. \(2022\)](#), respectively. Columns (1) and (3) report the baseline effects of TBTs, including interactions with product elasticity measures to control for heterogeneous trade responses. Columns (2) and (4) add an interaction with a dummy for whether the TBT was notified in advance. All regressions include HS2-country-year and firm-HS4-country fixed effects. Standard errors, clustered at the HS4-country-year level, are reported in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 12: Effects of TBTs, controlling for tariffs

DEPENDENT VARIABLE	Participation		Export (0)		Log Exports	
	(1)	(2)	(3)	(4)	(5)	(6)
log(1+tariff%)	0.031*** (0.006)	0.031*** (0.006)	-0.121** (0.061)	-0.104* (0.058)	-0.047** (0.019)	-0.047** (0.019)
TBT	-0.039*** (0.008)	-0.073*** (0.020)	-0.396*** (0.076)	-0.761*** (0.228)	0.015 (0.035)	0.001 (0.090)
TBT × Notified		0.042** (0.021)		0.433* (0.239)		0.017 (0.096)
Observations	7,615,002	7,615,002	7,550,366	7,550,366	4,196,684	4,196,684
R2	0.248	0.248			0.745	0.745
R2 adj.	0.201	0.201			0.682	0.682
HS2-Country-Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-HS4-Country FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The dependent variables are: participation (binary export indicator), export values including zeros (*Export (0)*), and log exports excluding zeros (*Log Exports*). Columns (3) and (4) are estimated on the full sample (export ≥ 0) using PPML; columns (5) and (6) are estimated on strictly positive exports (export > 0) using OLS. Tariff refers to the log of applied tariffs at the product-destination level. All regressions include HS2-country-time and firm-HS4-country fixed effects. Standard errors, clustered at the HS4-country-time level, are reported in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 13: PPML with HDFE

Dependent variable:	Exports ≥ 0 (PPML)		Exports > 0 (PPML)	
	(1)	(2)	(3)	(4)
TBT	-0.396*** (0.076)	-0.776*** (0.228)	-0.148** (0.067)	-0.263* (0.156)
TBT × Notified		+0.450* (0.238)		+0.136 (0.168)
Observations	8,082,903	8,082,903	3,316,077	3,316,077
HS2-Country-Time FE	Yes	Yes	Yes	Yes
Firm-HS4-Country FE	Yes	Yes	Yes	Yes

The table reports results from Poisson Pseudo-Maximum Likelihood (PPML) regressions with high-dimensional fixed effects, estimated using the `pmlhdfe` routine of [Correia et al. \(2019\)](#). The dependent variable is export value at the firm-product-destination-semester level. Columns (1)–(2) include all observations, including zeros; columns (3)–(4) restrict the sample to strictly positive export flows. All regressions include HS2-country-semester and firm-HS4-country fixed effects. Standard errors, clustered at the HS4-country-semester level, are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.3.2 Dynamic Effects of Notification on Export Participation

The table below reports the estimates corresponding to Figure 4. These estimates are obtained from the following specification:

$$y_{pdt} = \sum_{n=-A}^N (\alpha_n \text{NotifiedTBT}_{pd} + \beta_n \text{UnnotifiedTBT}_{pd}) \mathbb{1}\{\text{TBTtime}_{pdt} = n\} + \mu_{ipd} + \mu_{HS2,dt} + \epsilon_{pdt} \quad (15)$$

where NotifiedTBT_{pd} is defined as $\text{Notified}_{pd} \times \text{TBT}$, while The model builds on the identification strategy of the baseline equation and includes fixed effects μ_{ipd} and $\mu_{HS2,dt}$, thereby leveraging variation over time within firm–product–destination markets.

Table 14: Coefficients plotted in Figure 4 (dynamic specification)

Dependent variable: Participation	Estimate	Std. Error
$1\{\text{TBTtime}_{pdt} = -3\} \times \text{Notified}$	-0.001	(0.016)
$1\{\text{TBTtime}_{pdt} = -3\} \times \text{Unnotified}$	-0.025	(0.034)
$1\{\text{TBTtime}_{pdt} = -2\} \times \text{Notified}$	-0.019	(0.013)
$1\{\text{TBTtime}_{pdt} = -2\} \times \text{Unnotified}$	0.017	(0.032)
$1\{\text{TBTtime}_{pdt} = -1\} \times \text{Notified}$	-0.015	(0.014)
$1\{\text{TBTtime}_{pdt} = -1\} \times \text{Unnotified}$	-0.050	(0.033)
$1\{\text{TBTtime}_{pdt} = 0\} \times \text{Notified}$	-0.065***	(0.015)
$1\{\text{TBTtime}_{pdt} = 0\} \times \text{Unnotified}$	-0.089***	(0.033)
$1\{\text{TBTtime}_{pdt} = 1\} \times \text{Notified}$	0.001	(0.021)
$1\{\text{TBTtime}_{pdt} = 1\} \times \text{Unnotified}$	-0.068**	(0.027)
$1\{\text{TBTtime}_{pdt} = 2\} \times \text{Notified}$	0.010	(0.023)
$1\{\text{TBTtime}_{pdt} = 2\} \times \text{Unnotified}$	0.015	(0.024)
$1\{\text{TBTtime}_{pdt} \geq 3\} \times \text{Notified}$	-0.018*	(0.010)
$1\{\text{TBTtime}_{pdt} \geq 3\} \times \text{Unnotified}$	-0.044**	(0.018)
N		8 155 082
R^2		0.771
HS2–Country–Time FE		Yes
Firm–HS4–Country FE		Yes

Notes: Participation equals 1 if the firm records a positive export for product p to destination d in semester t . Robust standard errors clustered by (HS4, country, time) in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.4 Tariff-Equivalent Mapping: Derivation and Sensitivity

CES revenue mapping

Extending the theoretical framework, export revenues for firm i 's variety p to destination d at time t follow a CES structure in which both TBT compliance costs (d) and tariffs (τ) act as multiplicative iceberg barriers:

$$Y_{ipdt} = a_{pdt} \varphi_i^{\sigma-1} d_{pdt}^{1-\sigma} (1 + \tau_{pdt})^{1-\sigma}, \quad \sigma > 1, \quad (16)$$

where a_{pdt} captures destination–product demand shifters, φ_i is firm productivity, $d_{pdt} \geq 1$ is the ad valorem cost multiplier associated with TBTs, and τ_{pdt} is the ad valorem tariff rate (e.g., $\tau = 0.13$ for a 13% tariff). Under CES demand with constant markups, revenues scale with trade costs in the same way as profits, so the tariff-equivalent interpretation applies.

PPML specification and interpretation of $\hat{\beta}$

To account for both extensive and intensive margins, the empirical model is estimated by PPML in levels with high-dimensional fixed effects:

$$\mathbb{E}[Y_{ipdt} | \cdot] = \exp\left(\alpha \text{TBT}_{pdt} + \beta (\text{TBT}_{pdt} \times \text{Notified}_{pd}) + \mu_{HS2,dt} + \mu_{ipd}\right). \quad (17)$$

PPML retains zero flows and is consistent under heteroskedasticity, accommodating firm–product–destination triplets that switch in and out of exporting; thus coefficients reflect adjustments on both intensive (values) and extensive (participation) margins. The fixed effects μ_{ipd} absorb time-invariant firm–product–destination heterogeneity (including $\varphi_i^{\sigma-1}$ and match quality), while $\mu_{HS2,dt}$ absorb a_{pdt} and co-moving product–destination–time factors. The interaction coefficient $\hat{\beta}$ measures the revenue effect associated with receiving advance notification when a TBT is imposed (see Table 13; column (2) yields $\hat{\beta} = 0.450$).

Deriving the tariff-equivalent change

Let τ_0 denote the baseline tariff absent timely notification. Under (16), equating revenues without and with timely notification implies

$$(1 + \tau_0)^{1-\sigma} = \exp(\hat{\beta}) (1 + \tau_0 + \Delta\tau)^{1-\sigma}. \quad (18)$$

Solving (18) for the compensating ad valorem tariff increment $\Delta\tau$:

$$\Delta\tau = (1 + \tau_0) \left[\exp(\hat{\beta}) \right]^{1/(\sigma-1)} - (1 + \tau_0). \quad (19)$$

Expressing tariffs in percentage points $t = 100 \tau$ gives the tariff-equivalent change in percentage points:

$$\Delta t = (100 + t_0) \left(\exp\left(\frac{\hat{\beta}}{\sigma-1}\right) - 1 \right), \quad (20)$$

where $t_0 = 100 \tau_0$.

Sensitivity to σ and baseline tariffs

Table 15 reports implied tariff equivalents for benchmark elasticities $\sigma \in [3, 7]$ and baseline tariffs set to the median extra-EU rate (5.5%) and the median in TBT-affected markets (10.6%), using $\hat{\beta} = 0.450$.

The mapping in (20) assumes CES demand and constant markups so that revenues scale with trade costs as in (16). If markups vary with trade costs or if demand departs from CES, Δt remains an informative local translation but its level will depend on the specific alternative structure. The fixed effects in (17) absorb time-invariant productivity and

Table 15: Tariff Equivalent of Timely Notification under Alternative σ

Elasticity σ	Δt if $t_0 = 5.5\%$	Δt if $t_0 = 10.6\%$
3	26.6 pp	27.8 pp
4	17.2 pp	17.9 pp
5	12.7 pp	13.2 pp
7	8.3 pp	8.7 pp

Notes: Values computed from (20) with $\hat{\beta} = 0.450$. “pp” denotes percentage points. Magnitudes are illustrative and rely on CES demand with constant markups; they vary with σ and the baseline tariff.

match heterogeneity and control for product–destination–time demand shocks; the identifying variation for $\hat{\beta}$ comes from within–HS2–destination–time differences across products in the incidence and timing of notification conditional on TBT enforcement.