The Effect of Export Controls: Evidence from Japanese Customs Data*

Daisuke Adachi† Rinki Ito‡ Daiji Kawaguchi§ Yukiko U. Saito¶

September 3, 2026

Abstract

Export controls restrict trade in dual-use goods through technical thresholds that are revised almost every year, yet little is known about who bears the compliance burden. Using the universe of Japan's export declarations from 2014 to 2023, in which every licensed transaction carries a control-list code, we translate each revision into product-level treatments and estimate tightening and relaxation as separate arms in a stacked event-study design. The answer is asymmetric. Tightening produces no detectable adjustment among incumbent exporters, while relaxation activates exports—a benefit accruing almost entirely to small firms, whose shipments become more frequent while shipment size is roughly unchanged. The evidence points to compliance costs operating as per-shipment and fixed costs: sunk for incumbents, binding for small firms, and regressive in firm size—but visible only when the cost is removed, not when it is imposed.

JEL Classification: F13, F14, F52

Keywords: export controls, international trade, economic security, customs data

*This study is the result of joint research with the Policy Research Institute (PRI), which submitted an offer of use to the Ministry of Finance (MOF) based on the "Guideline on the utilization of Customs' import and export declaration data in a joint research with PRI," and received approval in February 2024. The views expressed in this research are those of the author's responsibility and do not represent the official views of the PRI or MOF.

†Aarhus University and RIETI

‡The Policy Research Institute, Ministry of Finance, Japan

§University of Tokyo, RIETI, and IZA

¶Waseda University and RIETI

1 Introduction

Export controls have moved to the center of economic statecraft—the pursuit of security objectives through economic means (Clayton et al., 2026; Mohr and Trebesch, 2025). The U.S. restrictions on exports of advanced semiconductors to China, and China's countervailing controls on critical minerals, are only the most visible instances of a rapidly expanding practice at the front line of great-power rivalry. Yet these dramatic, targeted measures are the tip of the iceberg of what export controls actually are.

The body of the institution is a permanent system of list-based controls, operated by the major supplier countries under multilateral coordination to prevent the proliferation of weapons of mass destruction and other sensitive military and dual-use technologies. What is controlled is delineated not by product categories but by technical specification thresholds, and these thresholds are revised almost every year as technology advances—tightened where new risks emerge, and relaxed where controls have become obsolete. The system delivers a security benefit, but it does so by imposing compliance costs on exporting firms: classification, licensing, and the maintenance of internal control systems. The question this paper asks is who bears this burden when the boundary of regulation moves. Which firms carry the cost of a tightening? Which firms capture the benefit of a relaxation? This distributional question is the counterpart to the security benefit and is of first-order importance for the design of the system, yet firm-level evidence is almost entirely absent.

The main reason is measurement. The product classification of export controls is designed independently of the HS classification used in trade statistics, and control status is determined by technical performance, so that controlled and non-controlled goods coexist within the same HS code. Semiconductors are the canonical case: the HS classification contains no subdivisions by performance, yet only high-performance devices are controlled. As long as one relies on trade statistics, therefore, it is impossible to determine which transactions are subject to control, and existing studies have been forced to assign treatment approximately through HS codes.[1] Under this adjacent-item problem, estimates mix con-

[1] Existing studies have coped with this constraint in three ways. First, when a regulatory episode targets a handful of narrowly defined items, the corresponding trade classification codes are enumerated directly, as in Makioka and Zhang (2024) for the three items of the 2019 Japan–Korea episode and Hayakawa et al. (2025) for the semiconductor manufacturing equipment controlled in 2023; the latter note that the matched codes

trolled with non-controlled goods and are attenuated—and comparing the incidence of the burden across firms is out of reach. This paper attacks the measurement problem at its source. We use the universe of Japan's export declaration data from 2014 to 2023. The decisive feature of these data is that every transaction exported under a license from the Ministry of Economy, Trade and Industry carries a five-digit code corresponding to the item of the control list (Appended Table 1 of the Export Trade Control Order), so that we can identify without approximation which transactions are controlled and which controlled item they fall under. This does not make the measurement problem disappear. Revisions are made to the provisions of the control codes, so the ministerial ordinance identifies only which codes an amendment pertains to; which of the HS products linked to those codes are actually affected still cannot be determined. We use sharp changes—jumps—in the number of controlled transactions around each revision to select, from the data, the HS products most likely affected. This mitigates rather than resolves the adjacent-item problem: it reduces the measurement error of treatment relative to the existing literature. Linking exporters to the Tokyo Shoko Research firm database through corporate numbers, we observe firm attributes such as size, which provides the basis for comparing the incidence of the burden across firms. Our analysis covers the revisions of the technical specifications of Japan's list-based controls—the controls under Appended Table 1 of the Export Trade Control Order. Japan participates in all four multilateral export control regimes and implements their agreements domestically by amending the ministerial ordinance that defines the technical specifications. Our sample consists of the revisions made through this channel between 2014 and 2023; country-specific measures such as sanctions and the end-use-based catch-all control are outside its scope. This setting has three advantages. First, because the revisions implement multilateral agreements, their content and timing are determined independently of individual

inevitably contain "neighboring products" outside the scope of the controls. Second, Hayakawa et al. (2023) construct treatment from the flag in the Export Statistical Schedule published by Japan Customs, which indicates that a nine-digit statistical code contains items potentially subject to the Export Trade Control Order—a property of the code rather than of the transaction, attached to roughly 45 percent of nine-digit codes accounting for about 70 percent of Japan's export value. Third, for the U.S. controls, Deseatnicov et al. (2024) construct a correspondence between the control-list classification (ECCN) and HS ten-digit codes, combining manual classification, an official German ECCN–HS correspondence, and machine-learning matching of control-list descriptions to product descriptions. All three approaches identify, at best, whether a product classification may contain controlled items; none observes whether a given transaction actually required a license.

Japanese firms' export performance. Second, the controls name no destination or firm and apply worldwide on the basis of technical specifications, so the effects we identify are those of the regulation itself rather than of a rerouting of destinations. Third, the specifications were revised nine times over the period, in both directions. We observe the routine operation of the system—not a single tightening episode—for both tightening and relaxation.

The empirical analysis proceeds in two steps. In the first step, we translate each revision into a treatment in the trade data: since the ministerial ordinance identifies only the control codes an amendment pertains to, we use the jumps in controlled transactions around each revision to identify the affected HS6 products. In the second step, we estimate a stacked event-study design at the firm $\times$ HS6 $\times$ quarter level, comparing the treated HS6 codes with control HS6 codes that belong to the same HS2 chapter and have a record of controlled transactions. Three devices secure identification: a between-firm comparison that excludes the other products of treated firms from the control group, a restriction to incumbents active before the estimation window, and a revision-by-revision pre-trend test that selects the revisions entering the main estimation.

Three findings emerge. First, tightening does not push incumbents out. Transaction counts are precisely estimated zeros, export values show no systematic decline, and treated incumbents are, if anything, 2 to 3 percentage points more likely to remain active than their controls—a relative persistence, not the creation of new exporters. Second, relaxation activates. Incumbents in relaxed products become 2 to 6 percentage points more likely to remain active, a differential that widens over the post-enforcement years, and their exports expand on the frequency margin. Third, the response to relaxation is concentrated among small firms: their transaction counts rise by 9 to 16 percent and their participation by 4 to 5 percentage points, while large firms show no intensive-margin response in either arm. The expansion occurs in the number of shipments with shipment size roughly unchanged—precisely the margin on which a per-transaction cost operates. These results give an asymmetric answer to the distributional question. The cost of a tightening does not fall on incumbent exporters: for firms that have already built compliance capability, the fixed cost is sunk, and bringing an additional product under control adds little marginal burden—the estimates are even consistent with the sunk cost sheltering incumbents from prospective entrants, although this cannot be separated from demand shifts in our incumbent panel. The benefit of a relaxation, by contrast, accrues to small firms, for whom per-shipment licensing costs bind; once those costs are removed, they ship more frequently and stay in the market longer. The burden of list-based controls is thus regressive in firm size—but its incidence becomes visible not when the cost is imposed, on incumbents for whom it is already sunk, but when it is removed. This asymmetry is invisible by construction in studies of single tightening episodes, and observing both directions of revision is what allows us to detect it.

Related literature. Empirical work on export controls is growing quickly, yet the question we address—how revisions to the technical thresholds that define control lists affect exporting firms—has remained essentially unstudied. The studies that come closest examine Japan's export controls using trade statistics: Hayakawa et al. (2023) detect little effect of Japan's own tightening on dual-use exports, Makioka and Zhang (2024) find that the 2019 measures toward South Korea reduced exports sharply only for hydrogen fluoride, and Hayakawa et al. (2025) document mixed, item-dependent effects of the 2023 controls on semiconductor manufacturing equipment. Because trade statistics are coarser than the control lists, however, these studies must assign treatment through product classifications, pooling controlled and uncontrolled goods within the same product code—an adjacent-item problem that Hayakawa et al. (2025) themselves emphasize—and each analyzes a single, salient episode. We mitigate the measurement problem where it arises: the license codes attached to each transaction identify controlled trade exactly, and our jump procedure narrows each revision down to the products most likely affected. This allows us to study the routine operation of the system—nine revisions over a decade—and to estimate tightening and relaxation as separate arms.

Most of the existing evidence on export controls instead concerns entity-based measures, above all the U.S. restricted lists. Crosignani et al. (2026) document broad supply-chain decoupling and sizable costs borne by U.S. suppliers following Entity List additions; Liu et al. (2024) exploit the item-level design of the 2007 China Rule, the closest methodological analog to our product-level treatment; and Hayakawa (2026) compares the trade effects of the U.S. regulations across major supplier countries. A related set of papers traces the effects of U.S. controls on third countries, and on Japanese firms in particular (Ando et al., 2024; Deseatnicov and Fukao, 2024; Hayakawa and Ito, 2025; Jinji et al., 2025). These measures target named firms or countries in exceptional, geopolitically charged circumstances; the list-based controls we study are applied worldwide, defined by technical specifications, and revised in both directions through multilateral agreements. Our estimates therefore speak to the ordinary machinery that every participating country operates continuously, rather than to exceptional confrontations.

Beyond export controls proper, our paper relates to work on restrictions that exporting countries place on strategic goods. Alfaro et al. (2025) show that China's 2010 rare-earth restrictions redirected innovation and reallocated supply in downstream industries worldwide, and Kee and Xie (2025) study Indonesia's nickel export ban; both trace the consequences of such measures for foreign, downstream economies. A larger literature on economic sanctions likewise documents trade destruction and deflection, together with the costs borne by firms in the sanctioning country (Felbermayr et al., 2020; Crozet and Hinz, 2020; Crozet et al., 2021). Sanctions and resource embargoes, however, are extraordinary and country-specific; export controls are a permanent, technical layer of regulation applied to all destinations. We complement this work by measuring the domestic incidence of that standing regime on the regulating country's own exporters. Finally, our interpretation builds on the literature on fixed and sunk costs of exporting (Baldwin and Krugman, 1989; Dixit, 1989; Roberts and Tybout, 1997; Melitz, 2003; Das et al., 2007). Both our descriptive evidence—fewer and larger shipments of controlled goods, and size-dependent selection into controlled varieties—and our treatment of tightening and relaxation as asymmetric arms follow from viewing licensing as a per-shipment and fixed compliance cost with hysteresis. Methodologically, we follow the stacked event-study approach (Gormley and Matsa, 2011; Cengiz et al., 2019; Baker et al., 2022).

The remainder of the paper is organized as follows. Section 2 describes the institutional background of export controls and their implementation in Japan. Section 3 introduces the data, documents basic patterns of controlled-goods trade, and describes the revisions we study. Section 4 identifies the products affected by each revision and sets out the empirical specification, and Section 5 presents the estimation results. The last section concludes.

2 Overview of export controls

2.1 Purpose of export controls

Export controls are a regulatory instrument designed to prevent goods and technologies that could contribute to the development of weapons of mass destruction (WMD) or the excessive accumulation of conventional weapons from reaching states or non-state actors of proliferation concern. The core challenge that export controls address is the dual-use problem: many items with entirely legitimate civilian applications—machine tools, sensors, advanced materials, chemicals, or semiconductor manufacturing equipment—can also serve as inputs into weapons programs. Because the military potential of an item is typically determined not by its product category but by its technical specifications (e.g., precision, purity, or performance thresholds), export controls operate by defining technical specification thresholds: items exceeding a given threshold require government screening and a license before they can be exported, while items below it can be exported freely. Export controls thus internalize a negative security externality of trade, but they do so by imposing compliance costs—classification, license applications, and internal compliance systems—on exporting firms. This trade-off between the security benefit and the private cost borne by firms is the central motivation of this paper.

However, export controls in recent years have come to serve not only these traditional security objectives but increasingly a strategic function as an instrument for protecting domestic industrial competitiveness and for exerting diplomatic and geopolitical pressure. The rivalry between the United States and China is emblematic. Motivated not only by concerns over military diversion but also by a desire to preserve its technological edge, the United States has imposed sweeping restrictions on exports of advanced semiconductors and semiconductor manufacturing equipment to specific countries—China in particular. In addition, through the Entity List, it designates particular foreign firms by name and requires an individual license for exports to them, effectively cutting off transactions (Crosignani et al., 2026). Unlike controls based on uniform technical thresholds applied worldwide, these measures are directed at specific destination countries or specific actors, and are therefore different in character. China, for its part, has tightened its export controls on critical minerals such as rare earths, rare metals, and gallium and germanium—a posture that functions as a countermeasure and a diplomatic lever, exploiting resources over whose supply China holds a chokehold (Alfaro et al., 2025). Export controls have thus increasingly come to operate as instruments of industrial policy and economic security, beyond their original framework of multilateral coordination for security purposes (Clayton et al., 2026; Mohr and Trebesch, 2025). These unilateral and strategic measures, however, are distinct—in both their objectives and their decision-making processes—from the controls this paper examines. What we analyze are amendments to Japan's list control: worldwide, specification-based controls that domestically implement the agreements reached in the four multilateral regimes. Controls of this kind are not designed to protect particular countries' industries or to target individual firms.[2] These developments nonetheless illustrate how the institution of export controls as a whole is broadening beyond security objectives toward economic and industrial ones, and thus form part of the backdrop against which we situate the traditional, security-oriented controls that this paper studies.

2.2 International export control regimes

A single country's export controls are easily undermined if the same item can be sourced from a less regulated competitor country; unilateral controls mainly divert procurement channels rather than block them. Preventing such circumvention requires that all major supplier countries control the same items at the same technical thresholds. This coordination is achieved through international export control regimes, in which participating states agree on common control lists and share information on proliferation risks. Figure 1 summarizes the architecture. The international layer consists of two complementary components.

[2] Japan itself has once imposed an export-control measure with a political dimension. In July 2019, citing a deterioration of trust in bilateral export-control cooperation, Japan revised its treatment of exports to South Korea of three materials essential to semiconductor manufacturing—fluorinated polyimide, photoresists, and hydrogen fluoride. Rather than changing any technical threshold, the measure altered the licensing arrangement: these three items were removed from eligibility for bulk licenses, so that each shipment to South Korea came to require an individual license (on the distinction between individual and bulk licenses, see Section 2.3). The three items were already listed items, and the measure applied only to goods meeting the specifications defined under the multilateral regimes—not to the product categories in their entirety. Separately, South Korea was also removed from the group of trusted destinations (the former "white country" group). Because this episode is a country-specific measure of exactly the kind we exclude from our analysis, it does not affect the amendments we study.

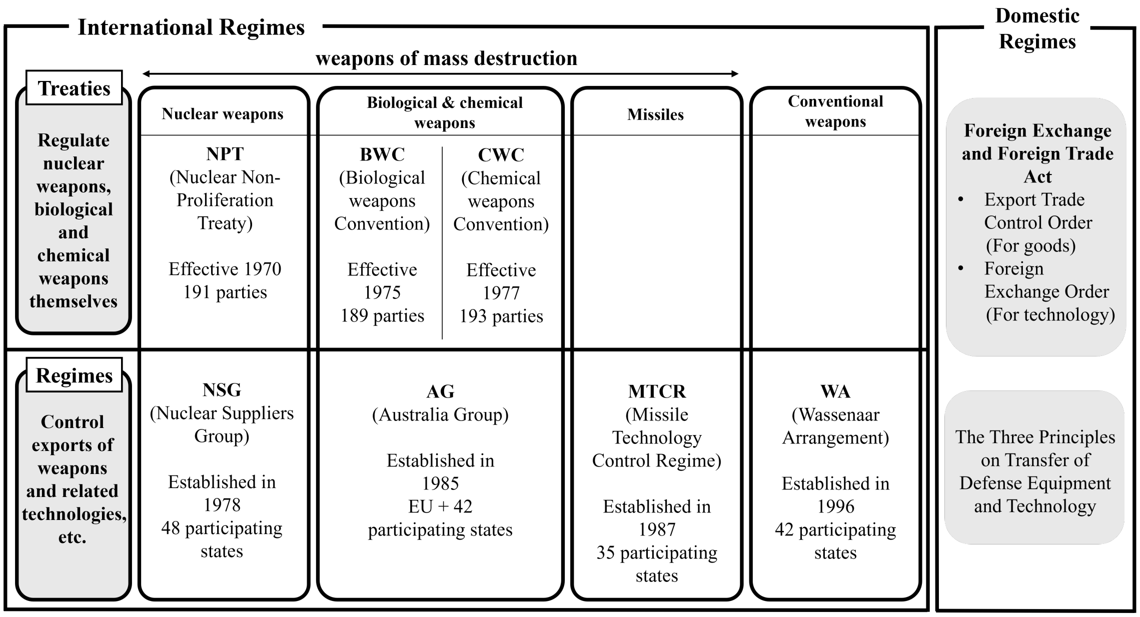

Figure 1: Overview of international export control regimes (as of July 2025)

Source: Security Export Inspection Office, Ministry of Economy, Trade and Industry.

Treaties—the NPT for nuclear weapons and the BWC and CWC for biological and chemical weapons—prohibit or regulate the weapons themselves and are near-universal in membership. The export control regimes, by contrast, are informal arrangements among supplier countries that control trade in the goods and technologies related to those weapons. There are four such regimes, each corresponding to a weapons category: the Nuclear Suppliers Group (NSG, established 1978, 48 participating states) for nuclear-related items; the Australia Group (AG, 1985, EU plus 42 states) for biological and chemical weapons-related items; the Missile Technology Control Regime (MTCR, 1987, 35 states) for missiles and their delivery systems; and the Wassenaar Arrangement (WA, 1996, 42 states) for conventional weapons and related dual-use goods and technologies. The regimes are not legally binding treaties: their agreed control lists acquire legal force only when each participating state transposes them into its domestic law. The lists are reviewed and updated regularly—typically through annual plenary meetings—to keep pace with technological change, adding items or tightening thresholds where new proliferation risks emerge, and deleting items or relaxing thresholds where controls have become obsolete or where items have become widely available outside the regimes.

2.3 Implementation in Japan

Japan participates in all four regimes and implements their agreements through the Foreign Exchange and Foreign Trade Act (FEFTA), as shown in the right panel of Figure 1. FEFTA delegates the operational details to two cabinet orders, corresponding to the two forms in which controlled items can cross borders: the Export Trade Control Order governs tangible goods, and the Foreign Exchange Order governs intangible technology transfers (e.g., technical data, software, or technical assistance).

Appended Table 1 of the Export Trade Control Order enumerates controlled goods on an item-by-item basis ("list control"); its full structure is shown in Figure 2. Items 1 through 15 of Appended Table 1 constitute the list control and can be organized into three broad layers. Item 1 covers arms themselves (military items) and corresponds to the Wassenaar Arrangement Munitions List (WA/ML). Items 2 through 4 cover WMD-related dual-use goods, each mapped to a specific regime: nuclear-related items (item 2, NSG), chemical weapons-related items (item 3, AG), biological weapons-related items (item 3-2, AG), and missile-related items (item 4, MTCR). Items 5 through 15 cover conventional arms-related dual-use goods—advanced materials, material processing, electronics, computers, communications and information security, sensors and lasers, navigation and avionics, marine equipment, and aerospace and propulsion, among others—items 5 through 13 correspond to categories 1 through 9 of the Wassenaar Arrangement dual-use list, while items 14 and 15 cover other munitions-related and particularly sensitive items. The structure of Appended Table 1 thus transposes the coverage of the four regimes almost directly into domestic law. Within items 1 through 15, each item is further subdivided into more specific sub-items defined by product type; as of June 2026, there are 240 such sub-items in total.[3] The list control is complemented by a catch-all control, shown as item 16 in Figure 2: it requires a license for non-listed items when the exporter knows, or is informed by the Ministry of Economy, Trade and Industry (METI), that the end use is WMD-related. Because the catch-all control is destination- and end-use-based, it lies outside the scope of this paper.

Crucially, each item of Appended Table 1 specifies only a product category; the technical specification thresholds that determine whether a particular good falls under control

[3] The number of sub-items varies slightly from year to year, but there are no substantial changes.

are defined in a METI ministerial ordinance (the "Ministerial Order Specifying Goods and Technologies," hereafter the Ministerial Order). An exporter must therefore assess, for each shipment, whether the good meets the technical specifications in the Ministerial Order, and if so, obtain an export license from METI. List control under Appended Table 1 applies to exports to all destinations worldwide.[4] Licenses come in two broad types: an individual license, which must be obtained for each transaction, and a bulk license, which grants blanket authorization for a fixed period so that no license is required on a transaction-by-transaction basis.[5] This distinction between individual and bulk licenses matters for the cost structure of compliance—an individual license entails a per-transaction cost that recurs with each shipment, whereas a bulk license entails a largely fixed, up-front cost of establishing an internal compliance program—and we return to it in Section 3.2.

2.4 Revisions and the scope of this paper

When the regimes revise their control lists, Japan implements the changes by amending the Ministerial Order. Each amendment changes the boundary of the controlled set: a tightening expands the specification range (bringing previously uncontrolled goods under control), and a relaxation narrows it (releasing previously controlled goods). Between 2014 and 2023, the Ministerial Order was amended nine times, in every year except 2016. Section 3.3 describes each amendment and its classification into tightening and relaxation at the level of individual control codes. Three scope restrictions follow from this institutional structure. First, we study controls on goods under the Export Trade Control Order, not technology transfers under the Foreign Exchange Order, because our data are customs export

[4]Besides Appended Table 1, the Export Trade Control Order also contains other appended tables, notably Appended Table 2, which lists goods subject to destination-specific controls (for example, measures targeting particular countries such as Russia and Belarus), and Appended Table 4, which specifies exemptions from the licensing requirement, such as those for low-value transactions. Because these are either destination-specific or value-based rather than the worldwide, specification-based list control we study, we exclude transactions under these tables from our dataset.

[5]A bulk license is available only to exporters that maintain an approved internal compliance program, and it comes in several kinds, of which the two main ones are the general bulk license and the special bulk license. The general bulk license covers a predetermined set of goods-destination combinations—in particular, destinations in the trusted country group of Appended Table 3 of the Export Trade Control Order, which comprises countries participating in the international export control regimes—while the special bulk license covers a broader set of combinations and is available to exporters that meet stricter internal-control requirements, including an on-site inspection by METI. Either type is valid for a period of up to three years.

Japanese Category	Type of Control	Classification of the Items			International Regimes
1	List Control	Military items	Arms		WA/ML
2		Dual-use items	WMD-related	Nuclear items	NSG
3				Chemical weapons	AG
3-2				Biological weapons	AG
4				Missiles	MTCR
5			Conventional arms-related	Advanced materials	WA Cat. 1
6				Material processing	WA Cat. 2
7				Electronics	WA Cat. 3
8				Computers	WA Cat. 4
9				Communication/ Information security	WA Cat. 5
10				Sensors and lasers	WA Cat. 6
11				Navigation/avionics	WA Cat. 7
12				Marine	WA Cat. 8
13				Aerospace/propulsion	WA Cat. 9
14				Other ML items	Except for WA/ML
15				Sensitive items	WA very sensitive
16	Catch-All Control	Items other than those under Categories 1 - 15			

Figure 2: The controlled item list

Source: CISTEC. (2015) "Overview of Japan's export control (fourth edition)"

declarations, which record goods shipments. Second, we study list control under Appended Table 1—the worldwide, specification-based control—and not country- or end-use-specific measures (such as sanctions or the catch-all control), so that the treatment we analyze is a change in the technical boundary of regulation applied uniformly to all destinations. Third, the amendments we use in estimation all implement agreements reached in the four multilateral regimes; their content and timing are therefore determined independently of individual Japanese firms' export performance, which supports the exogeneity underlying our identification strategy. The 2014 and 2023 revisions are not used in estimation—the latter because it added items to the control list at Japan's own discretion rather than implementing a regime agreement, and both because they lie at the ends of the sample period, where sufficiently long windows before and after the change cannot be secured (Section 3.3).[6]

3 Data and empirical evidence

3.1 Data

Export declaration data — Our primary data source is the export declaration data of the Ministry of Finance. The data cover the period from January 1, 2014 to December 31, 2023, capturing all export transactions from Japan during this period. The data are organized at the transaction level, and for each transaction they include the exporter, the exporter's corporate number, the name of the consignee, the destination country, the transaction value (in yen), and the volume, among other information. These data also allow us to identify re-export transactions—that is, foreign-origin goods that, having once been imported into Japan, are shipped abroad as they are—but we exclude such transactions from our analysis. A distinctive feature of these data is that the correspondence to export control regulation can be observed directly at the transaction level. Transactions exported under a license from METI carry a five-digit code corresponding to the item number in Appended Table 1, allowing us to identify precisely which controlled item a given transaction falls under. For example, item 7 of Appended Table 1 contains 32 sub-items; when a good subject to control under sub-item 1 of item 7 is exported under a license, the value 10701 is recorded in the variable for the Appended Table code in the declaration data. For transactions that do not require a license, by contrast, this variable is left blank. Furthermore, we can identify which type of license the export was made under—an individual license or one of the various bulk licenses. We can therefore determine exhaustively whether a given transaction falls under a controlled item, and if so, which item it belongs to and under which license it was exported.

This represents an advantage of our data over the existing literature on export controls. Existing studies observe exports only through product classifications, which carry no information on whether the goods traded are subject to control, and have therefore been

[6] Hayakawa et al. (2025) examine the trade effects of the 2023 revision using monthly trade statistics.

unable to capture the effects of export controls directly. Control status is determined by technical specifications and therefore cuts across the HS classification, so that controlled and non-controlled goods coexist within the same HS code; the institutional structure behind this disconnect is discussed in Section 4.1. Our data address this problem directly through the five-digit code described above.

Firm financial data — The export declaration data, however, are transaction-level data and do not contain basic attributes of the exporting firms themselves—information such as sales, number of employees, industry, and year of establishment. Such firm attributes are indispensable for analyzing how heterogeneity across firms shapes their responses to regulation. We therefore link the export declaration data to the firm database of Tokyo Shoko Research (TSR) using the corporate number. TSR compiles extensive financial and corporate information on Japanese firms, enabling us to observe the size and industry of each exporter and to conduct analyses that focus on firm heterogeneity. In this paper, we restrict the analysis to firms that could be linked to TSR.

3.2 Descriptive evidence on controlled-goods exports

This subsection documents what controlled-goods trade looks like in our data. We first describe the aggregate scale of controlled-goods exports and the characteristics of the firms that handle them, and then turn to what the comparison between controlled and non-controlled transactions reveals about the structure of compliance costs.

3.2.1 Aggregate patterns and exporter characteristics

Figure 3 contrasts all goods with controlled goods. Exports of controlled goods account for only 1.4–1.8 percent of total exports (panel (a)): the scale of controlled-goods trade is very small relative to non-controlled trade. Decomposing this small scale into the extensive and the intensive margin, the extensive margin shows that only 1,100–1,350 firms—3–4 percent of all exporters—handle controlled goods (panel (b)): controlled-goods trade is confined to a small set of firms. Moreover, unlike the number of exporters overall, this number has not recovered from its sharp drop in 2020. On the intensive margin, the average number of

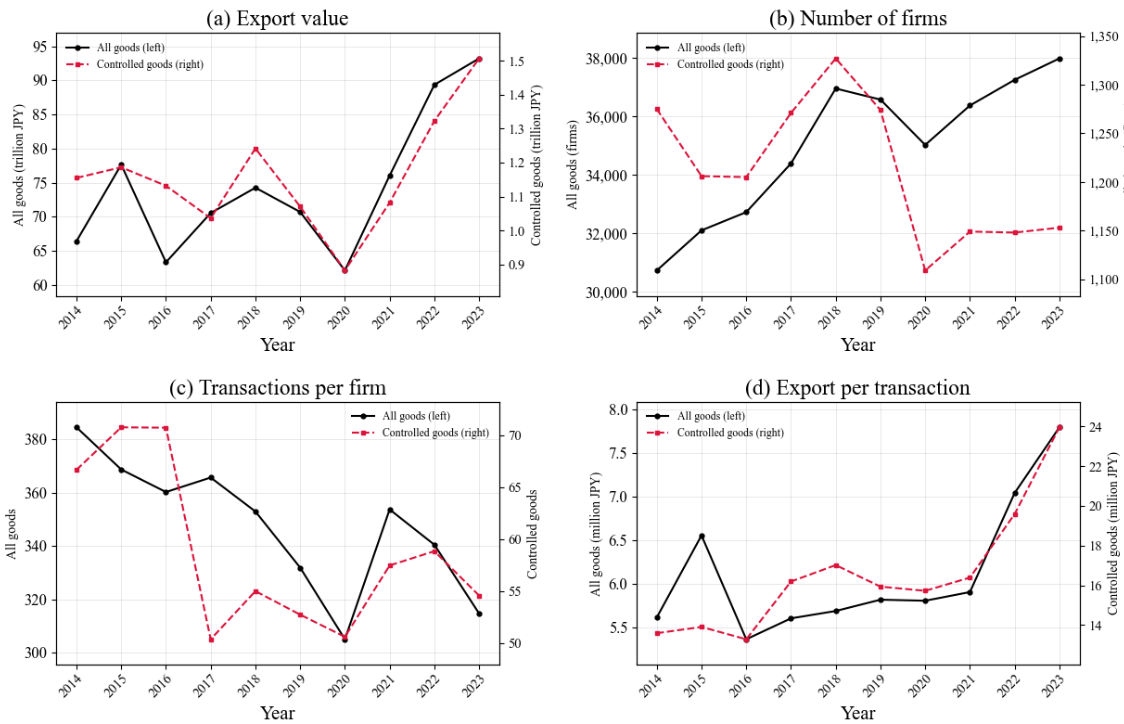

Figure 3: Trends in exports: all goods versus controlled goods, 2014–2023

Notes: In each panel, the solid black line refers to all goods (left axis) and the dashed red line to controlled goods (right axis). Panel (a) reports total export values, panel (b) the number of exporting firms, panel (c) the average number of transactions per firm, and panel (d) the average export value per transaction.

transactions per firm evolves similarly for the two groups (panel (c)), whereas export value per transaction differs markedly (panel (d)): a controlled transaction is worth about two and a half to three times as much as an average transaction, and the ratio rises over the period. What characterizes controlled-goods trade is thus not transaction frequency but the value per transaction—a premium consistent both with the technology-intensive nature of controlled goods and with per-transaction compliance costs discouraging small shipments.

Tables 1 and 2 describe this small set of firms. Firms exporting controlled goods export on average 95–105 distinct products per year, six to seven times the average of 15–18 products across all exporting firms (Table 1).

In the distribution, more than nine in ten exporting firms export fewer than 50 products (Table 2, Panel A), whereas controlled-goods exporters are concentrated in the upper bins (Panel B), and the propensity to handle controlled goods rises steeply with product scope: in 2023, 79 percent of the roughly 90 firms exporting 300 or more products handle controlled

Table 1: Average number of goods exported per firm

	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
All exporting firms	17.95	17.50	17.09	16.63	16.18	15.79	15.16	15.38	15.18	14.76
Controlled-goods exporting firms	104.50	105.30	105.06	103.28	103.38	99.47	97.50	96.52	94.81	94.45

Notes: Each cell reports the average number of goods exported per firm in the given year. "All exporting firms" covers all firms that export, and "Controlled-goods exporting firms" covers firms that export at least one controlled good.

goods. Controlled-goods trade is thus confined almost entirely to Japan's largest and most diversified exporters. At the same time, Panel C shows that about half of the firms exporting any controlled good export exactly one, most export five or fewer, and only 22–31 firms export more than 15. Controlled items occupy only a small corner of these large firms' portfolios.

Table 2: Number of exporting firms by the number of goods exported

	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Panel A. All exporting firms, by # of goods exported										
$0 < x < 50$	28,110	29,485	30,094	31,722	34,165	33,898	32,567	33,754	34,664	35,389
$50 \leq x < 100$	1,609	1,605	1,621	1,645	1,745	1,686	1,562	1,670	1,612	1,642
$100 \leq x < 300$	906	907	908	928	940	896	817	862	885	862
$300 \leq x$	109	108	105	98	105	97	81	89	92	90
Panel B. Controlled-goods exporting firms, by # of goods exported										
$0 < x < 50$	584	553	535	574	590	580	507	535	557	553
$50 \leq x < 100$	239	228	247	260	273	264	240	249	221	223
$100 \leq x < 300$	362	338	334	351	373	349	293	291	299	306
$300 \leq x$	90	87	89	86	91	81	69	74	71	71
Panel C. All exporting firms, by # of controlled goods exported										
0	29,459	30,899	31,523	33,122	35,628	35,303	33,918	35,226	36,105	36,830
1	676	604	628	641	681	637	576	583	572	573
$1 < x \leq 5$	450	455	435	487	508	489	416	435	454	454
$5 < x < 15$	119	116	115	114	111	126	90	108	100	100
$15 \leq x$	30	31	27	29	27	22	27	23	22	26

Notes: In Panels A and B, x denotes the number of goods a firm exports; in Panel C, x denotes the number of controlled goods a firm exports. Panel A covers all exporting firms, Panel B covers firms that export at least one controlled good, and Panel C classifies all exporting firms by the number of controlled goods they export (the row 0 therefore corresponds to firms exporting no controlled goods).

Table 3 looks inside the same firm and the same HS6 code, computing the share of controlled transactions for firm×HS6 pairs averaging at least 20 transactions per year. In more than nine in ten of the cells containing any controlled transaction, controlled and non-controlled transactions coexist: in 2023, of the 1,404 cells with any controlled transaction, only 80—fewer than 6 percent—consist entirely of controlled transactions. This pattern is stable across all ten years. Even conditioning on the firm and a six-digit product code, an HS-based classification cannot separate controlled from non-controlled trade—direct, micro-level evidence for the problem discussed in Section 3.1. Our measurement of treatment in Section 4 is designed around precisely this feature of the data.

Table 3: Distribution of the share of controlled transactions

Share (x, %)	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Panel A. # of observations										
0	60,585	62,036	63,084	64,524	65,866	65,828	66,080	67,756	68,359	67,324
$0 < x \leq 25$	1,176	1,195	1,114	1,106	1,110	1,060	964	1,020	1,036	1,072
$25 < x \leq 50$	192	168	160	123	123	121	121	113	127	129
$50 < x \leq 75$	84	92	112	78	76	63	57	68	68	60
$75 < x < 100$	80	84	75	67	77	79	78	68	71	63
100	51	60	61	56	63	73	70	78	70	80
Panel B. # of firms										
0	9,149	9,383	9,570	9,905	10,234	10,230	10,264	10,667	10,839	10,827
$0 < x \leq 25$	414	420	410	410	403	379	371	379	376	392
$25 < x \leq 50$	112	92	98	105	102	100	101	95	102	107
$50 < x \leq 75$	60	68	74	72	69	54	53	62	62	53
$75 < x < 100$	54	60	56	59	67	66	65	57	58	50
100	44	51	53	49	51	59	56	64	55	64

Notes: For each firm×HS6×year, we compute the share of controlled transactions, defined as the number of controlled transactions divided by the total number of transactions for that firm×HS6×year, and classify observations into share bins (x, in percent). Panel A reports the number of firm×HS6×year observations in each bin, and Panel B reports the number of distinct firms. The sample is restricted to firm×HS6 pairs whose average number of transactions per year over the ten-year period is at least 20.

These patterns—an export value per transaction several times higher, and the concentration among large, diversified exporters—are consistent both with controlled goods simply being higher-value products and with compliance costs shaping how, and by whom, controlled goods are exported. The next subsection turns to this question by comparing controlled and non-controlled transactions within the same firm and product.

3.2.2 Compliance Costs of Controlled Exports: Unit price, quantity, and frequency

The coexistence of controlled and non-controlled transactions within narrowly defined products allows us to characterize the cost structure of compliance. Table 4 compares the two types of transactions within firm × HS6 × destination × year cells that contain both.

Table 4: Compliance costs of controlled exports

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		$\ln p$			$\ln Q$		$\ln N$
	Baseline	License	Size	Baseline	License	Size	Baseline
Controlled	0.515***	0.548***	0.225**	−0.0257	−0.0575	0.341***	−0.700***
	(0.0929)	(0.0956)	(0.0901)	(0.147)	(0.156)	(0.0913)	(0.0444)
× Individual license		−0.331***			0.321**		
		(0.0903)			(0.154)		
× Large firm			0.319***			−0.403**	
			(0.109)			(0.173)	
Cell FE	✓	✓	✓	✓	✓	✓	✓
Year–month FE	✓	✓	✓	✓	✓	✓	
Observations	2,981,850	2,981,850	2,981,850	2,981,850	2,981,850	2,981,850	78,752

Notes: A transaction is *controlled* if it carries an Appended Table 1 code; the sample consists of list-controlled and unregulated transactions only. Columns (1)–(6) are estimated at the transaction level and a cell is a firm × HS6 × unit of measurement × destination × year; column (7) is estimated on data aggregated to the cell × control status level, conditional on annual export value, where a cell is a firm × HS6 × destination × year. $p = V/Q$ is the transaction unit value, Q the transaction quantity, and N the annual number of shipments. Estimation is restricted to cells containing both controlled and non-controlled transactions. *Large firm* equals one if the firm has 300 or more employees in the TSR data. Standard errors, two-way clustered by firm and HS6, are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Controlled transactions command unit values about 67 percent higher than non-controlled ones (column (1)),[7] but this premium does not reflect compliance costs passed through to prices. If it did, it should be largest under individual licenses, which require an application for every shipment; instead, the premium is only about 24 percent for individually licensed transactions against roughly 73 percent for bulk-licensed ones (column (2)), and it is concentrated among large firms—about 72 percent, versus 25 percent for smaller ones (column (3)). Since bulk-license holders and large firms are precisely the exporters of the high-end

[7] Percentage effects are obtained as $\exp(\beta) - 1$: the coefficient of 0.515 in column (1) implies $\exp(0.515) - 1 \approx 0.67$. The same transformation applies to the other percentages reported in this subsection; in the interacted specifications, the relevant coefficient for the interacted group is the sum of the main effect and the interaction term (for example, individually licensed transactions in column (2) correspond to $0.548 - 0.331 = 0.217$, or about 24 percent).

controlled segment, the premium points to quality, not cost pass-through. Where compliance costs do appear is in shipment size and frequency. Individually licensed shipments are about 38 percent larger than bulk-licensed ones (column (5)); in the size split, the enlargement of controlled shipments is concentrated among smaller firms—about 41 percent, with essentially no enlargement among large ones (column (6)); and, conditional on annual export value, controlled exports are shipped in about half as many transactions as non-controlled exports (column (7)). Firms facing a cost per license application consolidate exports into fewer, larger batches; firms with bulk licenses—disproportionately the large ones—do not need to. Compliance costs thus operate as per-shipment and fixed costs rather than as ad valorem variable costs, a structure that underlies our attention to firm heterogeneity in what follows.

3.3 Export control revisions

METI publishes information on the revisions to technical thresholds and related changes that are made pursuant to agreements reached in the international export control regimes.[8] For each revision, it provides an overview stating which provisions of the Ministerial Order were amended and whether the amendment constitutes a tightening of controls, a relaxation, both a tightening and a relaxation, or merely an operational change that clarifies the content of the regulation. From these publications, we extract the revisions to technical thresholds made between 2014 and 2023 and record, for each revision, which provisions were tightened and which were relaxed. However, what these publications record is, as a rule, which provisions of the Ministerial Order were amended, not which items of Appended Table 1 they correspond to. Because our analysis is based on the Appended Table 1 item numbers assigned in the export declaration data, we need to translate each revision into the items of Appended Table 1 to which it pertains. We therefore use the correspondence table published by METI to map the amended provisions of the Ministerial Order to the item numbers of Appended Table 1.[9]

Figure 4 summarizes the resulting classification. Between 2014 and 2023, the Ministerial

[8] `https://www.meti.go.jp/policy/anpo/law09-2.html`
[9] `https://www.meti.go.jp/policy/anpo/matrix_intro.html`

Order was revised nine times.[10] Each revision is dated by its enforcement date, with the promulgation date in parentheses, and the figure reports, at the level of Appended Table 1 item codes, how many codes were tightened and how many were relaxed in each revision. In the 2014 revision, for example, 21 item codes were tightened ("Tight."), one was relaxed ("Rel."), and one was both tightened and relaxed.[11]

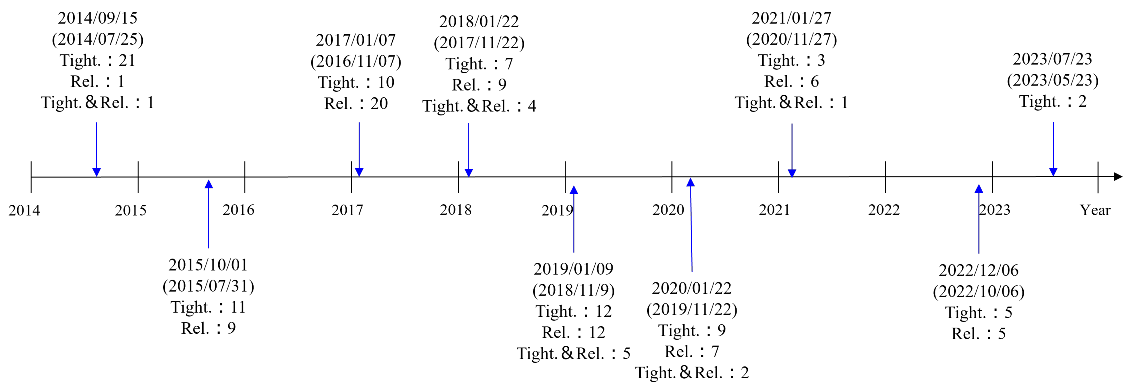

Figure 4: Dates of revisions between 2014 and 2023

Notes: Dates in parentheses indicate the promulgation date of each revision. "Tight." and "Rel." mean tightening and relaxation, respectively.

Of the nine revisions shown in Figure 4, the revision enforced in July 2023 differs from the other eight in its origin. It did not implement an agreement reached in the four multilateral regimes; instead, it newly brought manufacturing equipment for advanced semiconductors and related items under control at Japan's own discretion. In form, it was implemented—like the other revisions—as an amendment to the technical specifications of the Ministerial Order, and it applies to exports to all destinations without singling out any particular country; its origin, however, does not lie in a multilateral agreement. In addition, for the revisions at the two ends of the period, those of 2014 and 2023, sufficiently long aggregation windows before and after the change cannot be secured. For these reasons, we include these two revisions in our description of the policy changes but do not use them in the estimation in Section 5.

[10]The measure toward South Korea enforced on July 4, 2019 changed the licensing arrangement for three items rather than any technical threshold (see footnote 2 in Section 2.1); because it is an operational change, not a revision to technical specifications, it does not appear in the figure and is not treated in this paper.

[11]A single item code of Appended Table 1 covers a range of technical classifications, so that within the same code, some technologies may be tightened while others are simultaneously relaxed.

4 Empirical strategy

4.1 Conceptual issue: mapping

The first challenge in our empirical analysis is to map changes in regulation onto units that are observable in trade data. As described in Sections 2 and 3, revisions to technical specifications are made at the level of the provisions of the Ministerial Order, and the correspondence table published by METI allows us to identify which items of Appended Table 1 each amended provision corresponds to. Beyond that point, however, no official correspondence exists: there is no table linking each item of Appended Table 1 to the codes of the HS classification. Figure 5 illustrates this structure. The chain from technical specifications (the Ministerial Order) to control codes (Appended Table 1) is connected by government-issued correspondence tables, but the link from control codes to HS codes is unknown, and it is here that the chain breaks. This break is not accidental; it is rooted in the design of the institutions. First, jurisdiction differs. The classification of controlled items (Appended Table 1 and the Ministerial Order) is set by METI, which administers export controls, whereas the product classification of trade statistics (HS) is managed by the Ministry of Finance, which administers customs; no body is charged with producing a table that connects the two. Second, the principles of classification differ. HS codes are defined by product characteristics such as materials and uses, not by the technical performance thresholds that export controls turn on. Institutionally and conceptually, therefore, no one-to-one correspondence exists between control codes and HS codes.

Our data, however, provide an empirical way to bridge this gap. As described in Section 3.1, transactions exported under a license carry an Appended Table code, and the HS code declared for the same transaction is recorded alongside it. We can therefore observe directly, from the data, the set of HS codes under which trade has actually taken place for each control code. In the example of Figure 6, exports under control codes 10716 and 10717—both corresponding to Article 6, Item 17 of the Ministerial Order (equipment for manufacturing or testing semiconductor devices)—are observed across multiple HS codes, and this observed correspondence constitutes the empirical link between control codes and HS codes.

Even with this empirical link, however, we cannot fully identify the goods affected

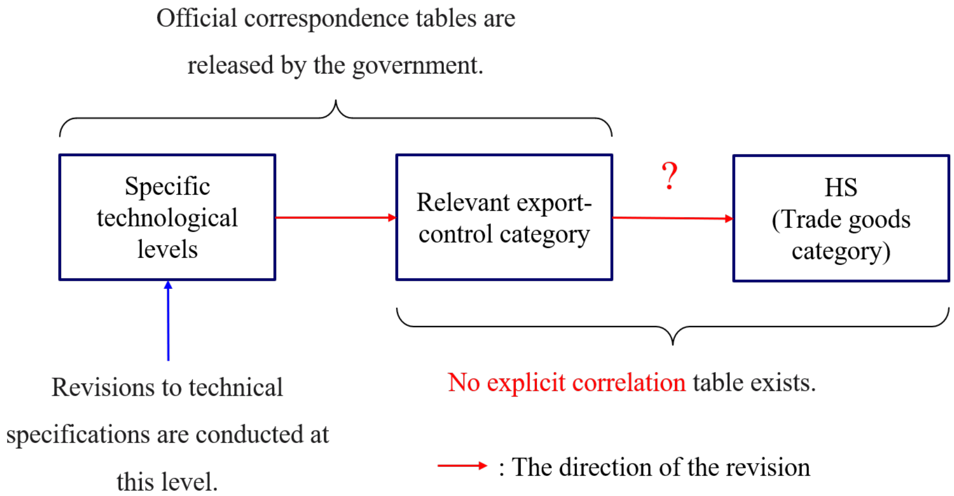

Figure 5: Mapping Export-control categories to HS codes

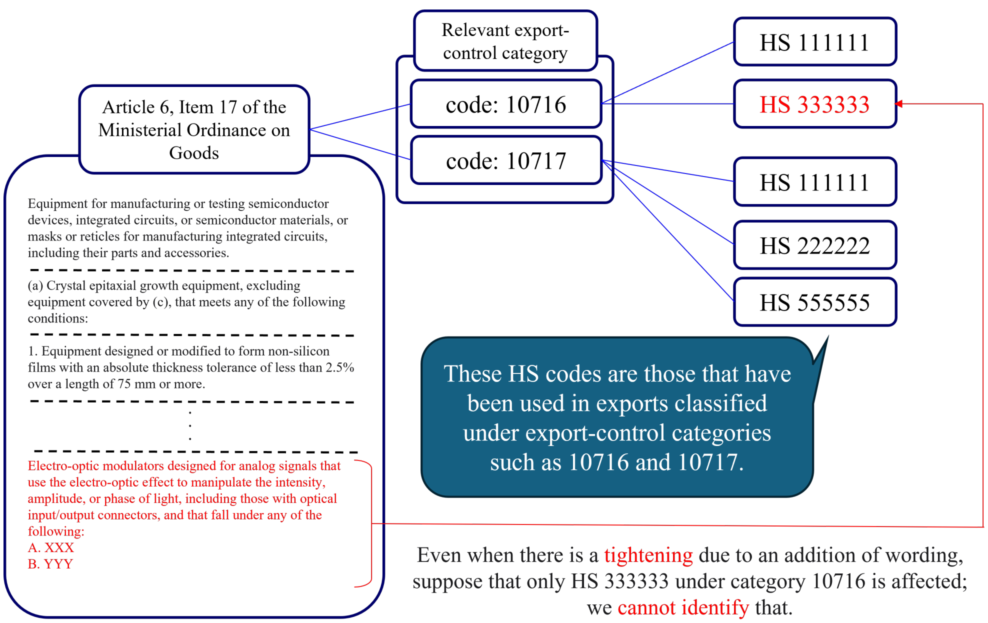

Figure 6: Affected HS codes by revision

by a revision. Figure 6 illustrates the limitation. Suppose that a revision adds wording to the technical specifications of control codes 10716 and 10717, tightening the regulation. From the data, we know the set of HS codes under which transactions have been recorded under codes 10716 and 10717. Yet the added wording may in fact capture only a subset of that set—say, only goods classified under HS 333333. Because HS codes do not distinguish technical performance, neither the data nor the legal text allows us to determine which of the HS codes linked to codes 10716 and 10717 have truly come under the new control. In short, the incidence of a revision can be traced precisely down to the level of control codes, but only partially to the level of HS codes. This limitation has a direct implication for the definition of treatment. Since we cannot identify one by one the HS codes affected by a revision, treatment must instead be constructed statistically from changes in the empirical correspondence between control codes and HS codes—that is, from how the set of HS codes in which controlled transactions are observed, and the volume of those transactions, changes around each revision. The next subsection defines our measurement of treatment based on this idea.

4.2 Identifying affected HS

Aggregation windows. As shown in the previous subsection, the incidence of a revision can be traced precisely down to the level of control codes, but which HS codes are truly affected cannot be determined from the legal texts. We therefore identify the affected HS6 codes in a data-driven manner, from changes in controlled transactions around each revision. Specifically, for each revision we set one-year aggregation windows on each side of the enforcement date, defining the year before enforcement as “Pre” and the year after as “Post” (Table 5). As discussed in Section 3.3, the 2014 revision does not allow a sufficiently long Pre window and the 2023 revision does not allow a sufficiently long Post window, so this aggregation covers the seven revisions from 2015 through 2022.

Jump classification. For each revision, we aggregate, for each HS6 $\times$ control-code pair, the number of controlled transactions in the Pre and Post windows, and we define the HS6 codes whose controlled-transaction counts change sharply (“jump”) between the two windows as the HS6 codes affected by that revision. A jump requires that the absolute change be at least 50 transactions and that the rate of change be at least 25 percent. For pairs with zero controlled transactions in the Pre window and positive transactions in the Post window, the rate of change is undefined; for these pairs we instead require at least 15 transactions in the Post window. Table 6 illustrates the criteria with hypothetical examples. The pair of HS

Table 5: Aggregation periods before and after each regulation change

Date of the regulation change	Aggregation period before the regulation ("Pre")	Aggregation period after the regulation ("Post")
2015/10/01	2014/10/01 – 2015/09/30	2015/10/01 – 2016/09/29
2017/01/07	2016/01/08 – 2017/01/06	2017/01/07 – 2018/01/06
2018/01/22	2017/01/22 – 2018/01/21	2018/01/22 – 2019/01/21
2019/01/09	2018/01/09 – 2019/01/08	2019/01/09 – 2020/01/08
2020/01/22	2019/01/22 – 2020/01/21	2020/01/22 – 2021/01/20
2021/01/27	2020/01/28 – 2021/01/26	2021/01/27 – 2022/01/26
2022/12/06	2021/12/06 – 2022/12/05	2022/12/06 – 2023/12/05

Notes: The 2014 and 2023 revisions are excluded because sufficiently long Pre and Post windows, respectively, are not available for them.

333333 and control code 10716 exhibits a jump under a tightening (215 to 400 transactions: an absolute change of 185 and a rate of change of 86 percent), and the pair of HS 666666 and code 10501 exhibits a jump under a relaxation (60 to 10: an absolute change of 50 and a rate of −83 percent). By contrast, the pair of HS 222222 and code 10717 is not classified as a jump: although its rate of change (43 percent) exceeds the threshold, its absolute change (9 transactions) does not.

Table 6: Examples of the jump classification

HS6	Export-Control Code	Period	Revision	# of Controlled Transactions
333333	**10716**	Pre	Tightening	**215**
333333	**10716**	Post	Tightening	**400**
222222	10717	Pre	Tightening	21
222222	10717	Post	Tightening	30
666666	**10501**	Pre	Relaxation	**60**
666666	**10501**	Post	Relaxation	**10**

Notes: This table illustrates the jump classification with hypothetical figures. Rows in bold indicate HS6 × control-code pairs classified as jumps.

Because an HS6 code is often linked to more than one control code, jumps may also arise in the same HS6 code through transactions under a different control code; the test is therefore conducted separately for each HS6 × control-code pair.[12] We exclude from the treatment group jumps whose interpretation is ambiguous: (i) HS6 codes in which jumps of opposite

[12] Appendix A.1 illustrates this many-to-many structure and the resulting attribution.

directions coexist—a positive jump under a tightened code together with a negative jump under a relaxed code—since the two effects cannot be separated; (ii) jumps in the direction opposite to that of the revision, which likely reflect variation unrelated to it; (iii) jumps under control codes subject to both tightening and relaxation in the same revision ("Tight. & Rel." in Section 3.3), for which the observed change cannot be attributed to either; and (iv) HS6 codes that jump more than once, whether across revisions or through multiple control codes within a revision, since recurring jumps prevent attributing treatment to a single revision and would contaminate the event window.

Control group. The control group is drawn from the same "regulated universe" as the treated group: the set of HS6 codes with positive controlled transactions during the one year preceding the enforcement date of the revision. HS6 codes in which no controlled transaction is ever observed differ substantially from those containing controlled goods, in both product characteristics and the composition of exporting firms, and are unsuitable as comparisons. Within the regulated universe, an HS6 code is assigned to the control group if (i) no jump is observed for it in any of the revisions from 2015 through 2022, paralleling the exactly-once requirement on the treatment side; (ii) it is not linked, through its transactions during the year before enforcement, to any control code amended in the revision in question or in any adjacent revision enforced within two years of it, since linked codes may be partially affected even without meeting the jump criteria; and (iii) it belongs to the same HS2 chapter as a treated HS6 code, ensuring that treated and control pairs are compared within narrowly defined product categories.

Table 7 reports, for each revision, the number of treated HS6 codes identified by the jump procedure and the number of control HS6 codes assigned by these criteria. In total, 19 HS6 codes are identified as treated under tightening and 39 under relaxation. No treated HS6 codes are identified on the tightening side for the 2020 and 2021 revisions, and none on the relaxation side for the 2018 and 2022 revisions; accordingly, the former two events do not contribute to the tightening arm of the estimation, and the latter two do not contribute to the relaxation arm.

Table 7: Number of treated and control HS6 codes by revision

Year of the revision		2015	2017	2018	2019	2020	2021	2022
Tightening	Treated HS6	4	3	2	6	0	0	4
	Control HS6	50	57	80	93	0	0	6
Relaxation	Treated HS6	6	27	0	3	1	2	0
	Control HS6	24	99	0	22	62	37	0

Notes: For each revision, the row "Treated HS6" reports the number of treated HS6 codes identified by the jump procedure, and the row "Control HS6" the number of control HS6 codes assigned by the criteria described in the text. Appendix Table 8 and 9 list the individual codes.

4.3 Empirical specification

To estimate the effects of the revisions, we construct a stacked event-study panel of firm–HS6 pairs at the quarterly frequency. Stacking estimates each revision in its own clean window and pools them, avoiding the contamination that arises in two-way fixed-effects designs when treatment timing is staggered (Gormley and Matsa, 2011; Cengiz et al., 2019; Baker et al., 2022). For each revision, the panel contains the treated and control HS6 codes defined in Section 4.2 over event time $\tau \in [-4, +4]$, where $\tau = 0$ is the quarter containing the enforcement date. The tightening and relaxation arms are stacked and estimated separately. For the participation dummy $active_{ihrt}$, which equals one if firm i records at least one transaction in HS6 code h in quarter t of event r, we estimate by least squares

$$active_{ihrt} = \sum_{\substack{\tau=-4 \\ \tau\neq-1}}^{+4} \beta_\tau \, T_{hr} \cdot \mathbf{1}\{t - e_r = \tau\} + \alpha_{ihr} + \lambda_{r,c(h),t} + \varepsilon_{ihrt}, \tag{1}$$

and for the number of transactions and the export value, $y_{ihrt} \in \{trans_{ihrt}, \, val_{ihrt}\}$, both of which contain a mass of zeros, we estimate by PPML

$$\mathbb{E}\left[y_{ihrt}\right] = \exp\left(\sum_{\substack{\tau=-4 \\ \tau\neq-1}}^{+4} \beta_\tau \, T_{hr} \cdot \mathbf{1}\{t - e_r = \tau\} + \alpha_{ihr} + \lambda_{r,c(h),t}\right), \tag{2}$$

where T_{hr} equals one if h is a treated code in event r; e_r is the enforcement quarter; and $\tau = -1$ is the base period. α_{ihr} is a firm $\times$ HS6 $\times$ event fixed effect, and $\lambda_{r,c(h),t}$ is an event $\times$ HS2 $\times$ quarter fixed effect, so that β_τ is identified from differential changes between treated and control codes within the same HS2 chapter, net of shocks common to that chapter. All outcomes are measured on total exports rather than on controlled transactions only, because a revision mechanically reclassifies transactions between the two categories even when actual behavior is unchanged. We date event time by the enforcement date, since promulgation precedes it by roughly two months and we find no anticipatory response in the intervening quarter. Standard errors are two-way clustered by HS6 code—the level at which treatment is assigned—and by firm, allowing for correlation across the products of the same firm (Cameron et al., 2011).

Three sample restrictions apply. First, control observations are drawn only from firms that export no treated HS6 code in the event, so that within-firm spillovers—such as the reallocation of compliance resources across products—do not contaminate the control group. Second, the sample is restricted to incumbents: firm–HS6 pairs with positive transactions during $\tau \in [-8, -5]$, the year preceding the estimation window.[13] Third, not every revision enters the main estimation. The stacked estimator averages revision-specific effects, so a single revision that violates parallel trends contaminates the pooled estimate rather than averaging out. We therefore audit identification revision by revision before pooling: for each revision and arm, we test whether the three pre-enforcement coefficients on export value are jointly zero, and retain the revisions that pass at the 5 percent level.[14] The tightening stack consists of the 2015, 2018, and 2019 revisions, and the relaxation stack of the 2015 and 2017 revisions.

[13]This restriction removes pairs whose first transactions occur in the four quarters before enforcement—often the very flows that triggered treatment status under jump detection. Appendix A.3 details the diagnostics. For the 2015 revision, the first quarter of this window, $\tau = -8$ (2013Q4), predates the start of the data in January 2014, so the incumbent condition is evaluated over $\tau \in [-7, -5]$ for that revision.

[14]Exclusion is based solely on the pre-trend test. The rejected revision–arm pairs are the 2017 and 2022 revisions in the tightening arm and the 2019, 2020, and 2021 revisions in the relaxation arm; notably, the 2019 revision is rejected in the relaxation arm but not in the tightening arm, indicating that rejection reflects the incidence of shocks on specific treated-product sets rather than calendar periods per se. Appendix Table 10 reports the full test matrix.

5 Results

5.1 Tightening

Figure 7a reports the baseline event-study estimates for the tightening arm, pooling the 2015, 2018, and 2019 revisions. The central result is an absence of adjustment among incumbent exporters. Transaction counts are precisely estimated zeros: coefficients range from -0.01 to 0.03 log points across all post-enforcement quarters, with confidence intervals that rule out declines larger than about 4 percent. Export values show no systematic response either; point estimates drift between 0.02 and 0.13 log points and are significant only sporadically at the 10 percent level, so that while large negative effects are firmly rejected, the value margin is less precisely estimated than the transaction margin. Participation, if anything, moves in the opposite direction from what a deterrence story for incumbents would predict: incumbents in treated codes are 2 to 3 percentage points more likely to remain active than incumbents in control codes from the second post-enforcement quarter onward—a relative persistence, not the creation of new exporters.

These results are informative about the structure of compliance costs. If tightening imposed significant variable or per-shipment costs on firms already exporting the affected products, incumbents should reduce shipment frequency, scale back values, or exit; they do none of these. This is consistent with Hayakawa et al. (2023), who detect little effect of Japan's own tightening on dual-use exports. The pattern is consistent with compliance costs that are fixed at the firm level and already sunk for incumbents; the positive participation coefficient is likewise consistent with that sunk cost deterring prospective entrants, though this cannot be distinguished from a demand shift toward treated products in the incumbent panel. An important caveat qualifies this reading. Because controlled and non-controlled transactions coexist within treated HS6 codes (Section 3.2), and outcomes are measured on total exports, the estimates average over incumbents who never handle the goods captured by the amended specifications. A true null for all incumbents and a negative effect on a small exposed subset, diluted by unexposed ones, are observationally equivalent in our design. The results therefore rule out large, broad-based adjustment among incumbents, but they do not establish that no incumbent bears a cost.

(a) Tightening

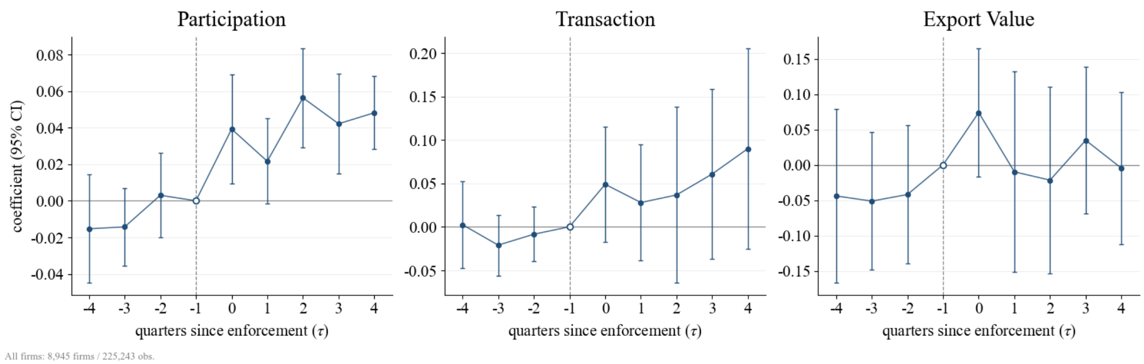

(b) Relaxation

Figure 7: Event-study estimates

Notes: The figure plots the coefficients β_τ from equations (1) and (2), with 95 percent confidence intervals. Panel (a) reports the tightening stack and Panel (b) the relaxation stack. The horizontal axis is event time in quarters relative to the enforcement date, and $\tau = -1$ is the base period (hollow marker). Participation is estimated by least squares; the number of transactions and the export value are estimated by PPML. Standard errors are two-way clustered by HS6 code and firm.

5.2 Relaxation

Figure 7b reports the corresponding estimates for the relaxation arm, pooling the 2015 and 2017 revisions. The distinctive result is on the participation margin. Incumbents in relaxed codes become 2 to 6 percentage points more likely to remain active than incumbents in control codes, and the differential widens over the post-enforcement years rather than fading—roughly twice the size of the corresponding response under tightening, and growing where that response is flat. The intensive margins, by contrast, show no robust movement in the pooled sample: transaction counts drift upward but remain insignificant, and export values are flat.

Taken together, the pooled results from the two arms share one feature and differ in one.

In neither arm do incumbents contract on any margin—list revisions, in either direction, do not push established exporters out or scale them back. Where the arms differ is in activation: relaxation draws incumbents into more continuous exporting in a way that tightening does not. Muted pooled intensive margins are, however, exactly what a per-shipment cost structure predicts if the response is concentrated among the subset of firms for whom those costs bind; the next subsection tests this composition directly.

5.3 Heterogeneity by firm size

The descriptive evidence in Section 3.2 suggested that the burden of compliance costs varies with firm size: controlled-goods trade is concentrated among Japan's largest and most diversified exporters, and the consolidation of controlled exports into fewer, larger shipments is driven by smaller firms. If building and maintaining compliance capability operates as a fixed cost, its burden is relatively heavier for smaller firms, and sensitivity to changes in that cost should be higher on their side. Figures 8a and 8b test this prediction for both arms.

In the relaxation arm, the response is concentrated among small firms. Their transaction counts rise immediately upon enforcement (+9 percent at $\tau = 0$) and expand monotonically thereafter, reaching +16 percent year out. Export values also jump on impact, by 15 to 20 percent in the first two quarters, before settling at positive but attenuating point estimates. Participation rises by 4 to 5 percentage points. Large firms, by contrast, show no systematic response on any margin: point estimates for transaction counts are, if anything, negative, and neither values nor participation move significantly. The muted intensive margins in the pooled estimates thus reflect dilution—unresponsive large firms averaging down responsive small ones. The contrast between transaction counts and export values carries additional information about the mechanism. The expansion among small firms occurs primarily on the frequency margin, with per-shipment size roughly unchanged. This is the event-study mirror image of the cross-sectional finding in Section 3.2.2 that controlled goods are shipped about half as frequently for the same annual value. When relaxation removes the costs the regulation had imposed, firms ship more often; that the response appears on the frequency margin itself suggests that the cost being removed operated at the transaction level.

In the tightening arm, no intensive-margin adjustment is observed for either size class:

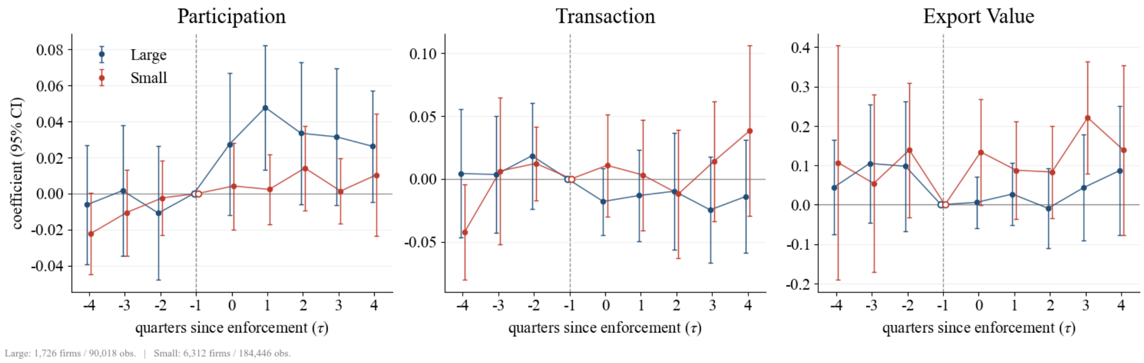

(a) Tightening

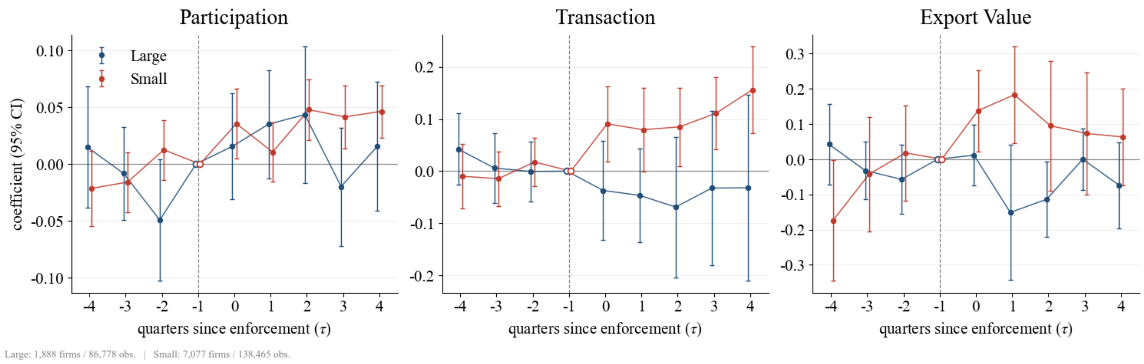

(b) Relaxation

Figure 8: Event-study estimates by firm size

Notes: See the notes to Figure 7. Estimates are obtained separately for large firms (300 or more employees in the TSR data, in blue) and smaller firms (in red); markers are offset horizontally for readability. Panel (a) reports the tightening stack and Panel (b) the relaxation stack.

transaction counts and export values are null for large and small firms alike. The only movement appears on the participation margin, among large firms (+3 to +5 percentage points)—the size-specific counterpart of the relative survival differential discussed in Section 5.1, and we interpret it in the same way. The absence of any movement in small-firm participation is itself notable: the concern that tightening selectively pushes out smaller incumbents—a regressive-burden scenario—finds no support among incumbents. In sum, the size decomposition reinforces the cost-structure interpretation. The response to relaxation is concentrated among small firms, while large firms move in neither direction. The burden of list-based controls is regressive in firm size, but its incidence becomes visible not when the cost is imposed, but when it is removed.

6 Conclusion

This paper asks who bears the burden of export controls when the boundary of regulation moves. Using the universe of Japan's export declarations, in which every licensed transaction carries a control-list code, we translate the revisions of the technical specifications over 2014–2023 into product-level treatments and estimate the effects of tightening and relaxation as separate arms in a stacked event-study design. The two arms give an asymmetric answer. Tightening does not push incumbent exporters out of the affected products: transaction counts and export values show no decline, and treated incumbents are, if anything, more likely to remain active than their controls. Relaxation, by contrast, activates—and its benefit accrues almost entirely to small firms, whose shipments become more frequent while shipment size is roughly unchanged. Together with the transaction-level evidence that controlled goods are shipped less frequently and in larger batches, these results are consistent with compliance costs operating as per-shipment and fixed costs rather than as ad valorem costs: sunk for incumbents, binding for small firms, and therefore regressive in firm size—but visible only when the cost is removed, not when it is imposed. Three limitations delimit the scope of these conclusions. First, because controlled and non-controlled transactions coexist within treated HS6 codes and outcomes are measured on total exports, our estimates average over incumbents unexposed to the amended specifications; the null under tightening rules out broad-based adjustment among incumbents but cannot exclude a cost borne by a small exposed subset. Second, our estimates concern incumbents in a within-country comparison, so they are silent on entry deterrence. Third, they say nothing about the security benefits that motivate the controls, and hence about the net desirability of any revision. Within that scope, however, the results carry a clear message for the design of export control regimes: the routine revision of technical thresholds redistributes real costs across firms, and relaxations—the neglected half of the system's operation—deliver measurable benefits to precisely the firms for whom compliance costs bind.

Appendix A Additional details on identification

A.1 The many-to-many correspondence between control codes and HS codes

This appendix supplements the jump classification of Section 4.2 by illustrating why the test is conducted at the level of HS6 $\times$ control-code pairs, and how a jump in a pair is attributed to an HS6 code. The correspondence between the two classifications is many-to-many. Panel (a) of Figure 9 illustrates this structure schematically. A single control code may be traded across multiple HS6 codes (Group A is traded under both HS I and HS II), and conversely a single HS6 code may involve multiple control codes (HS I contains controlled goods of both Group A and Group B). One-to-one correspondences, such as HS IV and Group D, do occur, but they are the exception rather than the rule. Changes in controlled transactions within an HS6 code therefore cannot be read off a single series: within HS I, transactions as Group A and transactions as Group B must be aggregated separately, and the jump test of Section 4.2 is applied to each of these control-code-specific series. Panel (b) illustrates the attribution. Suppose that Group A is tightened. If a positive jump is observed in HS I's transactions as Group A, HS I is classified as affected by the tightening of Group A; if a jump is also observed in HS II's transactions as Group A, HS II is likewise classified as affected. Conversely, if no jump appears in a given HS code's transactions as Group A, that HS code is classified as unaffected, even though it is linked to Group A. In the example shown, the tightening of Group A produces a jump in HS I but not in HS II, so only HS I is classified as affected. Attributing treatment only to those HS codes in which trade actually jumps is what gives an empirical answer to the mapping problem of Section 4.1: the ministerial ordinance identifies the amended control codes but not the affected HS codes, and the jumps locate the latter within the former. The example also illustrates the exclusion rules of Section 4.2. In HS I, the positive jump from the tightening of Group A coexists with a negative jump from the relaxation of Group B, so HS I is dropped from the treatment group under rule (i); HS II, HS III, and HS IV, each exhibiting at most a single jump, are unaffected by the rule. Rule (iv) operates analogously when the two jumps within

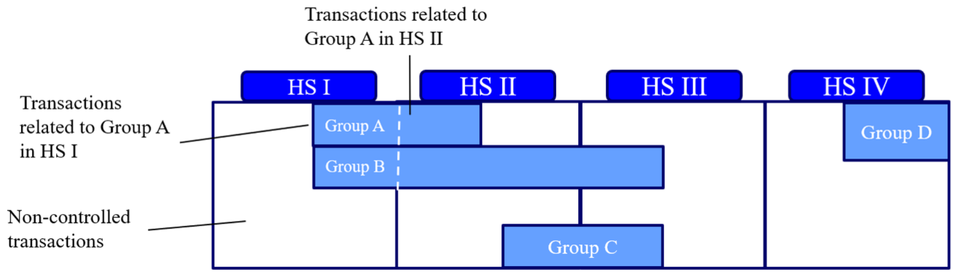

(a) Many-to-many correspondence between control codes and HS codes

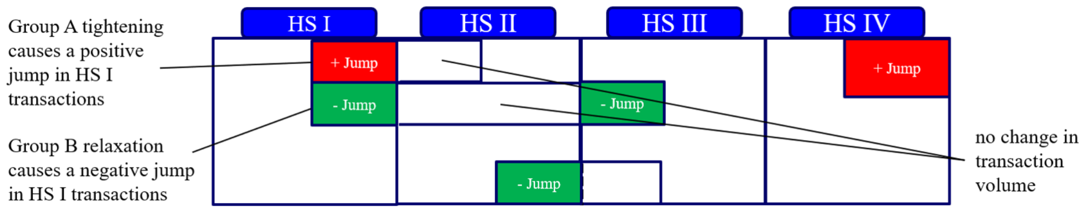

(b) Jumps in controlled transactions upon a revision

Figure 9: Attributing revisions to HS codes through jumps in controlled transactions

Notes: Panel (a) illustrates the many-to-many correspondence between control codes (Groups A–D) and HS6 codes (HS I–HS IV). Shaded boxes represent controlled transactions under each control code within each HS6 code, and the remaining white areas represent non-controlled transactions. Panel (b) illustrates jumps in controlled transactions around a revision. The tightening of Group A causes a positive jump in HS I, while the relaxation of Group B causes a negative jump in HS I: because jumps of opposite directions coexist in HS I, it is excluded from the treatment group. HS codes in which a single jump is observed, such as HS II, HS III, and HS IV, are retained.

an HS6 code arise under codes amended in the same direction or in different revisions.

A.2 Treated HS6 codes

Tables 8 and 9 list the individual HS6 codes identified as treated by the jump procedure of Section 4.2—19 codes under tightening and 39 under relaxation, corresponding to the counts in Table 7—with their product descriptions under the HS 2012 nomenclature.[15]

[15]Product descriptions are taken from the HS 2012 classification files published by the United Nations Statistics Division, `https://unstats.un.org/unsd/classifications/`.

Table 8: Treated HS6 codes: Tightening

Revision	HS6	Description
2015	841350	Pumps; reciprocating positive displacement pumps, n.e.c. in heading no. 8413, for liquids
	841370	Pumps; centrifugal, n.e.c. in heading no. 8413, for liquids
	848110	Valves; pressure reducing, for pipes, boiler shells, tanks, vats or the like
	848130	Valves; check (nonreturn) valves, for pipes, boiler shells, tanks, vats or the like
2017	842121	Machinery; for filtering or purifying water
	845891	Lathes; for removing metal, numerically controlled, other than horizontal lathes
	848610	Machines and apparatus of a kind used solely or principally for the manufacture of semiconductor boules or wafers
2018	281290	Halides and halide oxides of non-metals; excluding chloride
	842199	Machinery; parts for filtering or purifying liquids or gases
2019	370199	Photographic plates and film; (for other than colour photography), in the flat, sensitised, unexposed, with no side exceeding 255mm, of any material other than paper, paperboard or textiles
	370590	Photographic plates and film; exposed and developed, (other than cinematographic film or for offset reproduction)
	841430	Compressors; of a kind used in refrigerating equipment
	850132	Electric motors and generators; DC, of an output exceeding 750W but not exceeding 75kW
	852692	Radio remote control apparatus
	854232	Electronic integrated circuits; memories
2022	690390	Refractory ceramic goods; composition of which n.e.c. in heading no. 6903, other than those of siliceous fossil meals or similar earths
	710420	Stones; precious or semi-precious, synthetic or reconstructed, (not piezo-electric quartz), unworked or simply sawn, roughly shaped, (not strung, mounted or set), temporarily strung for transport convenience
	710490	Stones; precious or semi-precious, synthetic or reconstructed, (not piezo-electric quartz), worked or graded or not, (but not strung, mounted or set), temporarily strung for transport convenience, n.e.c. in item no. 7104.20
	820780	Tools, interchangeable; (for machine or hand tools, whether or not power-operated), for turning

Notes: The table lists the 19 HS6 codes identified as treated under tightening by the jump procedure of Section 4.2, corresponding to the counts in Table 7. Descriptions follow the HS 2012 nomenclature.

Table 9: Treated HS6 codes: Relaxation

Revision	HS6	Description
2015	390799	Polyesters; n.e.c. in heading no. 3907, saturated, in primary forms
	701939	Glass fibres; webs, mattresses, boards and similar non-woven products excluding mats and thin sheets
	761699	Aluminium; articles n.e.c. in heading 7616
	880330	Aircraft and spacecraft; parts of aeroplanes or helicopters n.e.c. in heading no. 8803
	900290	Optical elements; n.e.c. in heading no. 9002 (e.g. prisms and mirrors), mounted, being parts or fittings for instruments or apparatus, of any material (excluding elements of glass not optically worked)
	901390	Optical appliances and instruments; parts and accessories for articles of heading no. 9013
2017	391739	Plastics; tubes, pipes and hoses thereof, n.e.c. in item no. 3917.30
	392350	Plastics; stoppers, lids, caps and other closures, for the conveyance or packing of goods
	401693	Rubber; vulcanised (other than hard rubber), gaskets, washers and other seals, of non-cellular rubber
	690919	Ceramic wares; for laboratory, chemical or other technical uses, other than articles having a hardness equivalent to 9 or more on the Mohs scale or of porcelain or china
	731815	Iron or steel; threaded screws and bolts n.e.c. in item no. 7318.1, whether or not with their nuts or washers
	731819	Iron or steel; threaded screws, bolts and nuts n.e.c. in item no. 7318.1
	732690	Iron or steel; articles n.e.c.in heading 7326
	761610	Aluminium; nails, tacks, staples (other than those of heading no. 8305), screws, bolts, nuts, screw hooks, rivets, cotters, cotter-pins, washers and similar articles
	830710	Tubing; flexible, with or without fittings, of iron or steel
	841231	Engines; pneumatic power engines and motors, linear acting (cylinders)
	841410	Pumps; vacuum
	847170	Units of automatic data processing machines; storage units
	848180	Taps, cocks, valves and similar appliances; for pipes, boiler shells, tanks, vats or the like, including thermostatically controlled valves
	848410	Gaskets and similar joints; of metal sheeting combined with other material or two or more layers of metal
	850110	Electric motors; of an output not exceeding 37.5W
	850151	Electric motors; AC motors, multi-phase, of an output not exceeding 750W
	850432	Transformers; n.e.c. in item no. 8504.2, having a power handling capacity exceeding 1kVA but not exceeding 16kVA
	852990	Reception and transmission apparatus; for use with the apparatus of heading no. 8525 to 8528, excluding aerials and aerial reflectors
	853620	Electrical apparatus; automatic circuit breakers, for a voltage not exceeding 1000 volts
	853630	Electrical apparatus; for protecting electrical circuits, n.e.c. in heading no. 8536, for a voltage not exceeding 1000 volts

Continued on next page

Table 9 (continued)

Revision	HS6	Description
	853710	Boards, panels, consoles, desks and other bases; for electric control or the distribution of electricity, (other than switching apparatus of heading no. 8517), for a voltage not exceeding 1000 volts
	854140	Electrical apparatus; photosensitive, including photovoltaic cells, whether or not assembled in modules or made up into panels, light emitting diodes
	854420	Insulated electric conductors; co-axial cable and other co-axial electric conductors
	854710	Insulating fittings; of ceramics, for electrical machines, of insulating material only (except minor assembly parts), excluding those of heading no. 8546
	902610	Instruments and apparatus; for measuring or checking the flow or level of liquids
	902620	Instruments and apparatus; for measuring or checking pressure
	902730	Spectrometers, spectrophotometers and spectrographs; using optical radiations (UV, visible, IR)
2019	320820	Paints and varnishes; based on acrylic or vinyl polymers, dispersed or dissolved in a non-aqueous medium
	370790	Photographic goods; chemical preparations other than sensitised emulsions, put up in measured portions or put up for retail sale in a form ready for use
	902780	Instruments and apparatus; for physical or chemical analysis, for measuring or checking viscosity, porosity, expansion, surface tension or quantities of heat, sound or light, n.e.c. in heading no. 9027
2020	846241	Machine-tools; punching or notching machines (including presses), including combined punching and shearing machines, numerically controlled, for working metal
2021	340399	Lubricating preparations; other than for the treatment of textile and similar materials, not containing petroleum oils or oils obtained from bituminous minerals
	854519	Carbon electrodes; with or without metal, of a kind used for other than furnaces

Notes: The table lists the 39 HS6 codes identified as treated under relaxation by the jump procedure of Section 4.2, corresponding to the counts in Table 7. Descriptions follow the HS 2012 nomenclature.

A.3 The incumbent restriction and revision-by-revision pre-trend tests

Jump detection identifies treated codes from large changes in licensed transactions around enforcement. Part of what defines treatment is therefore the arrival of new trade flows: 42 percent of treated firm–HS6 pairs in the tightening arm (36 percent in the relaxation arm) record their first transaction within the four quarters preceding enforcement. These entrant pairs are not a random sample of entrants—they are the flows whose appearance contributed to treatment status, so including them regresses the treatment on itself. Figure 10 shows the consequence. Among pairs established before the detection window, the gap in mean export value between treated and control codes is flat throughout: it averages 15 million yen both before and after enforcement. Among pairs first active inside the window, the gap is zero by construction through $\tau = -5$ and then surges to 70 million yen by $\tau = -1$, accounting for 84 percent of the total treated–control gap on the eve of enforcement. The pre-trends in the unrestricted sample are thus carried entirely by entrant pairs; established pairs satisfy parallel trends from the outset.

Figure 10: Treated–control export value gap, by entry timing

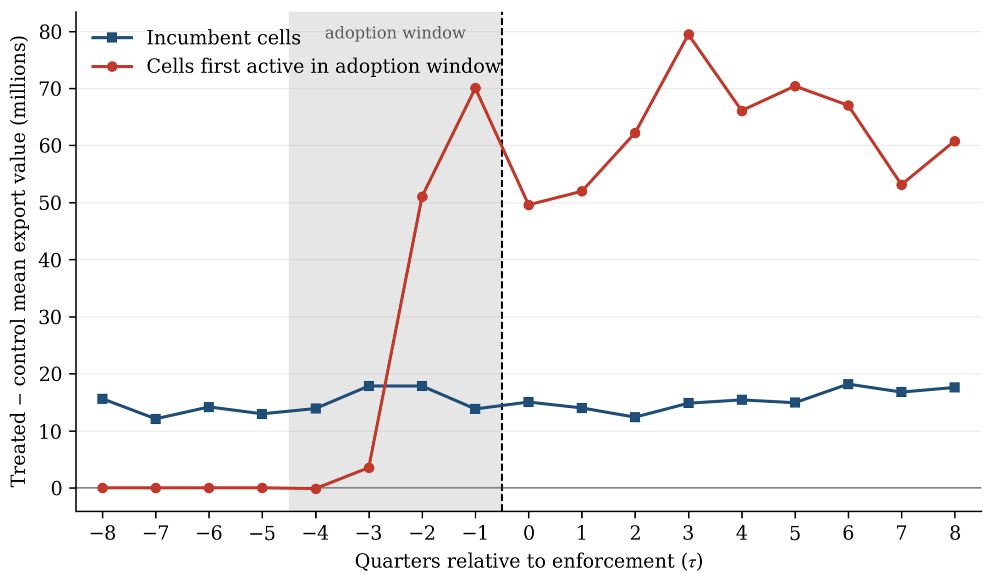

Notes: Tightening arm, all revisions, no incumbent restriction. Each line plots the difference between the mean export value of treated and control firm–HS6 pairs. *Incumbent cells* are pairs with positive transactions during $\tau \in [-8, -5]$, the year preceding the detection window; *cells first active in the adoption window* are pairs whose first transaction falls in $\tau \in [-4, -1]$ and are therefore zero by construction in earlier quarters. The shaded area marks the risk-set adoption window and the dashed line the enforcement date.

Figure 11 confirms this in the regression. Without the restriction, pre-enforcement coefficients on export value in the tightening arm reach -0.74 to -0.64 log points; with it, they collapse toward zero. In the relaxation arm the artifact loads on the participation margin—entrant pairs are inactive by construction in the distant pre-period—and the restriction removes it symmetrically. Conditioning on activity in $\tau \in [-8, -5]$, the pre-period outside the detection window, thus separates entry from selection rather than discarding inconvenient observations.

Figure 11: Event-study estimates with and without the incumbent restriction

(a) Tightening

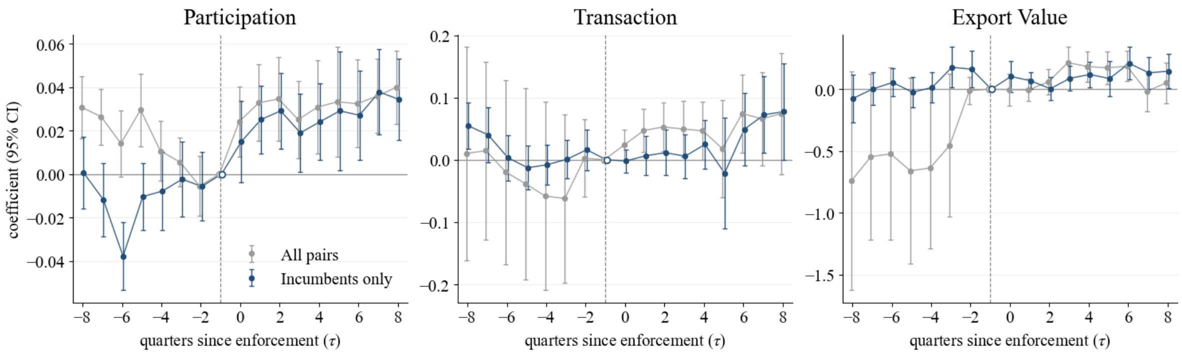

(b) Relaxation

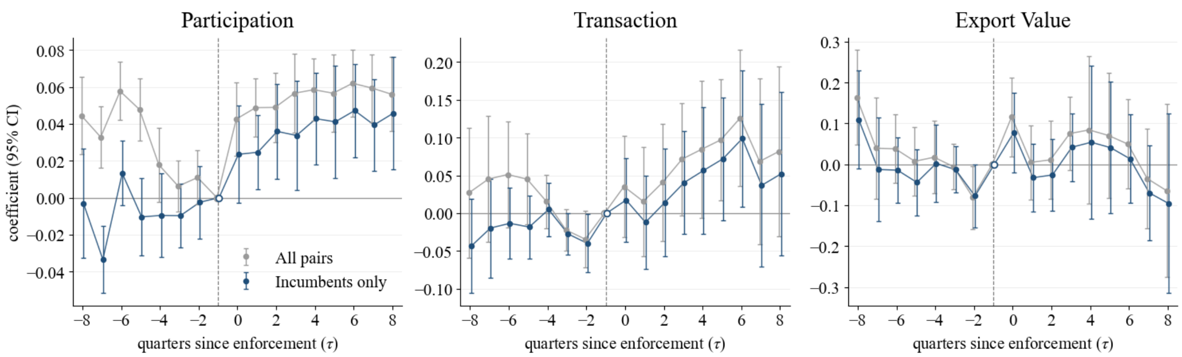

Notes: All revisions, $\tau \in [-8, 8]$, $\tau = -1$ omitted. The panel is unbalanced at the endpoints: $\tau = -8$ falls outside the data for the 2015 revision, and $\tau \geq +5$ falls outside the data for the 2022 revision (tightening arm). Each panel plots coefficients from the unrestricted sample and from the sample restricted to incumbents, with 95 percent confidence intervals. Export value is estimated by PPML and participation by OLS. Standard errors are two-way clustered by firm and HS6.

Even under the incumbent restriction, parallel trends need not hold for every revision. Table 10 reports the revision-by-revision tests described in Section 4.3: for each revision and arm, we test whether the three pre-enforcement coefficients ($\tau = -4, -3, -2$) are jointly zero, separately for participation, transactions, and export value. Based on the export-value test at the 5 percent level, the tightening stack retains the 2015, 2018, and 2019 revisions and the relaxation stack the 2015 and 2017 revisions. The rejections do not follow calendar time: the 2019 revision is rejected in the relaxation arm but retained in the tightening arm, indicating that violations reflect shocks to specific treated-product sets rather than period-wide disturbances.

Table 10: Revision-by-revision pre-trend tests

	p-value of joint pre-trend test					In main
Revision	Participation	Transactions	Export value	Obs.	Firms	stack
Panel A. Tightening arm						
2015	0.007	0.046	0.379	38,718	2,702	✓
2017	0.017	0.012	<0.001	62,829	3,519	
2018	0.076	0.014	0.363	80,694	4,940	✓
2019	0.077	0.578	0.189	155,052	6,018	✓
2022	0.190	<0.001	0.027	11,367	1,154	
Panel B. Relaxation arm						
2015	0.132	0.345	0.680	16,902	1,591	✓
2017	0.255	0.504	0.632	208,341	8,576	✓
2019	0.245	<0.001	<0.001	21,519	2,085	
2020	<0.001	<0.001	<0.001	130,491	5,029	
2021	<0.001	<0.001	<0.001	46,440	3,508	

Notes: For each revision and arm, the table reports the p-value of a joint Wald test that the three pre-enforcement coefficients ($\tau = -4, -3, -2$) in equations (1) and (2) are zero, estimated revision by revision on the same sample as the main estimation (between-firm controls, incumbents only). Participation is estimated by least squares; the number of transactions and the export value are estimated by PPML. Standard errors are two-way clustered by firm and HS6. Obs. and Firms report the number of firm–HS6–quarter observations and distinct firms in each revision-specific panel. A checkmark indicates that the revision enters the main stack; inclusion is based on the export-value test at the 5 percent level (Section 4.3). Revisions with no treated HS6 codes in an arm (Table 7) are omitted.

References

Alfaro, L., Fadinger, H., Schymik, J. S., and Virananda, G. (2025). Trade and industrial policy in supply chains: Directed technological change in rare earths. *NBER Working Paper*, No. 33877.

Ando, M., Hayakawa, K., and Kimura, F. (2024). Supply chain decoupling: Geopolitical debates and economic dynamism in East Asia. *Asian Economic Policy Review*, 19(1), 62–79.

Baker, A. C., Larcker, D. F., and Wang, C. C. Y. (2022). How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144(2), 370–395.

Baldwin, R., and Krugman, P. (1989). Persistent trade effects of large exchange rate shocks. *Quarterly Journal of Economics*, 104(4), 635–654.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2), 238–249.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics*, 134(3), 1405–1454.

Clayton, C., Maggiori, M., and Schreger, J. (2026). A framework for geoeconomics. *Econometrica*, 94(1), 105–136.

Correia, S., Guimarães, P., and Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *Stata Journal*, 20(1), 95–115.

Crosignani, M., Han, L., Macchiavelli, M., and Silva, A. F. (2026). Securing technological leadership? The cost of export controls on firms. *Journal of Financial Economics*, 175, 104192.

Crozet, M., and Hinz, J. (2020). Friendly fire: The trade impact of the Russia sanctions and counter-sanctions. *Economic Policy*, 35(101), 97–146.

Crozet, M., Hinz, J., Stammann, A., and Wanner, J. (2021). Worth the pain? Firms' exporting behaviour to countries under sanctions. *European Economic Review*, 134, 103683.

Das, S., Roberts, M. J., and Tybout, J. R. (2007). Market entry costs, producer heterogeneity, and export dynamics. *Econometrica*, 75(3), 837–873.

Deseatnicov, I., and Fukao, K. (2024). U.S. export controls and the restructuring of global values chains: An analysis of Japanese multinationals' exits from China. *RIETI Discussion Paper*, No. 24-E-082.

Deseatnicov, I., Fukao, K., Hayakawa, K., Ito, K., and Kucheryavyy, K. (2024). Technological decoupling between the U.S. and China. *Hitotsubashi University Institute of Economic Research Discussion Paper Series A*, No. 756.

Dixit, A. (1989). Entry and exit decisions under uncertainty. *Journal of Political Economy*, 97(3), 620–638.

Felbermayr, G., Kirilakha, A., Syropoulos, C., Yalcin, E., and Yotov, Y. V. (2020). The global sanctions data base. *European Economic Review*, 129, 103561.

Gormley, T. A., and Matsa, D. A. (2011). Growing out of trouble? Corporate responses to liability risk. *Review of Financial Studies*, 24(8), 2781–2821.

Hayakawa, K. (2026). The trade effects of the US export control regulations. *International Economics and Economic Policy*, 23(2), 1–24.

Hayakawa, K., and Ito, K. (2025). The collateral damage of US export control regulations on Japanese suppliers' exports to China. *RIETI Discussion Paper*, No. 25-E-061.

Hayakawa, K., Ito, K., Fukao, K., and Deseatnicov, I. (2023). The impact of the strengthening of export controls on Japanese exports of dual-use goods. *International Economics*, 174, 160–179.

Hayakawa, K., Kimura, F., and Yamanouchi, K. (2025). The trade effects of export control regulations in Japan. *Journal of the Japanese and International Economies*, 77, 101366.

Jinji, N., Ito, T., and Matsuura, T. (2025). The impact of US export controls on firms' export behavior in a third country: Evidence from Japanese customs data. *Discussion Paper*, No. E-25-008, Graduate School of Economics, Kyoto University.

Kee, H. L., and Xie, E. (2025). Nickel, steel and cars: Export ban and domestic value-added in Indonesia. *World Bank Policy Research Working Paper*, No. 11249.

Liu, E., Liu, Y., Makarin, A., and Wen, Y. (2024). Export controls and innovation in sanctioned countries. *Harvard Business School Working Paper*, No. 25-004.

Makioka, R., and Zhang, H. (2024). The impact of export controls on international trade: Evidence from the Japan–Korea trade dispute in the semiconductor industry. *Journal of the Japanese and International Economies*, 74, 101336.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6), 1695–1725.

Mohr, C., and Trebesch, C. (2025). Geoeconomics. *Annual Review of Economics*, 17.

Roberts, M. J., and Tybout, J. R. (1997). The decision to export in Colombia: An empirical model of entry with sunk costs. *American Economic Review*, 87(4), 545–564.

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