

BANCA D'ITALIA

EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

The cost of security: evidence on the impact of dual-use export controls on Italian exports

by Lilia Patrignani and Michele Savini Zangrandi

November 2026

Number 1085

BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

The cost of security: evidence on the impact of dual-use export controls on Italian exports

by Lilia Patrignani and Michele Savini Zangrandi

Number 1085 – November 2026

The series Occasional Papers *presents studies and documents on issues pertaining to the institutional tasks of Banca d'Italia and the Eurosystem. The* Occasional Papers *appear alongside the* Working Papers *series which are specifically aimed at providing original contributions to economic research.*

The Occasional Papers *include studies conducted within Banca d'Italia, sometimes in cooperation with the Eurosystem or other institutions. The views expressed in the studies are those of the authors and do not involve the responsibility of the institutions to which they belong.*

The series is available online at www.bancaditalia.it .

ISSN 1972-6643 (online)

Designed by the Printing and Publishing Division of Banca d'Italia

THE COST OF SECURITY: EVIDENCE ON THE IMPACT OF DUAL-USE EXPORT CONTROLS ON ITALIAN EXPORTS

by Lilia Patrignani and Michele Savini Zangrandi[1]

Abstract

Restrictions on exports of dual-use goods – items that can serve both civilian and military purposes – have expanded substantially in recent years, becoming a central instrument of economic security policy. In Italy, a major exporting economy, dual-use regulations require the screening of approximately one third of total exports to assess their potential military use. This paper provides the first causal evidence on the firm-level effects of dual-use export controls. Using granular customs data covering all Italian exporters from 2017 to 2023, we exploit the staggered addition of products to the EU's control list in a difference-in-differences framework. We find that dual-use controls reduce exports significantly in the short run, but have no lasting effect on aggregate values. Instead, they induce a persistent reallocation of trade, as firms offset permanent losses in established markets by shifting trade towards politically aligned partners, where compliance costs are lower and licensing requirements less stringent. In light of the likely further tightening of security-related controls, these findings point to the importance of firms' compliance capacity and of the efficiency of the licensing process in mitigating the trade costs of export controls.

JEL Classification: F13, F52, F14, L51.

Keywords: export controls, dual-use goods, economic security, international trade, geopolitical fragmentation, firm-level trade, compliance costs, trade diversion.

DOI: 10.32057/0.QEF.2026.1085

[1] Bank of Italy. The views expressed are personal and do not necessarily reflect the views of the Bank of Italy or the European Central Bank. We are grateful to Alessandro Amato, Xavier Arnes-Novau, Alessandro Borin, Giovanni Brignone, Stéphane Chardon, Francesco Paolo Conteduca, Antonio della Guardia, Lorenzo Donatelli, Riccardo Ferretti, Adriano Francescangeli, Marcello Irlando, Viktor Ivanov, Fabrizio Leone, Michele Mancini, Manuela Mandarà, Valeria Secchini, Filippo Sevini, Elena Stringa, Cristina Versino, Gabriele Zanardo, Matteo Zardin. All remaining errors are our own.

1. Introduction

Recent years have witnessed a vast expansion in restrictions to the export of dual-use items. These are goods, software and technology that can be used both for civilian and military purposes.[2] The quintessential modern-day example of a dual use item is semiconductor chips, which can power domestic appliances as well as advanced weaponry. Other examples range from high-strength carbon fibers (used both in sports equipment and in aerospace) to advanced optical fibers (used in telecommunications and missile guidance) or fermenters (used for vaccine production and the development of biological weapons).

Given their peculiar nature, trade in dual-use items is subject to security restrictions. These measures, known as export controls, take the form of licensing requirements: trade in controlled items is not prohibited *tout-court* but needs to be approved by competent authorities. Growing international tensions and the increasing military application of modern technologies have led to a vast expansion in controls, which now constitute a central tenet of security policy. The expansion was led by the United States, which imposed sweeping controls on semiconductors and chip-making equipment and steadily widened their scope, but has since been taken up by other major jurisdictions, notably China and the EU.

These developments have relevant implications for exporting firms. Although only a small share of export licenses is denied, compliance with export control regulations entails costs and delays. Securing an export license is not a mere administrative formality, but a complex process involving multiple corporate functions. Exporters must assess whether their products fall under the regulation, which requires detailed technical judgment (Chatelus and Heine, 2016). Firms must then assess possible end uses, and screen all parties involved against applicable controls, including extraterritorial ones. European business associations highlight that compliance with EU dual-use export controls can "overstrain the financial and human resources of many companies".[3] Noncompliance however can result in both criminal and administrative sanctions against the firm and its management. The tensions generated by the regime are indicative of the quest for managing heighted security concerns in open interconnected economies (European Commission, 2024).

Against this backdrop, systematic evidence on the trade effects of export controls remains scarce, especially for the structural, treaty-based controls that affect the bulk of dual-use trade. This paper addresses this gap by evaluating the impact of dual-use export controls on Italian exports. We consider this to be particularly valuable

---

[2] EU Regulation 821/2021 defines dual use items as "*items, including software and technology, which can be used for both civil and military purposes, and includes items which can be used for the design, development, production or use of nuclear, chemical or biological weapons or their means of delivery, including all items which can be used for both non-explosive uses and assisting in any way in the manufacture of nuclear weapons or other nuclear explosive devices*".

[3] Joint statement by European business on the recast of the EU Regulation on Dual-Use Export Controls, Business Europe, 2020.

evidence, given Italy's role as one of the world's largest exporters, with a strong manufacturing base that is deeply interconnected with the European and global economy.

Our empirical strategy exploits the staggered introduction of new products into the EU dual-use list. Over our sample period (2017-2023), EU Regulation 2021/821 requires that the list of export-restricted products is updated exclusively by transposing decisions adopted within multilateral export control regimes. This means that the timing of product additions to the list is largely driven by multilateral negotiations, not by EU-specific strategic or political considerations, which gives us a credible source of identification. Using firm-level Italian customs data, we estimate the short- and medium-term effects of licensing requirements on merchandize export flows via a Poisson pseudo-maximum likelihood implementation of the Extended Two-Way Fixed Effects (ETWFE) estimator of Wooldridge (2023, 2025).[4]

Our contribution is threefold. First, we provide the first causal evidence on the impact of EU dual-use controls. Our results show that EU controls leads to a short-term decline in exports of controlled items of approximately 10–20% in the first three quarters after listing; the effect fades over the medium term, with no significant long-run effect. Importantly, our estimates capture the effect of "structural" controls – that is, controls embedded in treaty-based and multilateral export-control regimes – rather than controls imposed as sanctions, which have been the primary focus of the existing literature.

Second, we show that firms are able to compensate for export losses by shifting to new markets. Decomposing the aggregate listing effect into intensive and extensive margins reveals that firms compensate for losses in existing markets by expanding exports to new destinations, particularly to "trusted", politically aligned partners. This pattern is consistent with the broader logic of export deflection documented in the sanctions literature, where political restrictions generate lasting reallocation of trade flows across destinations (Haidar, 2017; Crozet and Hinz, 2020).

Third, and closely linked to the last point, we show that the simplified licensing procedure for exporting to politically aligned countries helps firms cushion the impact of listing. Our results show that listing a product leads to a nearly 8 percentage points increase in the share of exports to trusted destinations relative to never-listed products, mostly driven by the United States.

Taken together, the results indicate that EU export controls generate temporary adjustment costs and a durable reallocation of trade toward politically aligned partners, without creating permanent aggregate barriers to trade. While the aggregate trade value recovers, the composition and geography of trade are permanently altered.

The remainder of the paper is structured as follows. Section 2 reviews the related literature. Section 3 describes the export control framework and its implementation in the EU. Section 4 presents data and descriptive evidence. Section 5 describes the empirical strategy and baseline results. Section 6 presents the robustness

[4] While our empirical analysis focuses on goods trade, export controls apply also software and technology transfers, including technical assistance, know-how, and other forms of intangible technology exchange. Our estimates should therefore be interpreted as lower bounds of the overall economic impact of controls, as they do not capture restrictions operating through services.

checks. Section 7 decomposes the results into intensive and extensive margins and documents the reorientation of exports toward trusted destinations. Section 8 concludes.

2. Literature

This paper contributes to a growing body of evidence on the tradeoffs between economic openness and security. In the widest sense, our work relates to the field of geoeconomics (Strange, 1988; Luttwak, 1990), namely the study of how international political rivalry shapes economic policy and outcomes (Mohr and Trebesch, 2024). More specifically it connects to the literature on economic statecraft, defined as the use of economic means to pursue foreign policy objectives (Baldwin, 2020), and on weaponized interdependence in the modern economy (Farrell and Newman, 2019).

While export controls have received extensive attention from the international relations and think tank community (Rasser, 2020; Gehrke and Ringho, 2023; Reinsch et al., 2024; Szczepański, 2025), the academic literature dedicated specifically to export controls remains more limited. The closest body of work is the literature on trade sanctions, which has documented their effects across a range of settings – including the impact of US and EU sanctions on Russia (Crozet and Hinz, 2020; Crozet et al., 2021; Görg et al., 2024; Kohl et al., 2024; Jäkel et al., 2024), of export sanctions on Iranian exporters (Haidar, 2017), and of bilateral embargoes on firm-level trade and domestic operations (Aytun et al., 2024). A consistent finding across this literature is that while sanctions reduce trade on targeted routes, firms partially offset losses through diversion to alternative markets, with broader effects on employment and supply chains.

The literature on export controls proper is more limited, and mostly focused on the recent wave of US semiconductor-related controls on China (Savini Zangrandi et al., 2026). This empirical literature can be organized into three strands.

The first examines the impact on the target country or firms. Hayakawa (2024) and Liu et al. (2024) find significant negative effects of US controls on Chinese imports of controlled products, with firm-level evidence pointing to lower profits and productivity alongside adaptation through increased R&D and patenting. Hu et al. (2024) and Alfaro et al. (2025) similarly document innovation responses to US end-user controls and Chinese rare-earth restrictions, respectively. The evidence is not uniform, however: Hayakawa et al. (2023; 2025) and Makioka and Zhang (2023) find mixed effects of Japanese controls on trade with China and South Korea.

The second strand focuses on extraterritoriality, documenting how US controls spill over to affect the exports of third countries such as South Korea and the Netherlands to China (Hayakawa et al., 2023; 2024; Kim and Cho, 2024).

The third strand studies the effects on sender countries. Crosignani et al. (2024) show that US firms with ties to sanctioned Chinese entities suffer significant losses in revenue, profitability and market capitalization, with

negative spillovers to non-designated partners. Moon and Jang (2023), by contrast, find that Korean export controls are associated with productivity gains.

Finally, it is worth noting the work of Alekseev and Lin (2026), which study the optimal export policy for dual-use goods in a model that embeds a military contest between states – to our knowledge the only theoretical contribution dedicated specifically to dual-use controls.

In this context, our paper contributes to the third strand of the literature, dedicated to the study of the impact of export controls on sender countries. This strand has so far focused on unilateral controls, often limited to semiconductors or rare earths rather than the full range of dual-use goods, while little attention has been devoted to the structural, treaty-based dual-use controls that operate continuously rather than in response to specific crises. Our paper fills this gap, providing the first causal, firm-level evidence on the trade effects of EU structural dual-use controls, using data on exporters from one of the Union's largest economies.

3. The export control framework, and its implementation

3.1. General framework

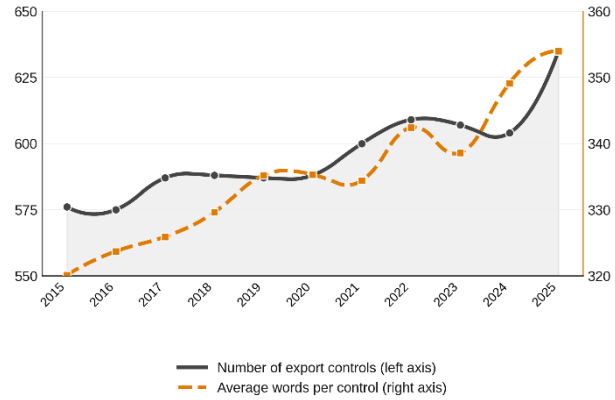
Modern-day export controls are rooted in a set of international treaties and multilateral arrangements aimed at containing the spread of potentially dangerous items and technologies (see Appendix A, Table 4). While these arrangements constitute the basis of most national controls, they are often supplemented with unilateral measures.

The past decade has seen a vast expansion in export controls across major jurisdictions. In the United States, over the last ten years, the number of items under US export controls grew by around 10 percent, with individual controls becoming increasingly complex (Figure 1). In 2020, China introduced a formal export control framework, subsequently expanded to cover advanced technologies, rare earth elements and critical materials, with notable extraterritorial provisions. A similar dynamic applies to EU export controls, governed by Regulation 2021/821: between 2012 and 2023, the number of controlled items increased by nearly 10 percent,[5] with roughly 300 CN8-level products added to the dual-use list (Figure 2). In 2025, the European Commission further expanded the scope of its controls beyond the prescriptions of the multilateral arrangements to which EU controls had traditionally been anchored.[6] Across the board, export controls are evolving from a peacetime tool of strategic containment to one of active coercion, deployed alongside economic sanctions.

[5] Identified by the number of entries under Annex 1 of the Regulation. Annex 1 of Reg. 821/2021 contains the list of dual use products identified through specific dual use product codes. In order to facilitate comparison with the customs nomenclature, the Commission also publishes a correlation table, which identifies CN codes correlated with each dual-use product. In line with the Commission, we refer to such products as dual use-related.

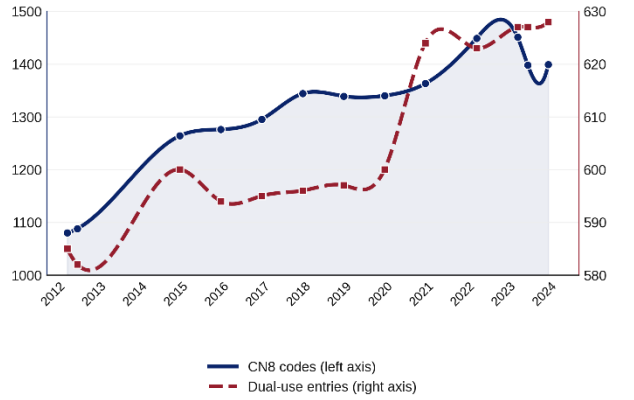
[6] EU consolidates export control powers to bypass Russia, Financial Times, 2025.

Figure 1: Items in the US export control list



Notes: "Number of export controls" counts the individual entries in the US Commerce Control List (each identified by an Export Control Classification Number, ECCN), while "average words per control" measures the mean length of each control entry, used as a proxy for the complexity of individual controls. Source: Code of Federal Regulation, authors elaboration.

Figure 2: Items in the EU dual-use list



Notes: The dual-use entries (right axis) count the items listed in Annex I of Regulation 2021/821 as updated by successive delegated acts; the CN8 codes (left axis) count the product codes appearing in the European Commission correlation tables that map customs nomenclature to those entries. Dates refer to the entry into force of the relevant delegated act updating the Annex.

3.2. The EU framework

The export of dual-use items from the EU is governed by Regulation 2021/821 (previously 428/2009, henceforth "the Regulation"). The Regulation creates a unified legal framework for all member states, while allowing additional national controls on non-listed items. Any item listed in the over 300 pages of Annex I of the Regulation requires an export license for destinations outside the EU.[7] Dual-use items span a wide range of products, from advanced electronics and cryptographic systems to chemical substances and aerospace components, reflecting the cross-cutting nature of the dual-use sector (see Appendix A, Table 5).[8]

The EU updates its dual-use control list on a yearly basis through delegated acts which incorporate changes agreed in international treaties and multilateral export control regimes in the year before.[9] As the classification provided in Annex 1 of the regulation can be difficult to navigate, the European Commission also compiles a "correlation table" (generally published a couple of months after the delegated act) which maps customs (CN) codes with the dual use item codes contained in the regulation.

The Regulation requires the export of controlled items to be approved by National Authorities, which issue export licenses. The EU dual-use system provides several types of export authorizations with differing levels of scrutiny. General authorizations facilitate exports to low-risk markets without requiring a license for each

[7] Dual-use items can move freely within the EU's single market, except for a subset of very sensitive items listed in Annex IV of the Regulation.

[8] Dual-use items are not limited to the list set out in the Regulation: items not explicitly listed can also be controlled if they may be used for weapons of mass destruction or in violations of arm embargoes ("catch-all" clause). This provision imposes a strong due diligence obligation: exporters are expected to screen end users and end-uses carefully and watch for red flags, indicating potential misuse.

[9] The timeline is typically the following: multilateral control regimes update their control lists toward the end of a given year. The subsequent autumn, the European Commission issues its draft delegated acts for potential objections by the Council. Absent objections, the delegated act are formally adopted by the end of the year.

transaction. The most important one is the General Export Authorization (GEA) EU001, which allows exports of most dual-use items to a set of allied, "trusted" countries (including all G7 members, plus Australia, Iceland, Liechtenstein, New Zealand, Norway, and Switzerland). Italy grants a similar treatment to Argentina, Turkey, and South Korea through a national general license.[10] All other situations require case-by-case authorization, either an individual license for a single transaction or a global license for ongoing business with multiple customers.[11] In Italy, licenses are issued by the *Unità per le Autorizzazioni dei Materiali d'Armamento* (UAMA) of the Ministry of Foreign Affairs.

The EU dual-use regulation places significant obligations on exporters. The compliance process starts with product classification, typically by consulting EU correlation tables. These are only indicative: firms must still verify technical specifications to confirm dual-use status (Chatelus and Heine, 2016). If a product's CN code appears on the table, exporters need either a license or a declaration that it falls outside the Regulation. Ascertaining dual-use status can be complex, requiring detailed understanding of the controls, including their exceptions and exemptions, and of the item's technical specifications; firms often rely on external consultants to interpret these thresholds.[12] Beyond product-level checks, firms must obtain and screen end-user documentation from foreign clients, including verifying that counterparties are not sanctioned and that the stated use is legitimately civilian. Legal teams must update contractual clauses (typically to protect the firm against risks from license delays or denials) and prepare documentation for the licensing authority, while maintaining records of all dual-use transactions for at least five years.[13]

As non-compliance can result in security breaches, penalties are severe. These include fines, trade restrictions, and criminal liability, depending on the jurisdiction. The Italian law prescribes criminal penalties of 2 to 6 years imprisonment for unauthorized exports or false declarations, alongside administrative fines ranging from €25,000 to €250,000.[14]

---

[10] For detail on the implementation of Reg. 821/2012 across member states see: Information Note – Regulation (EU) 2021/821 of the European Parliament and of the Council setting up a Union regime for the control of exports, brokering, technical assistance, transit and transfer of dual-use items: Information on measures adopted by Member States in conformity with Articles 4, 5, 6, 7, 8, 9, 11, 12, 22 and 23 (OJ C, C/2024/5881, 02.10.2024, ELI: http://data.europa.eu/eli/C/2024/5881/oj.

[11] Licenses are typically processed within 10 to 30 days, but processing times can increase significantly – up to the statutory maximum of 180 days – if documentation is incomplete or inaccurate, or if additional checks are required.

[12] An additional layer of complexity lays in the fact that a product whose CN code does not appear in the correlation tables may still be subject to dual-use controls if one of its key components is listed as a dual-use item.

[13] The European Commission strongly recommends the adoption of Internal Compliance Programs (ICPs), structured policies and procedures for managing export controls, which are mandatory in Italy for global licenses. A well-implemented ICP reduces the risk of violations and can speed up licensing, as authorities tend to prioritize firms with robust compliance systems.

[14] Administrative fines also apply for record-keeping failures, or for failing to file mandatory reports on exports under general licenses. Notably, firms are required to submit regular reports of their exports of licensed products, even in periods with nil exports. Beyond legal sanctions, export violations can damage an exporter's relationships with licensing authorities (potentially hampering the ability to obtain export licenses in the future), as well as with banks and other counterparties.

4. Data and descriptive evidence on dual-use sector

Our analysis combines information on EU dual use export control classifications with firm-level customs data for Italy.

Dual use-related products are identified by the EU Commission correlation tables,[15] which map items listed in Annex I of the Regulation to Combined Nomenclature (CN) product codes, namely the classification used in EU customs declarations. The tables provide the official reference to determine which product codes may fall within the scope of the Regulation. This mapping is necessarily approximate: correlation tables serve only as an initial screening tool, and actual applicability depends on technical product specifications that must be verified on a case-by-case basis (see Section 3). Notably, just a fraction of dual use-related items turns out to be genuinely dual-use and require an export license. Licensed dual use trade only accounts for 2 percent of total extra-EU exports.[16] The wide gap between dual use-related and effectively dual use products is a key feature of the EU export controls regime, which requires all firms to evaluate whether their products fall under the regulation, regardless of whether a licence is ultimately required. Firms in this position face important decisions, and their behaviour is precisely the object of our analysis.

Italian exports are identified by administrative records collected by the Italian Customs and Monopolies Agency, covering the universe of Italian trading firms engaged in extra-EU trade, at quarterly frequency from 2017Q1 to 2023Q4. The data includes the firm's tax ID, the CN8 product code, the partner country, and the value of exports.[17]

The combination of these datasets shows three important characteristics of dual use-related trade.

First, dual use-related trade is large. Between 2021 and 2023, nearly 30 percent of Italy's extra-EU exports involved dual use-related products. More than half of exporters (106,000 firms) ship at least one dual use-related product. Such products are not confined to a narrow set of specialised industries: while the machinery and electronics sector accounts for half of this value, controlled products are also present across chemicals, mineral products, transport equipment and most other major sectors, making compliance a concern for firms across the economy (Figure 3).

Second, dual use-related trade is growing fast. Between 2017 and 2023, Italian dual use-related exports grew at an average annual rate of about 8 percent – nearly twice as fast as total goods exports – increasing from 40bn to more than 70bn euros (Figure 4.a).

---

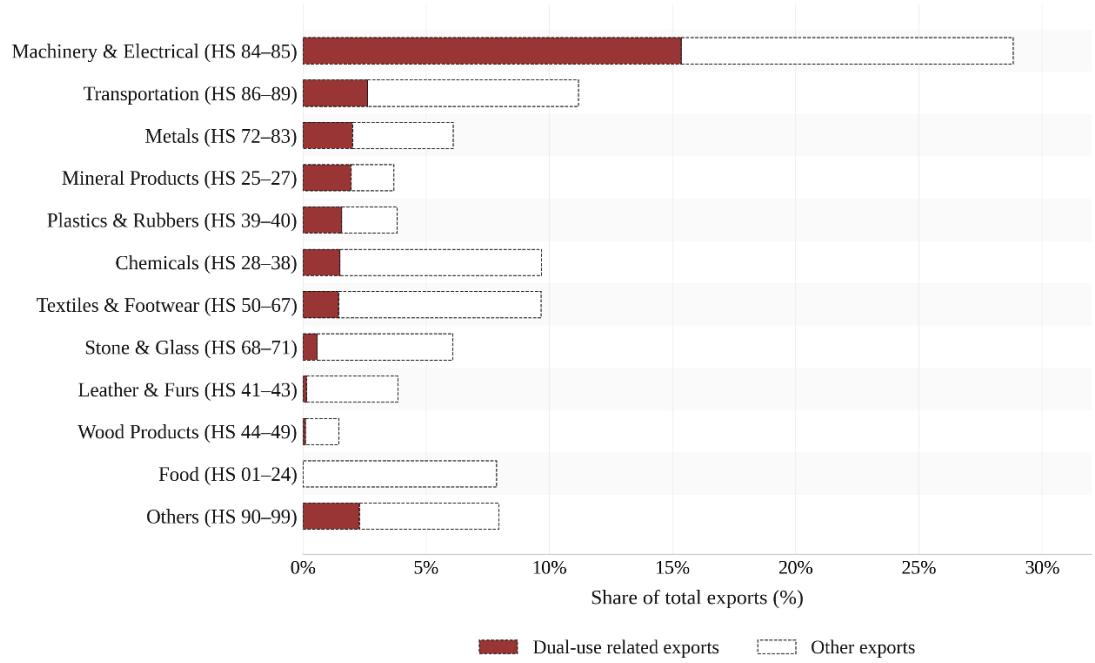
[15] Correlation tables are available at: https://circabc.europa.eu/ui/group/0e5f18c2-4b2f-42e9-aed4-dfe50ae1263b/library/c3d06bd7-6ef0-4771-bbd7-f92b976ae9a0.

[16] EC report on the implementation of Regulation (EU) 2021/821 on dual-use goods and technologies.

[17] CN8 codes are updated annually to reflect changes in the Combined Nomenclature. To ensure comparability over time, we use the CN8 correlation tables provided by Eurostat to harmonize all codes to the 2017 classification.

Third, dual use related trade has a strong geopolitical component. Following the Di Casola et al. (2023) index of geopolitical affinity,[18] we note that export growth over the same period is primarily driven by Western and Neutral blocks, which together account for the 75% of dual use related export value (the US alone accounts for 20 percent). Exports to the Eastern bloc grew in 2018 but declined thereafter, reverting to 2017 levels by 2023. In terms of licenses, the share of dual-use exports to "trusted" countries (covered by EU GEA 001 or the Italian General Licence) accounted for around 25 percent of dual use-related trade in 2017 and increased significantly over the period, both in absolute value and relative to flows to other partners, reaching almost 30 percent in 2023 (Figure 4.b).

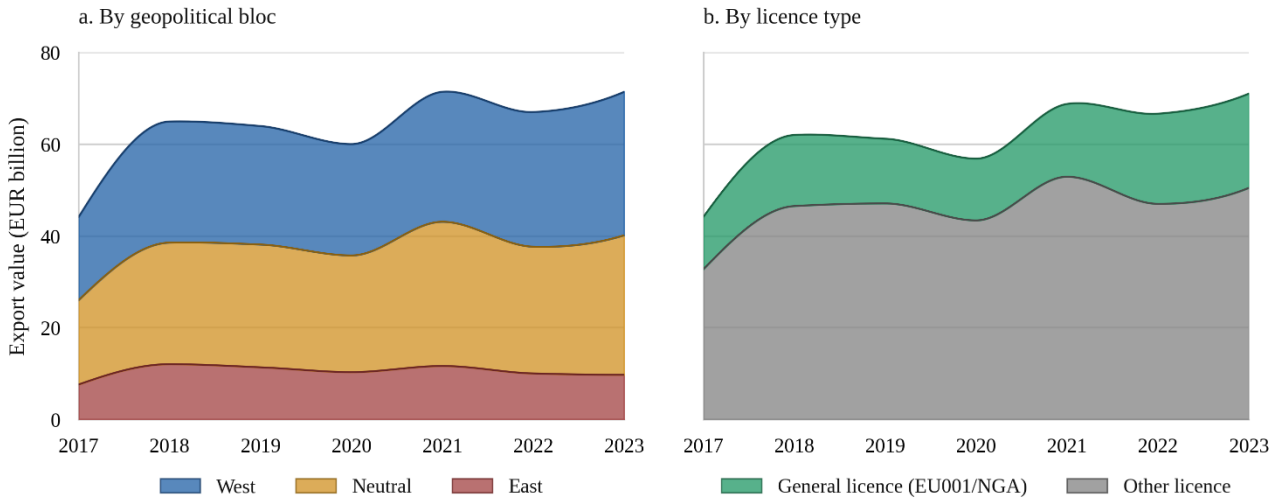
Figure 3: Dual use-related exports by sector (% total exports)



Notes: The figure shows the share of total Italian exports accounted for by dual-use related products (red) and other products in the same sector (white), over total Italian exports in 2021–2023. Dual use-related exports refer to CN8 product codes appearing in the EC correlation tables linking the CN nomenclature to the control list in EU Regulation 2021/821.

[18] The index measures a country's alignment with the US versus China/Russia based on sanctions, military imports, official lending, and UN voting. Countries with an index below 0.25 are assigned to the West, those above 0.75 to the East, and the remainder to the Neutral bloc. For countries not covered by the index, we use the country list provided by Capital Economics (available at: https://www.capitaleconomics.com/fracturing-dashboard), which shows near-perfect correspondence in the overlapping sample. Overall, the classification includes 81 Western countries, 95 Neutral countries, and 56 Eastern countries.

Figure 4: Italian dual use- related exports by geopolitical bloc and licence type



Notes: Panel a. shows the value of Italian exports of dual-use products disaggregated by geopolitical bloc of the destination country. Countries are classified as West (NATO members and close allies), East (Russia, China, and other adversarial states), and Neutral (all remaining countries), following the classification of Di Casola et al. (2023). Panel b. disaggregates dual-use exports by licence regime: 'General licence (EU001/NGA)' refers to exports covered by the EU general export authorisation EU001 or by Italian national general authorisations; 'Other licence' includes exports under individual licences or other general authorisations. Export values are expressed in EUR billion at current prices.

5. Baseline estimates

This section describes the empirical strategy and the results for the baseline estimates on the effect of dual-use listing on firms' exports.

5.1. Empirical strategy

To assess the impact of export controls on trade in dual-use related items, we implement a staggered difference-in-differences (DiD) design, which exploits the staggered timing with which products are added to the EU dual-use control list.

Our identification strategy relies on the timing and exogeneity of the updates to Annex I of the Regulation, which transposes into EU law the control lists adopted by multilateral export control regimes and international treaties. Under Article 17 of the Regulation, the European Commission is empowered to update Annex I exclusively to reflect changes adopted by these regimes. Thus, the EU cannot expand the Annex unilaterally, regardless of strategic or political circumstances, like geopolitical shocks such as Russia's annexation of Crimea in 2014 and its full-scale invasion of Ukraine in 2022.[19]

Between 2017 and 2023, 192 new products were added to the EU dual-use list and never removed, forming 8 treatment cohorts, where $g \in G = \{t_0, \dots, T\}$, where $t_0 = 2017$ Q4 and $T = 2023$ Q4 (Figure 5). Cohorts are defined by the quarter in which a product is included in the EU dual-use control list. Of the 192 newly listed

[19] As noted in section 1, in 2025 the European Commission assumed a more proactive role in managing export controls, enabling the EU to enforce restrictions on sensitive technologies not listed in multilateral agreements. However, since our analysis only considers updates up to 2023, all Annex I updates we exploit in our design are exogenously constrained by multilateral decisions.

products, 83 belong to machinery and electrical equipment (Figure 6) – including marine diesel engines, metal-working machines capable of manufacturing weapons components, industrial furnaces used in nuclear material processing, and electronic components for military systems. A further 28 are chemical products, covering precursors to explosives agents, as well as ammonium nitrate and fluorine gas used in uranium enrichment. The remainder includes precision instruments (HS 90-99, like military-grade optical lenses and lasers), metals (such as specific alloys used in nuclear reactor vessels and centrifuges), and textiles (almost exclusively bulletproof garments).

Figure 5: Newl dual-use products by treatment cohort

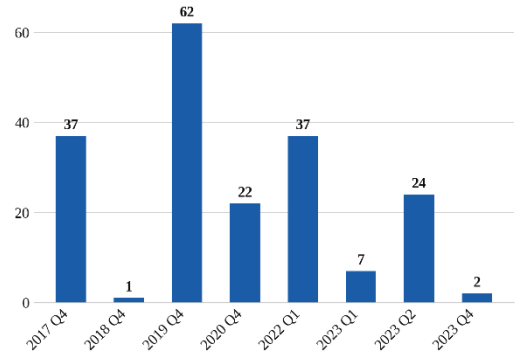
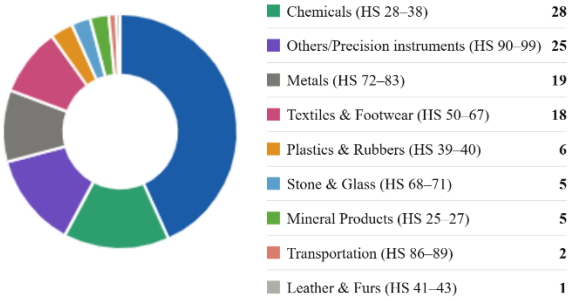


Figure 6: Sectoral composition of new dual-use products



Notes: figures report CN8 product codes added to the EC dual-use correlation tables between 2017 and 2023, expressed in the 2017 CN nomenclature. Codes added to the dual-use list but subsequently removed before end-2023 are excluded.

We estimate treatment effects using a Poisson pseudo-maximum likelihood (PPML) version of the Extended Two-Way Fixed Effects (ETWFE) model proposed by Wooldridge (2023, 2025). The ETWFE framework extends the conventional two-way fixed effects DiD model by allowing the average treatment effect on the treated (ATT) to vary flexibly across treatment cohorts g and over time t. This flexible specification avoids the well-documented problems of "bad controls" and negative weighting that arise in traditional TWFE estimators under staggered treatment adoption and heterogeneous treatment effects. Our main specification is as follows.[20]

$$\mathbb{E}[Y_{p,t,d,f} \mid \xi_{p,d,f}, \xi_t] = \exp\left\{\alpha + \xi_{p,d,f} + \xi_t + Trusted_d\,\xi_t + \sum_{g\in G}\sum_{\substack{t \\ t-g\neq -1}} \theta_{g,t}\, D_{g,t} + \sum_{g\in G}\sum_{\substack{t \\ t-g\neq -1}} \beta_{g,t}\, D_{g,t}\, Trusted_d\right\} \quad (1)$$

The outcome variable in equation (1), $Y_{p,t,d,f}$, is the value of exports from Italian firm f of product p (at the CN8 level) to non-EU destination d, observed at quarter t. The unit of observation is the firm-product-destination combination. The specification includes time fixed effects (ξ_t) and firm–product–destination fixed effects ($\xi_{p,d,f}$), which absorb all time-invariant heterogeneity at the unit level. The treatment cohort variable G_p is defined as the quarter in which product p is first added to the EU dual-use list. Products never listed are

[20] Regression is estimated using the jwdid STATA command proposed by Yotov, Nagengast and Rios-Avila (2026). Note that the ETWFE specification is numerically equivalent to an OLS regression with cohort fixed effects ξ_g instead of firm–product–destination fixed effects $\xi_{p,d,f}$, as shown in Wooldridge (2025).

assigned $G_p = 0$ and constitute the control group. The indicator $D_{g,t}$ takes value 1 if the product is first treated in period $g \in G = \{t_0, \dots, T\}$ (i.e., if $G_p = g$) and current time $= t$, and 0 otherwise. We adopt a leads-and-lags (event study) specification, including indicators $D_{g,t}$ for each cohort g at all relative periods $(t - g)$ except for the period immediately before treatment (i.e., when $(t - g) = -1$), which serves as the reference.

Our parameters of interest are the cohort- and time-specific treatment effects $\theta_{g,t}$. Under the Poisson specification, $\theta_{g,t}$ is naturally interpreted as a log change (semi-elasticity). Accordingly, $e^{\theta_{g,t}} - 1$ gives the percentage change in exports induced by being added to the dual-use list for cohort g at time t, which can be interpreted as the percentage average treatment effect on the treated (ATT). Formally,

$$\theta_{g,t} = \log\mathbb{E}[Y_t(g) \mid G = g] - \log\mathbb{E}[Y_t(0) \mid G = g]$$

This framework allows for heterogeneity in treatment effects across products, firms, destinations and time, implying that the impact of being added to the dual-use list may differ across cohorts and evolve dynamically after treatment. We allow treatment effects to vary by destination group by interacting the cohort-time treatment indicators $D_{g,t}$ with a dummy variable $Trusted_d$, which captures differences in export control regimes across destination countries. EU export control legislation provides general export authorizations that relax licensing requirements for less sensitive dual-use items exported to a set of trusted countries, comprising all G7 members and additional destinations, further expanded under Italian general. Accordingly, $Trusted_d$ equals one for trusted destinations (namely, Argentina, Australia, Canada, Iceland, Japan, Liechtenstein, New Zealand, Norway, Korea, Switzerland, Turkey, and the United States) and zero otherwise. The interaction allows us to assess whether the impact of being added to the dual-use list is attenuated for exports to countries benefiting from more permissive licensing regimes. The term $Trusted_d\, \xi_t$ allows for flexible destination-group-specific time effects. Standard errors are clustered at the unit level.

The average effect of being treated for s periods is obtained as a weighted average of the cohort-time effects $ATT(g, g + s)$, where the weights are given by the relative size of each cohort:

$$ATT(s) = \sum_{g:\, g+s \le T} \frac{n_g}{\sum_{h:\, h+s \le T} n_h} \cdot ATT(g,\, g + s)$$

where T is the final period in the sample window. The overall ATT, across all post-treatment periods and for all treated cohorts, is obtained as a weighted average of the $ATT(g, t)$'s, where weights account for both cohort size and the number of periods each cohort contributes:

$$ATT = \sum_g \lambda_g \sum_t ATT(g, t), \qquad \lambda_g = \frac{n_g}{\sum_h n_h\, T_h}$$

where n_g is cohort size and T_g is the number of post-treatment periods cohort g is observed.

Our estimation sample consists of all firm-product-destination combinations for which at least one export transaction is recorded over the 2017–2023 period. For each such combination, we construct a balanced quarterly panel covering all 28 quarters from 2017Q1 to 2023Q4, assigning a value of zero to quarters with no recorded exports.[21] The final sample consists of 114 million firm-product-destination-quarter observations, 7,679 products, of which 192 (2.5%) are newly listed dual-use items, and 180,000 firms.[22] As a robustness check, we verify that our main results are not driven by the smallest trading relationships by re-estimating the baseline specification after dropping firm-product-destination combinations whose total export value over the entire sample period falls below €2,000 and €5,000 respectively.

Because our panel is balanced, the ETWFE estimator is numerically equivalent to the group-time average treatment effect estimator proposed by Callaway and Sant'Anna (2021). We adopt the Wooldridge formulation because it naturally extends to nonlinear models, including Poisson regression (Wooldridge, 2023, 2025).[23] Identification relies on two standard assumptions: no anticipation and parallel trends. The no-anticipation assumption requires that, for each cohort g, expected export outcomes prior to treatment are unaffected by future listing decisions. This assumption is plausible given most firms do not directly monitor regulatory developments within multilateral export control regimes. The parallel trend assumption requires that, absent the inclusion in the dual-use list, the average export outcome for treated and never-treated products would have followed the same trajectory. As we are using a Poisson specification, this means no systematic difference in percentage changes between treatment and control (i.e. parallel growth rates), measured relative to the last pre-treatment period.

5.2. Baseline results

Column (1) of Table 1 reports the baseline results, presenting the average treatment effect on the treated (ATT) across all cohorts and over the entire post-treatment period.

Overall, we do not find evidence of a significant long-run reduction in exports of newly controlled products, as the estimated ATT is small in magnitude and statistically insignificant. When we distinguish between destination groups (trusted vs non-trusted countries), the ATT is even marginally positive for trusted countries.

---

[21] We restrict the analysis to products added to the list and never subsequently removed, and we exclude products that are already listed at the beginning of the sample period. Finally, to isolate the effects of export controls under normal regulatory functioning, we exclude from the sample destinations subject to comprehensive sanctions regimes, namely Russia, Iran, North Korea, Belarus and Myanmar. We identify sanctioned destinations using the European Commission's list of countries subject to restrictive measures on dual-use exports, available here: https://www.sanctionsmap.eu/#/main. For completeness, we report in Appendix A the estimates obtained when these countries are included; the results are largely unchanged.

[22] The distribution of total export value per unit (firm-product-destination combination) over the full sample period is highly right-skewed, as is typical of firm-level trade data.

[23] PPML has three properties that make it particularly suited to our setting. First, it accommodates the large share of zero observations that arises mechanically from the balanced panel structure (more than 75% of observations record zero exports). Second, it is robust to the severe right-skewness of the trade value distribution, as unlike OLS on log-transformed outcomes PPML does not require the error term to be homoskedastic and yields consistent estimates under a broad class of distributional assumptions (Silva e Tenreyro, 2006). Third, the Poisson specification delivers treatment effects that are naturally interpreted as percentage changes. Relative to level-based specifications, the Poisson model therefore implies a common-trends assumption in growth rates rather than in levels, which is a more plausible assumption for outcomes like exports.

When we re-estimate the specification after excluding the smallest and most marginal flows (Table 1, columns (2) and (3)), the results remain broadly stable and consistent with the baseline. The negative (though insignificant) effect for non-trusted destinations becomes slightly larger in magnitude as economically negligible trading relationships are progressively excluded.[24]

Table 1: Overall ATT

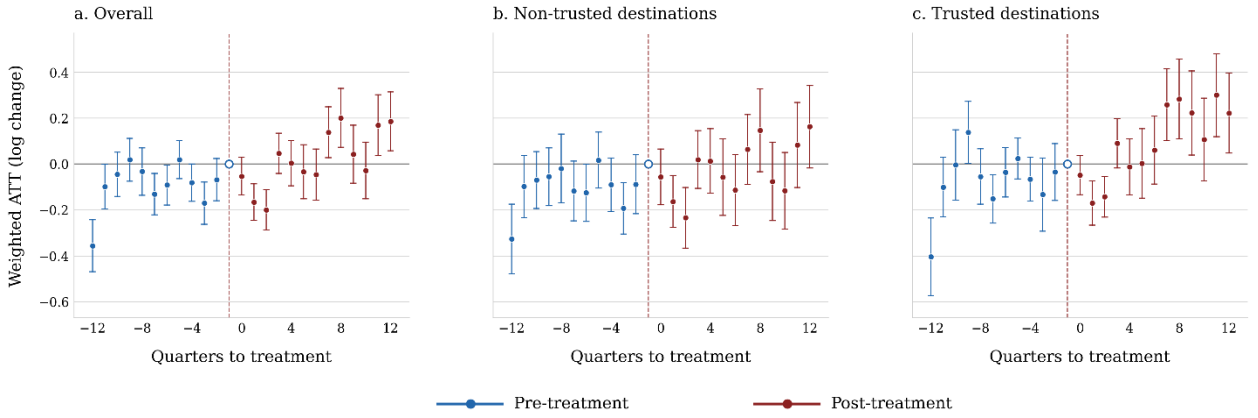
	(1)	(2)	(3)
	All units	*Excl. units with <2,000 EUR exports*	*Excl. units with <5,000 EUR exports*
Overall ATT	−0.001	−0.005	−0.011
	(0.045)	(0.046)	(0.046)
ATT - Trusted destinations	0.111*	0.109*	0.107*
	(0.063)	(0.063)	(0.063)
ATT - Non-trusted destinations	−0.070	−0.076	−0.086
	(0.062)	(0.063)	(0.064)
Observations	114,126,516	94,804,416	70,607,600
Units	4,075,947	3,385,872	2,521,700
PPML-ETWFE in all specifications.			

Notes: The dependent variable is the value of Italian exports. Aggregated ATTs are computed as weighted averages of cohort–time ATTs via estat simple (post-jwdid command in STATA), with unconditional standard errors. Column (1) uses the full estimation sample. Column (2) excludes units with total exports below €2,000 over the full sample period. Column (3) excludes units with total exports below €5,000. Trusted is a dummy equal to one for friendly destination countries. Standard errors in parentheses are clustered at the unit level. * p<0.10, ** p<0.05, *** p<0.01.

When looking at the event-study estimates, however, we observe that in the quarters immediately following listing exports of newly controlled products fall by about 10 to 20 percent (Figure 7.a).[25] This decline is visible both for non-trusted (Figure 7.b) and trusted destinations (Figure 7.c), but it is stronger for non-trusted. The effect is concentrated in the first 3 quarters after listing (9 months), as firms adjust to compliance requirements, and then tends to fade. The first quarter corresponds to the formal update of the Regulation. While some firms, particularly larger or more specialized exporters, may closely monitor regulatory changes and react immediately, many firms might rely on correlation tables updates (typically occurring with some delay after the regulatory update) to identify potentially controlled products. Consistent with this, the event-study estimates show a relatively smaller effect in the quarter of the regulatory update and a stronger decline in the subsequent two quarters. This suggests that while a subset of firms adjusts promptly, most of exporters begin to react only after the updated classification becomes operational.

[24] Results are also robust to the exclusion of smaller and noisier cohorts, namely 2018Q4 and 2023Q4, from the sample.
[25] For the event study estimates when excluding marginal trading relationships see Appendix A.

Figure 7: Effect of dual-use listing on exports – event study estimates



Notes: Event-study estimates of the effect of dual-use listing on Italian export values. ATTs are aggregated as weighted averages of cohort time specific estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Panel a. reports the overall ATT across all non-sanctioned destinations. Panel b. restricts to non-trusted destinations; panel c. to trusted destinations (Argentina, Australia, Canada, Iceland, Japan, Liechtenstein, New Zealand, Norway, Korea, Switzerland, Turkey, and the United States). Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period (t-g= 1).

To assess whether a particular sector drives the short-run export decline, we replicate the baseline specification separately for Machinery and Electrical Equipment (HS 84–85) - the largest group of newly listed products - and for all remaining sectors combined (see Table 6 in Appendix B). The overall ATT in the first three quarters following listing is negative and significant for Machinery and Electrical Equipment, while it is negligible and statistically indistinguishable from zero for all other sectors. The export decline in the first 3 quarters is therefore primarily driven by Machinery and Electrical Equipment, consistent with the greater technical complexity of compliance requirements for these products, which typically require detailed technical specification assessments and demanding end-user screening. Within this sector, the decline is concentrated in flows to non-trusted destinations, for which no general licences are applicable, while the effect for trusted countries is smaller and not significant. The null result for the remaining sectors suggests that for less technically complex products (such as basic chemical compounds or textiles) dual-use listing does not generate significant short-run trade disruptions.

6. Robustness checks

6.1. Propensity score matching and entropy balancing

To strengthen the baseline DiD analysis and mitigate concerns related to non-random treatment assignment across products, we complement our empirical strategy with a propensity score matching (PSM) procedure at the product level. Products are matched on two pre-treatment characteristics that are strongly related to both export performance and the probability of being included in the dual-use list: technological intensity and product complexity. Technological intensity is captured by a binary indicator for high- and medium-high-

technology manufacturing sectors, following the Eurostat classification,[26] while product complexity is measured by the Product Complexity Index (PCI) from the Observatory of Economic Complexity.[27] Propensity scores are estimated using a logit model in which an ever-treated indicator is regressed on these covariates. Products are classified as ever-treated if they are added to Annex I at any point during the sample period (i.e., $G_p > 0$), and as never-treated otherwise. The Eurostat technology classification is time-invariant, while PCI values refer to the pre-sample year 2016.[28] The estimated coefficients indicate that both PCI and high-tech status are strong and statistically significant predictors of treatment, confirming their relevance for the selection process (Table 2).

Table 2: Propensity score estimation: logit results

	(1)
	P(dual-use)
Product Complexity Index	0.626***
	(0.122)
High-tech share	1.197***
	(0.203)
Constant	−4.249***
	(0.152)
Observations	5,480
Pseudo R^2	0.098
Estimator	Logit

Notes: The table reports coefficients from the logistic regression used to estimate the propensity score for matching. The dependent variable is a binary indicator equal to one if the product is ever added to the EU dual-use control list. The same propensity score is used across all matching specifications (kernel, NN5, NN10). Standard errors in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

We implement PSM using a kernel matching estimator and nearest-neighbor matching. Kernel matching (with an Epanechnikov kernel) assigns weights to control products based on their distance in propensity score space, using all controls within the common support region. Nearest-neighbor matching pairs each treated product with its n closest controls in propensity score space, using a more restricted set of controls; we report results for $n = 5$ (NN5) and for $n = 10$ (NN10). Prior to matching treated and control products differ substantially along both metrics (PCI and high-tech share), with normalized differences well above the conventional threshold of 0.25 that signals imbalance (Rubin, 2001; Stuart, 2010). After matching, covariate balance improves markedly across all specifications, and differences in covariate means between treated and control

[26] Available at: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Glossary:High-tech_classification_of_manufacturing_industries. Eurostat defines high-technology and medium-high-technology sectors at the NACE Rev. 2 level. We convert NACE codes into the CN8 product nomenclature using the official Eurostat correspondence table.

[27] PCI data are available at: https://oec.world/en/rankings/pci/hs6/hs02?tab=ranking

[28] When the 2016 value is not available for a specific product code, we use the most recent available prior-year value to limit the number of missing observations. Products for which no PCI value can be retrieved are excluded from the matching sample, resulting in 2,204 products excluded out of 7,684.

groups are not statistically significant (Table 3). The resulting matched sample is then used to re-estimate the DiD model.

Table 3: Covariate balance before and after propensity score matching

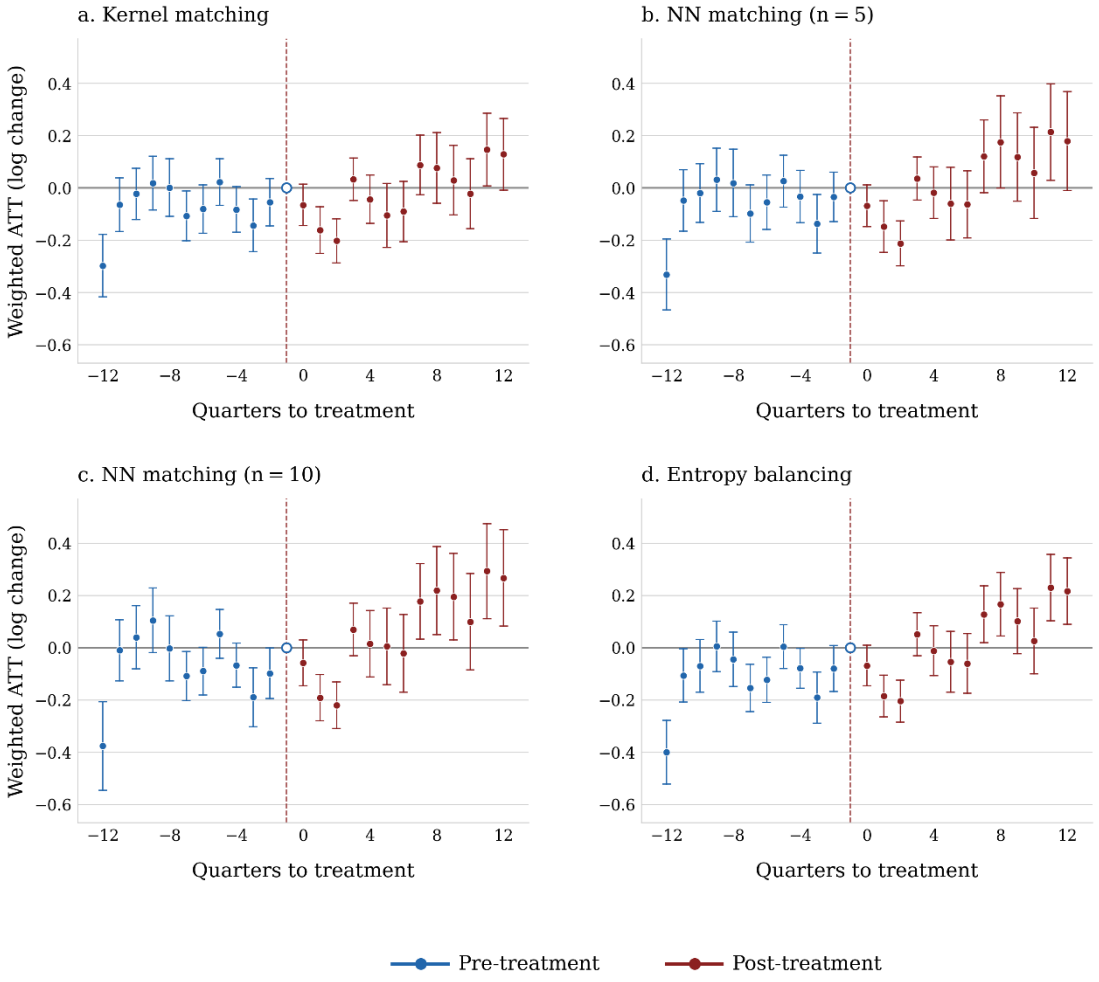
	Treated	Control (unmatched)	Control Kernel	Control NN (n = 5)	Control NN (n = 10)
Means					
Product Complexity Index	0.54	−0.21	0.46	0.50	0.49
High-tech share	0.73	0.31	0.74	0.75	0.75
Standardised difference					
Product Complexity Index	–	0.78	0.05	0.04	0.05
High-tech share	–	0.91	−0.05	−0.05	−0.05
Treated units	161		156	161	161
Control units		5,319	2,084	592	980

Notes: The table reports covariate balance before and after propensity score matching. Treated units are products added to the EU dual-use control list; control units are never-listed products. Means are reported for each group before matching and for matched control units under each matching algorithm. The standardised difference is computed as the difference in means divided by the pooled standard deviation of the two groups. Kernel matching uses an Epanechnikov kernel. NN matching uses nearest-neighbour matching with 5 and 10 neighbours respectively.

We complement these robustness checks with entropy balancing as an additional verification. Entropy balancing (Hainmueller, 2012) reweights control products so that the weighted covariate means of the control group exactly match those of the treated group. Unlike PSM, it does not discard any control unit and requires no choice of kernel or bandwidth, making the resulting balance less sensitive to methodological choices.

When we re-estimate the effects using propensity score matching (PSM) and entropy balancing, the main results remain unchanged. The event-study estimates all display a negative and statistically significant drop in exports in the first three quarters following listing, followed by a gradual recovery (Figure 8). The short-run decline is of similar magnitude to that estimated in the baseline specification, suggesting that the result is not driven by differences in pre-treatment characteristics between treated and control products. Over the entire period, the ATT becomes insignificant across all specifications (Table 7 in Appendix C), confirming the absence of a permanent effect on aggregate export values. All specifications except kernel matching confirm the negative sign for non-trusted destinations and the positive sign for trusted found in the baseline. Under kernel matching the coefficient for trusted destinations turns slightly negative, though insignificant.

Figure 8: Effect of dual-use listing on exports – matching results



Notes: Event-study estimates of the effect of dual-use listing on Italian export values estimated with propensity score matching and entropy balancing weights. ATTs are aggregated as weighted averages of cohort–time specific estimates with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Panel a. reports results under kernel matching (Epanechnikov kernel); panel b. under nearest-neighbour matching with n = 5 neighbours; panel c. under nearest-neighbour matching with n = 10 neighbours; panel d. under entropy balancing. The propensity score used in panels a.–c. is estimated via logistic regression on Product Complexity Index and high-tech share (see Table 2); entropy balancing in panel d. directly reweights the control group to match the means of the same covariates among treated products.. Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period (t−g = −1).

6.2. No-anticipation assumption

To assess the validity of the no-anticipation assumption, we conduct two additional robustness exercises. First, we redefine the treatment date as four quarters prior to the official regulation update, reflecting the fact that EU legislation may take up to one year to incorporate additions decided at the multilateral level. If firms do not generally anticipate regulatory changes, this redating should not substantially alter the ATT estimates. Even under this conservative re-dating, the short-run export decline remains negative and significant in the quarters corresponding to the actual listing date, and becomes small and insignificant thereafter, confirming the dynamics observed in the baseline (Table 8 in Appendix C).[29]

[29] For the event study estimates see Appendix C.

Second, we exclude firms that are likely to be large and actively involved in the dual-use sector (and therefore more likely to anticipate regulatory changes) by removing firms affiliated with the two main industry associations for which membership lists are publicly available, namely AIAD (Italian Federation for Aerospace, Defense and Security Industries) and UCIMU (Italian Association of Machine Tool, Robots and Automation Manufacturers). Also in this case, the baseline estimates are not similarly substantially altered (Table 8 in Appendix C). This suggests that neither anticipatory behavior nor the presence of large and sophisticated exporters drives the main findings.

7. Other results: decomposing intensive and extensive margins, and export reorientation

7.1. Intensive and extensive margins

To understand the channels through which export controls influence trade, we separately examine the intensive and extensive margins of firm-level exports.

The intensive margin captures the effect on the value of exports along existing trade relationships. To isolate this, we estimate the staggered DiD specification described in equation (1), restricting the sample to firm-product-destination combinations that report strictly positive exports prior to the product's inclusion in the dual-use list. This allows us to assess how listing affects existing export relationships, barring the creation of new ones.

The extensive margin instead captures changes in the probability that a firm exports a given product to a new destination. We retain the difference-in-differences design of equation (1), but redefine the outcome as a binary indicator equal to one if firm f records strictly positive exports of product p to destination d in quarter t, and zero otherwise. Because the outcome is binary, we estimate a linear probability model by OLS, so that the conditional mean is linear rather than exponential. The coefficients $\theta_{g,t}$ are accordingly interpreted as the percentage-point change in the probability of an active export relationship induced by listing.

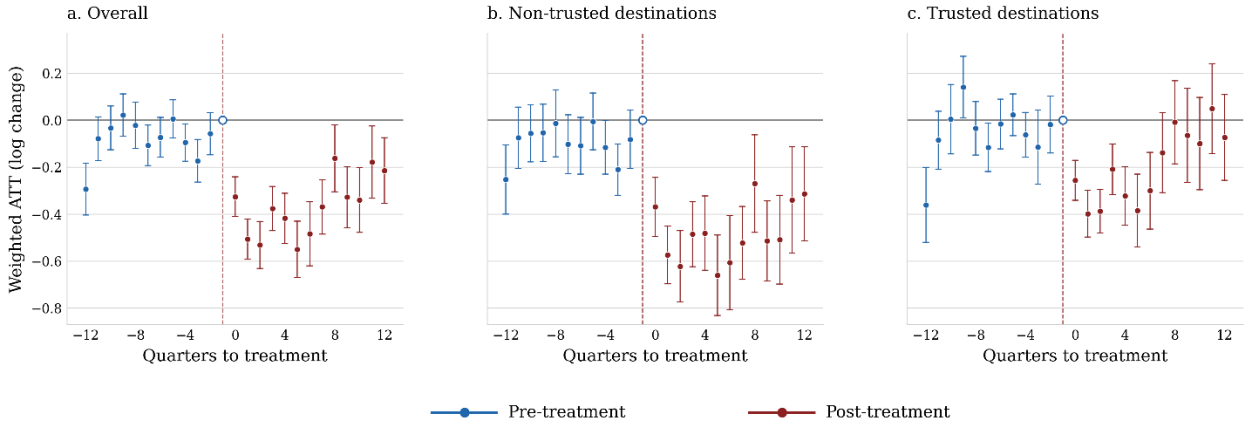
Table 8 presents the results for the intensive margin, estimated via Poisson PPML on the subsample of firm–product–destination combinations with strictly positive exports in the pre-listing period. The baseline results indicate that dual-use listing reduces export values along pre-existing trade routes by 34% on average, and this effect is persistent and does not recover over the post-treatment window, as shown by the event-study estimates (Figure 9a.).[30] There are significant differences for flows towards trusted compared to non-trusted destinations. For non-trusted partners the effect is larger in magnitude and more persistent: export values fall by 31% on average in the listing quarter, than a bit more (-44%) in the first quarter after listing (the one in which correlation tables are updated), and remain in the range of −24% to −48% throughout the post-treatment period (Figure 9.b). Instead, for trusted destinations the negative effect is lower in magnitude, with coefficients

[30] Cohort 2018Q4, which includes just one designated product, is excluded from all these specifications due to separation problems in the Poisson likelihood which prevent convergence.

roughly half the size of those for non-trusted destinations in the short run, and converging toward zero in the medium term (Figure 9.c).[31] The results are robust to the exclusion of economically negligible trading relationships and to propensity score matching (Table 8, Appendix D.1). The differential between trusted and non-trusted destinations is also preserved across specifications, with non-trusted destinations consistently showing larger negative effects.

Therefore, for established trade relationships, export controls seem to introduce a permanent burden for each subsequent shipment of listed products.

Figure 9: Intensive margin - event-study estimates (baseline)

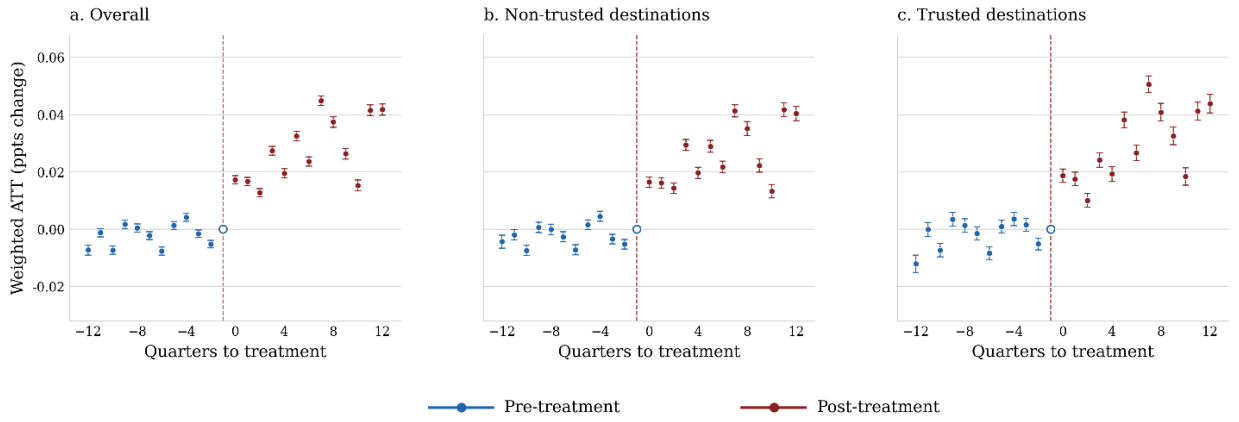


Notes: Event-study estimates of the effect of dual-use listing on Italian export values along pre-existing trade routes. The sample is restricted to firm–product–destination combinations with strictly positive exports in the pre-listing period. Cohort 2018Q4 is excluded due to separation problems in the Poisson likelihood. ATTs are aggregated as weighted averages of cohort–time specific estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Panel a. reports the overall ATT across all non-sanctioned destinations. Panel b. restricts to non-trusted destinations; panel c. to trusted destinations (Argentina, Australia, Canada, Iceland, Japan, Liechtenstein, New Zealand, Norway, Korea, Switzerland, Turkey, and the United States). Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period (t−g = −1).

Regarding the extensive margin, the baseline results indicate that dual-use listing increases the probability of a firm exporting a listed product to any destination by 2.2 percentage points on average (+24% relative to the pre-treatment mean, see Table 8 appendix D.2). This effect is positive and significant across all destination groups, with a slightly larger coefficient for trusted destinations relative to non-trusted ones, and it is permanent, as shown by the event-study estimates (Figure 10). The results are robust to the exclusion of economically negligible trading relationships and to restricting the sample to units with at least two quarters of positive exports. If anything, the magnitude increases slightly when the most marginal units are removed, suggesting that the extensive margin effect is not driven by transitory or negligible trade flows. The estimated ATTs remain positive and significant also across all matching specifications, confirming the robustness of the baseline finding. [32]

[31] For the event study estimates of all robustness checks specifications for the intensive margin see Appendix D.1.
[32] For the event study estimates of all robustness checks specifications for the extensive margin see Appendix D.2.

Figure 10: Extensive margin - event-study estimates (baseline)



Notes: ATTs are aggregated as weighted averages of cohort–time estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Panel a. reports the overall ATT across all non-sanctioned destinations. Panel b. restricts to non-trusted destinations; panel c. to trusted destinations. Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period (t−g=−1). Standard errors clustered at the unit level.

Taken together, the intensive and extensive margin results reveal that long-run aggregate null reflects the net outcome of two offsetting forces: a permanent reduction in trade flows along pre-existing routes, partially compensated by the formation of new trade relationships.

7.2. Export reorientation

To examine whether firms reallocate exports across destinations, we construct two measures of how oriented a firm's exports are toward trusted partners, both computed at the firm-product-quarter level. The first one is the "value share", that is – the share of a given product (in value) that a firm exports to trusted destinations. As an alternative, we use the "count share", namely – the share of trusted destinations (in count) that a firm serves for a given product. The two outcomes answer different questions: the value share captures how a firm's sales are concentrated across trusted versus non-trusted markets, while the count share captures the composition of the destination network. Together they tell us whether firms actively redirect sales toward lower-risk, more permissively licensed destinations, and whether they do so by shifting volumes, by adding or dropping markets, or both. The effect is estimated using the same staggered DiD framework as in equation (1), with firm-product pairs as the unit of observation. Formally:

$$\mathbb{E}\left[Y_{p,t,f} \mid \xi_{p,f}, \xi_t\right] = \exp\left\{\alpha + \xi_{p,f} + \xi_t + \sum_{g \in G} \sum_{\substack{t \\ t-g \neq -1}} \theta_{g,t}\, D_{g,t}\right\}$$

where

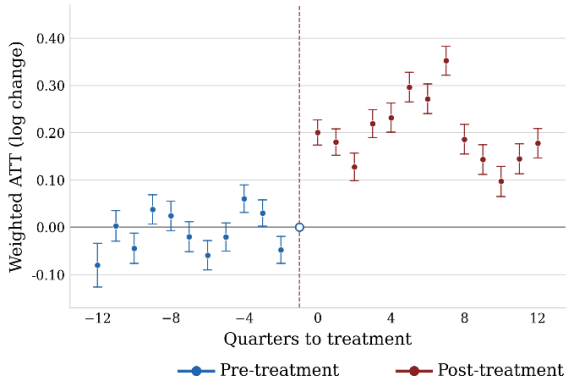
$$Y_{p,t,f} = \frac{\sum_{d:\, Trusted_d = 1} Y_{p,t,d,f}}{\sum_d Y_{p,t,d,f}}$$

or, alternatively

$$Y_{p,t,f} = \frac{\sum_{d:Trusted_d=1} \mathbf{1}\left(Y_{p,t,d,f} > 0\right)}{\sum_d \mathbf{1}\left(Y_{p,t,d,f} > 0\right)}$$

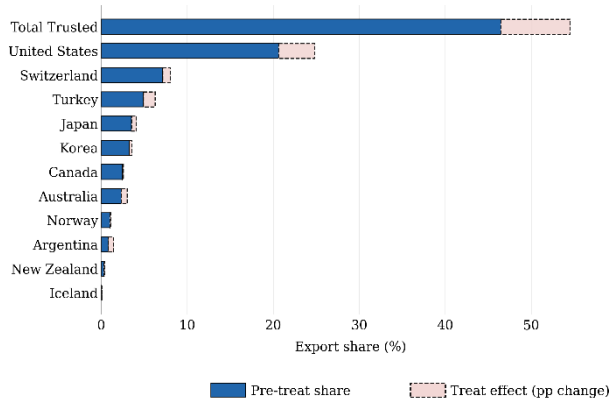
We find evidence that, following the inclusion in the dual-use list, the composition of firms' exports changes persistently toward trusted partners. On average, the share of value shipped to trusted countries increases by 17% (see Table 10, Appendix D.3), corresponding to 8 percentage points - a shift from about 46% to 54% of export value on average (Figure 11). At the country level, the effect is driven primarily by the United States, where the average export share rises by 4.2 percentage points for listed products (Figure 12). We also observe a significant effect for Turkey, though the country's role may be more complex: while listed among trusted partners under the Italian national general license, Turkey may also act as a regional gateway facilitating access to nearby non-trusted markets (Borin et al., 2023). When we re-estimate the effect on the share of exports directed toward trusted destinations using propensity score matching and entropy balancing, the result holds.[33]

Figure 11: Effect on export share to trusted countries



Notes: The dependent variable is the ratio of export values to trusted destinations to total export values, computed at the firm–product level. ATTs are aggregated as weighted averages of cohort–time specific estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Trusted destinations are Argentina, Australia, Canada, Iceland, Japan, Liechtenstein, New Zealand, Norway, Korea, Switzerland, Turkey, and the United States. Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period (t−g = −1). Standard errors clustered at the firm–product level.

Figure 12: Effect on export share, by country



Notes: The figure reports, for each trusted destination country, the pre-treatment export share for treated units (solid bar) and the estimated treatment effect expressed in percentage points (dashed bar). The treatment effect is estimated via a separate Poisson PPML regression for each destination, using the ETWFE framework (Wooldridge, 2023, 2025) with the share of exports toward that destination as the dependent variable. Liechtenstein is excluded as the model is not estimable due to insufficient observations.

These findings are confirmed when we consider an alternative measure based on the share of trusted destinations relative to the total number of destinations served. After listing, the average share of trusted destinations increases by 17% (Table 11, appendix D.3), suggesting that firms are not only concentrating sales

[33] The event study estimates for propensity score matching and entropy balancing specifications are presented in Appendix D.3.

in existing trusted markets, but also expanding the breadth of their trusted-country export network. The resulting ATT is close to our estimate based on value shares.

A large share of firm-product pairs export exclusively to either trusted or non-trusted destinations throughout the sample period. For these observations, the value-based and count-based outcome measures coincide by construction (both being always equal to zero or one), which could explain why the two estimates are so close. To address this, we exclude from the sample units whose share is constant at zero or one throughout. The ATT remains stable in this restricted sample (Table 11, column 2), confirming a similar increase in both the number of trusted destinations and the value exported to trusted partners.

Taken together with the baseline result presented in section 5.2, this shows that even as total exports return to their pre-listing trajectory after the initial quarters, the share directed to trusted destinations increases permanently. This is consistent with the intensive and extensive margin results: the decline in export flows along pre-existing relationships is much larger for non-trusted destinations, while the probability of exporting increases following the listing, with the effect slightly more pronounced for destinations covered by general authorizations. One plausible explanation is that transactions with trusted partners involve lower administrative burdens and more predictable approval outcomes. Buyers in these countries are also typically familiar with export-control requirements and can readily provide end-use certificates, technical documentation, and compliance assurances. More in general, exporters face legal and reputational risks in case of non-compliance, which may further incentivize trade with destinations perceived as low diversion risk and less likely to trigger catch-all provisions or sudden policy tightening. As a result, firms may increasingly prioritize markets with stable regulatory expectations, especially for technologically complex products.

8. Conclusion

Growing international tensions and the increasing military application of modern technologies have led to a vast expansion in export controls, which are now a central tenet of security policy alongside economic sanctions. In the EU, between 2012 and 2023, the number of controlled items increased by nearly 10 percent, with approximately 300 CN8-level products added to the list of dual use-related products.

This expansion has relevant implications for firms. Using firm-level customs data for Italy, a major manufacturing economy and one of the world's largest exporters, we show that dual use-related trade is large and rapidly growing: it accounts for almost 30% of Italy's extra-EU exports and involves more than a hundred thousand firms. Compliance is costly, as securing export licenses requires firms to establish dedicated internal structures, analyse the Regulation, the item to be exported and the counterparts involved. In this paper, we provide the first causal evidence of the effects of EU dual-use export controls on firms' exports.

Our results can be summarized in two main findings.

First, the inclusion of a new product in the dual-use list generates a sizable but temporary decline in its overall exports, concentrated in the first three quarters after listing, driven by compliance frictions and licensing delays. The effect is primarily driven by exports in Machinery and Electrical Equipment, which arguably requires the closest scrutiny to ascertain its dual-use nature. In the medium-term, instead, export controls have no lasting effect on aggregate exports of controlled items.

Second, export controls cause a permanent re-orientation of trade towards "trusted" destinations, with lighter licensing requirements. We find that listing causes a large and permanent loss along the intensive margin, with exports on existing routes falling by about a third and never recovering. Firms are able to compensate for these losses along the extensive margin, by opening new trade relationships, particularly with politically aligned partners. The net result is a permanent reorientation of trade toward "trusted" destinations that benefit from lighter licensing requirements: the share of exports going to these partners rises by around 8 percentage points and does not revert.

Taken together, these results suggest that EU dual-use controls generate short-term adjustment costs and do not create permanent barriers to aggregate trade volumes. However, they induce a lasting reallocation of trade toward politically aligned partners. While the evaluation of the effectiveness or benefits of export controls are beyond the scope of this paper, this appears to be a reasonable cost to pay for achieving security objectives.

In light of the likely further tightening of security-related controls, however, it will be important to keep such costs in check: the wider the perimeter of controlled goods becomes, the more these short-run adjustment costs are likely to accumulate across firms and sectors. In this sense, our findings point to the importance of (i) firms' compliance capacity and (ii) the efficiency of the licensing process.

On the first point, our evidence suggests that a large share of firms remains unaware of their obligations under export control regulation. Indeed, most firms seem to begin the compliance process only after the correlation tables are updated and the customs system flags the affected products in the portal. Clearer outreach and training of firms may go a long way in reducing compliance costs, particularly for firms outside the sectors traditionally associated with dual-use trade (such as aerospace and defense) which are less likely to anticipate that their products will fall under control and to have compliance procedures already in place.

On the second point, two specific measures may improve the efficiency of the licensing process. First, correlation tables may benefit from improved specificity. Most products whose CN codes appear in the tables ultimately turn out not to require a licence, so the current design generates a large volume of compliance checks that impose costs without resulting in actual authorisations. Reducing the false-positive rate in the tables would reduce the compliance burden on firms. Second, ensuring that licensing authorities have funding, staff and IT systems that are fit for processing an increasing number of applications. As control lists expand, licensing capacity might become a bottleneck if the volume of applications requiring individual review grows substantially. The administrative reforms in e-licensing underway in several Member States represent a promising direction for maintaining the system's efficiency under growing demand.

9. Bibliography

Alekseev, M. and X. Lin (2026). Export Policy for Dual-Use Goods. Working paper.

Alfaro, L, H. Fadinger, J. S. Schymik, G. Virananda (2025). Trade and Industrial Policy is Supply Chians: Directed Tecnological Change in Rare Earths. Working Paper 33877. NBER Working Paper Series.

Aytun, U., J. Hinz, and C. Özgüzel (2024). "Shooting down trade: Firm-level effects of embargoes." Kiel Working Paper No. 2277.

Baldwin, D. A. (2020). Economic Statecraft: New Edition. Princeton, NJ: Princeton University Press.

Borin, A., G. Cappadona, F. P. Coteduca, B. Hilgenstock, O. Itskhoki, M. Mancini, M. Mironov, E. Ribakova (2023). The impact of EU sanctions on Russian imports. VoxEU

Callaway, B., and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. Journal of econometrics, 225(2), 200-230.

Chatelus, R. and P. Heine (2016). Rating Correlations Between Customs Codes and Export Control Lists: Assessing the Needs and Challenges. Strategic Trade Review.

Crosignani, M., L. Han, M. Macchiavelli, and A. F. Silva (2024). Geopolitical Risk and Decoupling: Evidence from U.S. Export Controls. Federal Reserve Bank of New York Staff Report No. 1096. New York: Federal Reserve Bank of New York, April 2024.

Crozet, M. and J. Hinz (2020). Friendly Fire: The Trade Impact of the Russia Sanctions and Counter-Sanctions. Economic Policy, 35(101), 97–146.

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

Di Casola, P., T. den Besten, and M.M. Habib (2023). Geopolitical fragmentation risks and international currencies. In The international role of the euro. European Central Bank.

European Commission (2024). White Paper on Export Controls. COM (2024) 25 final, Brussels.

Farrell, H., and A. Newman (2019). Weaponized Interdependence: What Global Economic Networks Shape State Coercion. International Security 44 (1): 42–79.

Gehrke, T. and J. Ringho (2023). The Power of Control: How the EU Can Shaper the New Era of Strategic Export Restrictions. European Council on Foreign Relations. Policy Brief.

Görg, H., A. Jacobs, and S. Meuchelböck (2024). Who is to suffer? Quantifying the impact of sanctions on German firms." Journal of Economic Behavior & Organization, 228, 106767.

Haidar, J.I. (2017). Sanctions and export deflection: evidence from Iran. Economic Policy, 32, 319–35.

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

Hayakawa, K. (2024) The Trade Effects of the US Export Control Regulations. IDE Discussion Paper N. 911. Institute of Developing Economies.

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

Hu, H., S. Yang, L Zeng, X. Zhang (2024). U.S.–China trade conflicts and R&D investment: evidence from the BIS entity lists. Humanities and Social Sciences Communications.

Jäkel, I. C., S. Østervig, and E. Yalcin (2024). The effects of heterogeneous sanctions on exporting firms: Evidence from Denmark. Review of International Economics, 32, 161–190.

Kim, H. and J. Cho (2024). The Impact of US Export Controls on Korean Semiconductor Exports. KDI Journal of Economic Policy.

Kohl, T., M. van den Berg, and L. Franssen (2024). Going Dutch? Firm exports and FDI in the wake of the 2014 EU-Russia sanctions. Review of International Economics, 32, 190–222.

Liu, X., Y. Liu, and J. Wen (2024). The Consequences of Export Controls in Target Countries. Working Paper 25-004. Harvard Business School.

Luttwak, E. N. (1990). From Geopolitics to Geo-Economics: Logic of Conflict, Grammar of Commerce. The National Interest, no. 20 (Summer): 17–23.

Makioka, R. and H. Zhang (2023). The Impact of Export Controls on International Trade: Evidence from the Japan–Korea trade dispute in the semiconductor industry. RIETI Discussion Paper Series 23-E-017.

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

Moon, M. H., and Y. J, Jang (2023). The Effects of Strategic Goods Control on Productivity: The Case of Korea. East Asian Economic Review.

Rasser, M. (2020). Rethinking Export Controls: Unintended Consequences and the New Technological Landscape. Center for New American Security.

Reinsch, W. A., T. Denamiel and M. Schleich (2024). Optimizing U.S. Export Controls for Critical and Emerging Technologies: Working with Partners. Center for Strategic and International Studies.

Rubin, D. B. (2001). Using propensity scores to help design observational studies: application to the tobacco litigation. Health Services and Outcomes Research Methodology, 2(3), 169-188.

Savini Zangrandi, M., L. Patrignani, and A. Schiavone (2026). Between Sanctions and Subsidies: Geoeconomic Lessons from the Semiconductor Ecosystem. Journal of Geoeconomics, 1.

Silva, J. S., and S. Tenreyro (2006). The log of gravity. The Review of Economics and statistics, 641-658.

Strange, S. (1988). States and Markets. London: Pinter.

Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. Statistical science: a review journal of the Institute of Mathematical Statistics, 25(1), 1.

Szczepański, M. (2025). Dual-use export controls as tools of EU economic security. From coordination to a proactive EU approach. European Parliamentary Research Service. Briefing.

Wooldridge, J. M. (2023). Simple approaches to nonlinear difference-in-differences with panel data. The Econometrics Journal, 26(3), C31-C66.

Wooldridge, J. M. (2025). Two-way fixed effects, the two-way mundlak regression, and difference-in-differences estimators. Empirical Economics, 69(5), 2545-2587.

Nagengast, A. J., F. Rios-Avila, and Y. V. Yotov (2026). The European single market and intra-EU trade: an assessment with heterogeneity-robust difference-in-differences methods. Economica, 93(369), 298-331.

APPENDIX A. Dual-use export control framework

This appendix details the institutional architecture underpinning the EU dual-use export control framework discussed in Section 3. Table 4 lists the international treaties and multilateral arrangements from which EU dual-use controls derive. Table 5 reports the ten technical categories and five product groups used to classify controlled items under the Regulation.

Table 4: International Treaties and Multilateral Arrangements for Export Controls

International Treaties	
UNSC Resolution 1540	Prevent non-state actors from acquiring nuclear, chemical, or biological weapons.
Nuclear Non-Proliferation Treaty	Prevents nuclear weapons proliferation and promotes peaceful nuclear use and disarmament.
Chemical Weapons Convention	Bans chemical weapons and controls toxic chemicals and precursors.
Biological Weapons Convention	Prohibits biological and toxin weapons.
Multilateral Arrangements	
Nuclear Suppliers Group	Coordinates controls on nuclear materials and technology.
Australia Group	Coordinates controls on chemical and biological weapons-related exports.
Missile Technology Control Regime	Coordinates controls on missiles and WMD delivery systems.
Wassenaar Arrangement	Coordinates controls on conventional arms and dual-use technologies.

Table 5: Dual Use categories

Category	***Description***
0	Nuclear materials, facilities and equipment
1	Special materials and related equipment
2	Materials processing
3	Electronics
4	Computers
5	Telecommunications and information security
6	Sensors and lasers
7	Navigation and avionics
8	Marine
9	Aerospace and propulsion
Product Group	***Description***
A	Systems, Equipment and Components
B	Test, Inspection and Production Equipment
C	Materials
D	Software
E	Technology

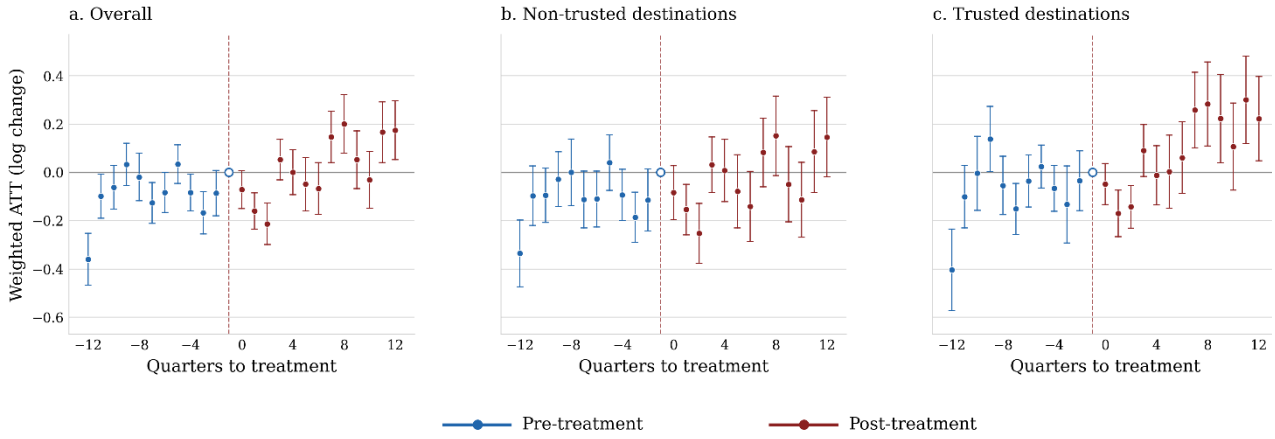
The European Regulation identifies ten categories of dual-use items, further divided into five product groups.

Within each category-product pair, controlled items are specified in minute technical detail. By way of example, a high-precision pressure sensor that can find application in the chemical or pharmaceutical industry, but also in uranium enrichment, may fall into category 2 (materials processing) and product group B (test, inspection and production equipment), under dual use code 2B230. This code identifies pressure transducers that incorporate specified materials (such as certain aluminum, nickel alloys or fluoropolymers in sensing elements or seals) and that achieve defined performance thresholds, including very high accuracy within specified low-pressure ranges.

APPENDIX B. Additional results for baseline specification

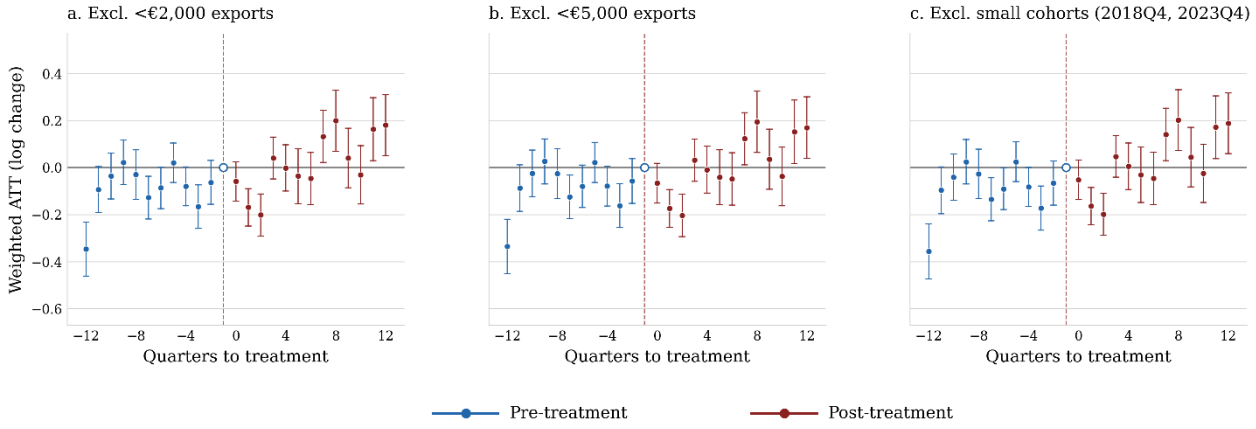
This appendix reports robustness checks for the baseline results presented in Section 5.2. Figure 13 replicates the event study estimates including all extra-EU destinations, among which countries subject to comprehensive sanctions regimes that are excluded from the baseline sample. Figure 14 examines sensitivity to sample composition, progressively excluding marginal trading relationships and the two smallest treatment cohorts. Table 6 decomposes the short-run ATT by broad sector, distinguishing machinery and electrical equipment, the largest group of newly listed products, from all remaining sectors.

Figure 13: Results including all destinations - event study estimates



Notes: This specification includes all extra-EU destinations, including countries subject to comprehensive sanctions regimes (Russia, Belarus, Iran, North Korea, and Myanmar), which are excluded from the baseline specification. ATTs are aggregated as weighted averages of cohort–time specific estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Panel a. reports the overall ATT across all destinations. Panel b. restricts to non-trusted destinations; panel c. to trusted destinations. Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period ($t-g = -1$). Standard errors clustered at the unit level.

Figure 14: Robustness to sample composition - event study estimates



Notes: ATTs are aggregated as weighted averages of cohort–time specific estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Panel a. excludes firm–product–destination combinations with total exports below €2,000 over the full sample period. Panel b. excludes units with total exports below €5,000. Panel c. excludes the two smallest treatment cohorts (2018Q4 and 2023Q4). All panels report the overall ATT across all non-sanctioned destinations. Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period (t−g = −1). Standard errors clustered at the unit level.

Table 6: ATT by sector – first 3 quarters after designation

	(1) *Machinery & electrical*	**(2)** *Other sectors*
ATT (first 3 quarters)	−0.182***	−0.032
	(0.067)	(0.036)
ATT (first 3 quarters) - Trusted destinations	−0.019	−0.044
	(0.084)	(0.040)
ATT (first 3 quarters) – Non-trusted destinations	−0.253***	−0.023
	(0.089)	(0.054)
Observations	21,217,112	92,909,404
Units	757,754	3,318,193
Estimator: PPML-ETWFE in all specifications.		

Notes: The dependent variable is the value of Italian exports. The ATT is computed as a weighted average of cohort–time specific estimates restricted to the first three post-treatment periods, via estat simple (post-jwdid command in STATA), with unconditional standard errors. Column (1) restricts the sample to machinery and electrical equipment products. Column (2) covers all remaining sectors. Standard errors in parentheses are clustered at the unit level. * p<0.10, ** p<0.05, *** p<0.01.

APPENDIX C. Matching results

Table 7 reports the overall ATT estimates from the baseline specification re-estimated using propensity score matching and entropy balancing weights. Across all four specifications, kernel matching, nearest-neighbour matching with n = 5 and n = 10, and entropy balancing, the overall ATT is small and statistically indistinguishable from zero, consistent with the baseline finding of no significant long-run effect on export values.

Table 8 reports two additional robustness checks addressing potential concerns about anticipation behaviour and selection into treatment. Panel A reports ATTs over the full post-treatment period; Panel B restricts to the first three quarters following the actual designation date. Column (2) re-dates the treatment four quarters earlier than the official regulation update, allowing for up to one year of anticipatory adjustment. Column (3) excludes firms affiliated with AIAD and UCIMU, the associations most likely to track regulatory changes proactively. In both specifications, the short-run ATT remains negative and significant, while the long-run effect is not significant, consistent with the baseline. These results indicate that neither anticipatory behaviour nor the presence of large, sophisticated exporters drives the main findings.

Table 7: Overall ATT – matching results

	(1)	**(2)**	**(3)**	**(4)**
	Kernel	*NN (n = 5)*	*NN (n = 10)*	*Entropy balancing*
Overall ATT	−0.031	0.008	0.049	0.014
	(0.045)	(0.057)	(0.061)	(0.044)
ATT – Trusted destinations	−0.045	0.026	0.142	0.053
	(0.074)	(0.116)	(0.129)	(0.072)
ATT – Non-trusted destinations	−0.022	−0.002	−0.007	−0.009
	(0.057)	(0.058)	(0.057)	(0.055)
Observations	44,538,424	19,375,720	26,517,988	101,399,844
Units	1,590,658	691,990	947,071	3,621,423
PPML-ETWFE in all specifications.				

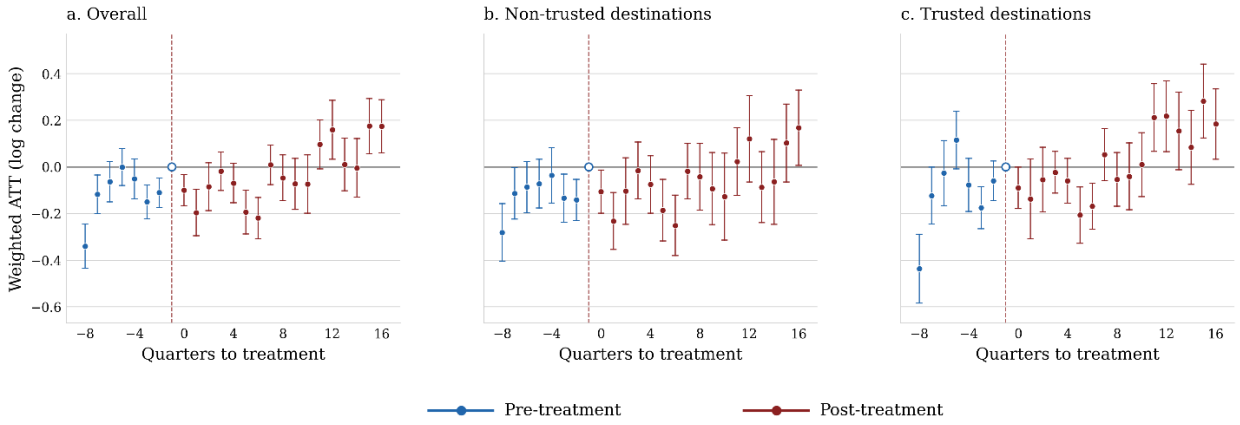
Notes: The dependent variable is the value of Italian exports. Aggregated ATTs are computed as weighted averages of cohort–time ATTs via estat simple (post-jwdid command in STATA), with unconditional standard errors. Specifications (1)-(3) use propensity score matching weights estimated from a logistic regression of dual-use listing status on Product Complexity Index and high-tech share (see Table 2). Column (1) uses kernel matching weights with an Epanechnikov kernel. Columns (2) and (3) use nearest-neighbour matching weights with n = 5 and n = 10 neighbours respectively. Column (4) uses entropy balancing weights. Trusted is a dummy equal to one for friendly destination countries. Standard errors in parentheses are clustered at the unit level. * p<0.10, ** p<0.05, *** p<0.01.

Table 8: Overall ATT - no-anticipation and excluding industry associations

	(1)	(2)	(3)
	Baseline	*Pre-dated (−4 quarters)*	*Excl. AIAD/UCIMU*
Panel A: Overall ATT – full post-treatment period			
Overall	−0.001	−0.025	0.022
	(0.045)	(0.041)	(0.045)
Trusted destinations	0.111*	0.056	0.138**
	(0.063)	(0.054)	(0.063)
Non-trusted destinations	−0.070	−0.076	−0.049
	(0.062)	(0.058)	(0.061)
Panel B: ATT – first 3 quarters after designation			
Overall	−0.139***	−0.161***	−0.114***
	(0.035)	(0.039)	(0.036)
Trusted destinations	−0.121***	−0.145**	−0.094**
	(0.037)	(0.047)	(0.038)
Non-trusted destinations	−0.151***	−0.170**	−0.126**
	(0.052)	(0.055)	(0.054)
Observations	114,126,516	113,743,532	113,170,876
Units	4,075,947	4,062,269	4,041,817
Estimator: PPML-ETWFE in all specifications.			

Notes: The dependent variable is the value of Italian exports. Aggregated ATTs are computed as weighted averages of cohort–time ATTs via *estat simple* (post-*jwdid* command in Stata), with unconditional standard errors. Column (1) reports the baseline estimates. Column (2) re-dates the treatment four quarters earlier than the official regulation update. For Column (2), Panel A reports the ATT for periods $t-g \geq 4$ of the re-dated specification, corresponding to the post-treatment period relative to the actual listing date; Panel B reports the ATT for periods $t-g = 4, 5, 6$ of the re-dated specification, corresponding to the first three quarters after the actual listing date. Column (3) excludes firms affiliated with AIAD and UCIMU, the industry associations whose members are most likely to monitor and anticipate regulatory changes. Standard errors in parentheses are clustered at the unit level. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Figure 15: Pre-dated treatment specification – event study estimates



Notes: Treatment date is shifted four quarters earlier than the official regulation update. ATTs are aggregated as weighted averages of cohort–time specific estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Panel a. reports the overall ATT across all non-sanctioned destinations. Panel b. restricts to non-trusted destinations; panel c. to trusted destinations. Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period ($t-g = -1$). Standard errors clustered at the unit level.

APPENDIX D. Robustness results for intensive and extensive margin specifications and export reorientation

D1. Intensive margin

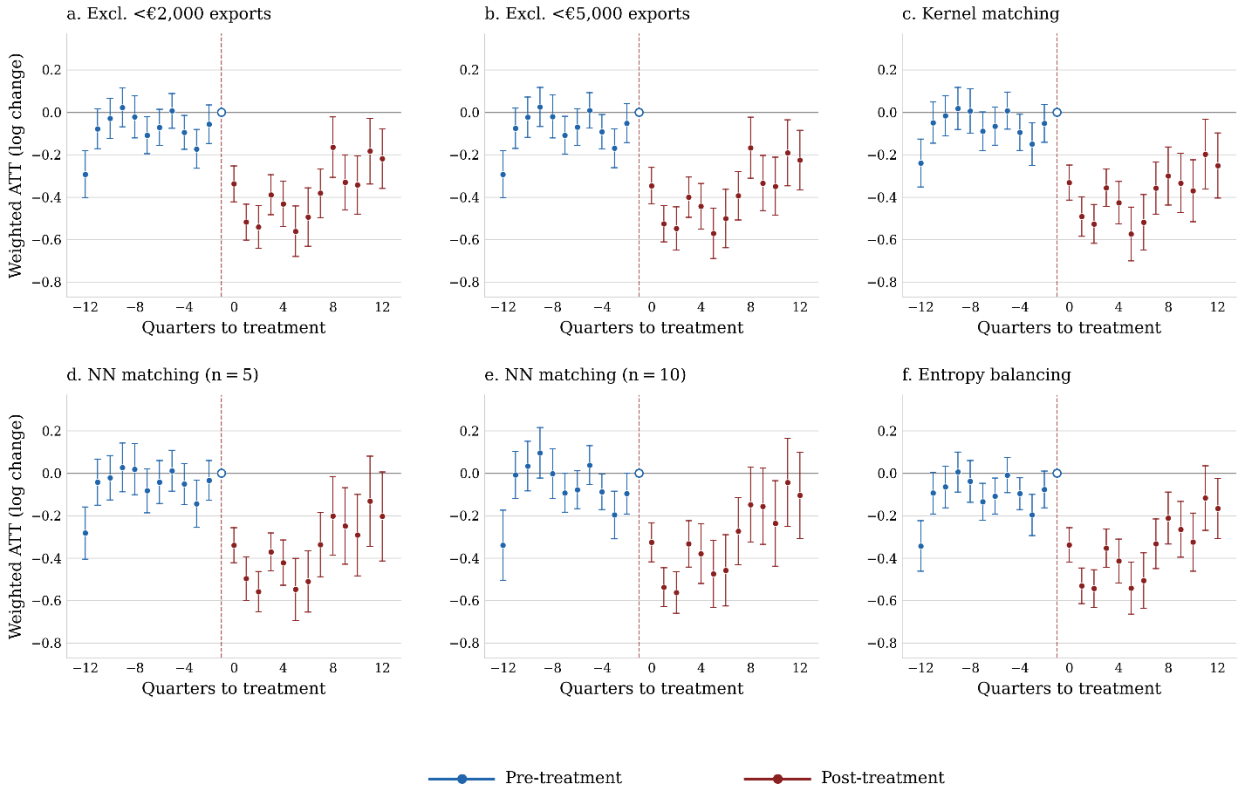
Table 8 reports the full set of results for the intensive margin specification. The baseline estimate (column 1) indicates a large and highly significant negative ATT of −34%, reflecting the persistent decline in export volumes along pre-existing relationships following designation. This result is robust across all alternative specifications: excluding firm-product-destination combinations with total exports below €5,000 (column 2) slightly increases the magnitude of the effect, consistent with compliance costs weighing more heavily on established relationships. Propensity score matching and entropy balancing specifications (columns 3–6) yield ATTs ranging from −29% to −35%, all statistically significant. The asymmetry between trusted and non-trusted destinations is also robust, with the negative effect on non-trusted destinations consistently larger in magnitude than on trusted destinations across all specifications. Figure 16 confirms that the event-study pattern is stable across specifications.

Table 8: Intensive margin: baseline and robustness

	(1)	(2)	(3)	(4)	(5)	(6)
	All units	*Excl. <€5,000 exports*	*Kernel matching*	*NN (n = 10)*	*NN (n = 5)*	*Entropy balancing*
Overall ATT	−0.419***	−0.436***	−0.424***	−0.346***	−0.397***	−0.390***
	(0.045)	(0.045)	(0.048)	(0.065)	(0.060)	(0.046)
ATT – Trusted destinations	−0.211***	−0.216***	−0.313***	−0.119	−0.250**	−0.217***
	(0.063)	(0.064)	(0.078)	(0.137)	(0.121)	(0.076)
ATT – Non-trusted destinations	−0.556***	−0.582***	−0.497***	−0.495***	−0.492***	−0.504***
	(0.061)	(0.062)	(0.060)	(0.060)	(0.061)	(0.058)
Observations	110,367,724	68,469,948	41,382,180	23,236,136	16,154,868	98,117,992
Units	3,948,883	2,447,641	1,477,935	829,862	574,781	3,504,214
Estimator: PPML-ETWFE in all specifications.						

Notes: The dependent variable is the value of Italian exports. The sample is restricted to firm–product–destination combinations with strictly positive exports in the pre-listing period. Cohort 2018Q4 is excluded from all specifications due to separation problems in the Poisson likelihood. Aggregated ATTs are computed as weighted averages of cohort–time ATTs via estat simple (post-jwdid command in STATA), with unconditional standard errors. Column (1) uses the full estimation sample (excluding cohort 2018Q4). Column (2) additionally excludes units with total exports below €5,000 over the full sample period. Columns (3)–(5) use propensity score matching weights (kernel Epanechnikov, NN n = 10, NN n = 5). Column (6) uses entropy balancing weights. Trusted is a dummy equal to one for friendly destination countries. Standard errors in parentheses are clustered at the unit level. * p<0.10, ** p<0.05, *** p<0.01.

Figure 16: Intensive margin - event-study estimates (robustness)



Notes: The sample is restricted to firm–product–destination combinations with strictly positive exports in the pre-listing period. Cohort 2018Q4 is excluded from all specifications due to separation problems in the Poisson likelihood. ATTs are aggregated as weighted averages of cohort–time specific estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Panels a. and b. report results after excluding firm–product–destination combinations with total exports below €2,000 and €5,000 respectively over the full sample period. Panels c.–e. use propensity score matching weights (kernel Epanechnikov, NN n = 5, NN n = 10). Panel f. uses entropy balancing weights. Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period (t−g = −1). Standard errors clustered at the unit level.

D2. Extensive margin

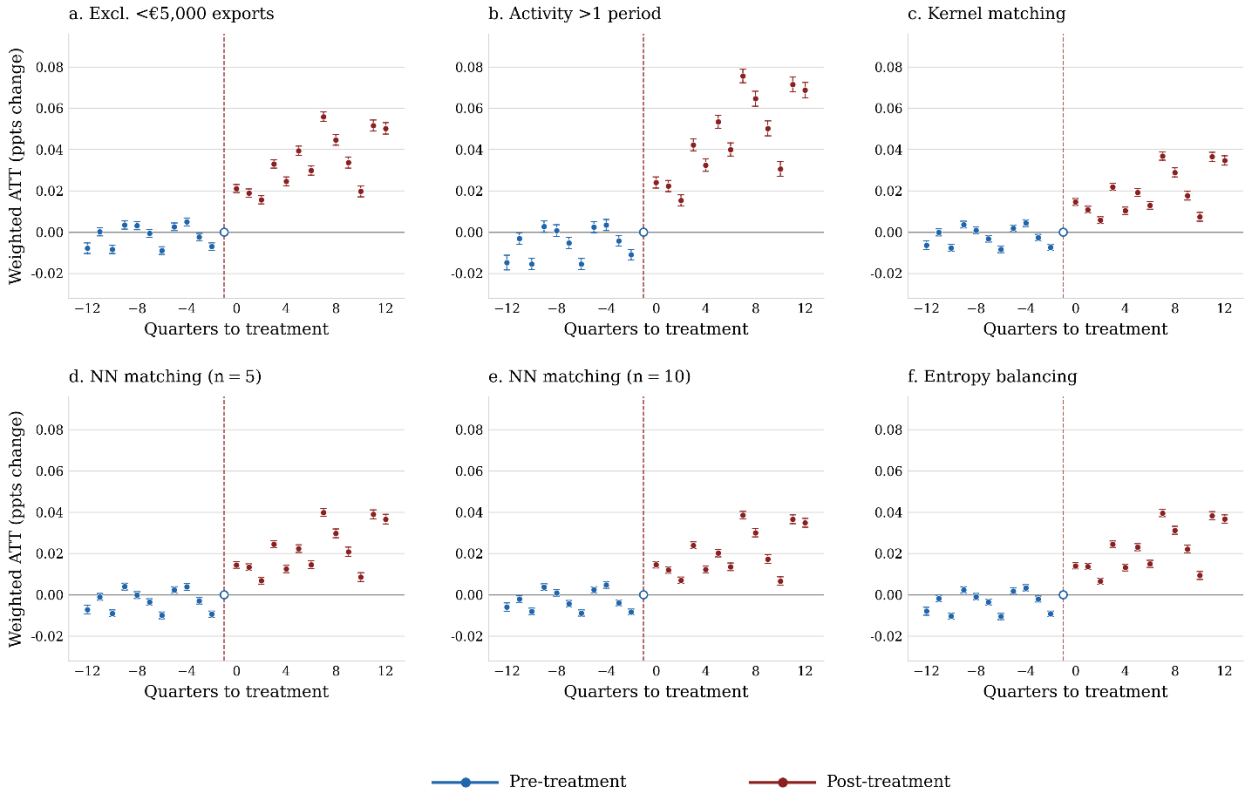
Table 9 reports robustness checks for the extensive margin specification. The baseline ATT of 2.2 percentage points (reflecting the increase in the probability of exporting following designation) is positive and highly significant across all seven specifications. Excluding small trading relationships (column 2) and restricting to units active in at least two quarters (column 3) yields a slightly larger effect of 2.7 and 3.8 percentage points respectively, consistent with the effect being more pronounced for more established relationships. Propensity score matching and entropy balancing specifications (columns 4–7) yield ATTs of 1.5–1.8 percentage points, smaller than the baseline but still highly significant. Figure 17 shows that the event-study pattern is stable across specifications.

Table 9: Extensive margin - baseline and robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	All units	*Excl. <€5,000 exports*	*Activity >1 period*	*Kernel matching*	*NN matching (n=10)*	*NN matching (n=5)*	*Entropy balancing*
Overall ATT	0.022*** (0.001)	0.027*** (0.001)	0.038*** (0.001)	0.015*** (0.001)	0.016*** (0.001)	0.017*** (0.001)	0.017*** (0.001)
ATT - Trusted destinations	0.024*** (0.001)	0.031*** (0.001)	0.041*** (0.002)	0.015*** (0.001)	0.015*** (0.001)	0.017*** (0.001)	0.017*** (0.001)
ATT - Non-trusted destinations	0.021*** (0.001)	0.025*** (0.001)	0.035*** (0.002)	0.015*** (0.001)	0.016*** (0.001)	0.018*** (0.001)	0.018*** (0.001)
Observations	114,126,516	70,607,600	51,918,692	44,538,424	26,517,988	19,375,720	101,399,844
Units	4,075,947	2,521,700	1,854,239	1,590,658	947,071	691,990	3,621,423
Method: LPM in all specifications.							

Notes: The dependent variable is a binary indicator equal to one if the firm exports product p to destination d in quarter t. Aggregated ATTs are computed as weighted averages of cohort–time ATTs via estat simple (post-jwdid command in STATA), with unconditional standard errors. Column (1) uses the full estimation sample. Column (2) excludes units with total exports below €5,000 over the full sample period. Column (3) restricts to units with exports recorded in at least two quarters. Columns (4)–(6) use propensity score matching weights: kernel matching uses an Epanechnikov kernel; NN matching uses nearest-neighbour matching with 10 and 5 neighbours respectively. Standard errors in parentheses are clustered at the unit level. * p<0.10, ** p<0.05, *** p<0.01.

Figure 17: Extensive margin - event-study estimates (robustness)



Notes: ATTs are aggregated as weighted averages of cohort–time specific estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). All panels report the overall ATT across all non-sanctioned destinations. Panel a. excludes firm–product–destination combinations with total exports below €5,000 over the full sample period. Panel b. restricts to units with exports recorded in at least two quarters. Panels c.–e. use propensity score matching weights (kernel Epanechnikov, NN n = 5, NN n = 10). Panel f. uses entropy balancing

weights. Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period (t−g = −1). Standard errors clustered at the unit level.

D.3 Exports to trusted destinations

Table 10 reports robustness checks for the effect of designation on the share of exports directed toward trusted destinations, estimated using PPML. The baseline ATT of 16.8% is positive and highly significant, indicating a substantial reorientation of export flows toward trusted partners following designation. The result is robust across all matching and reweighting specifications, with ATTs ranging from 9% to 10% under propensity score matching and entropy balancing. Figure 18 shows that the event-study pattern for the export share is positive and persistent across all matching specifications.

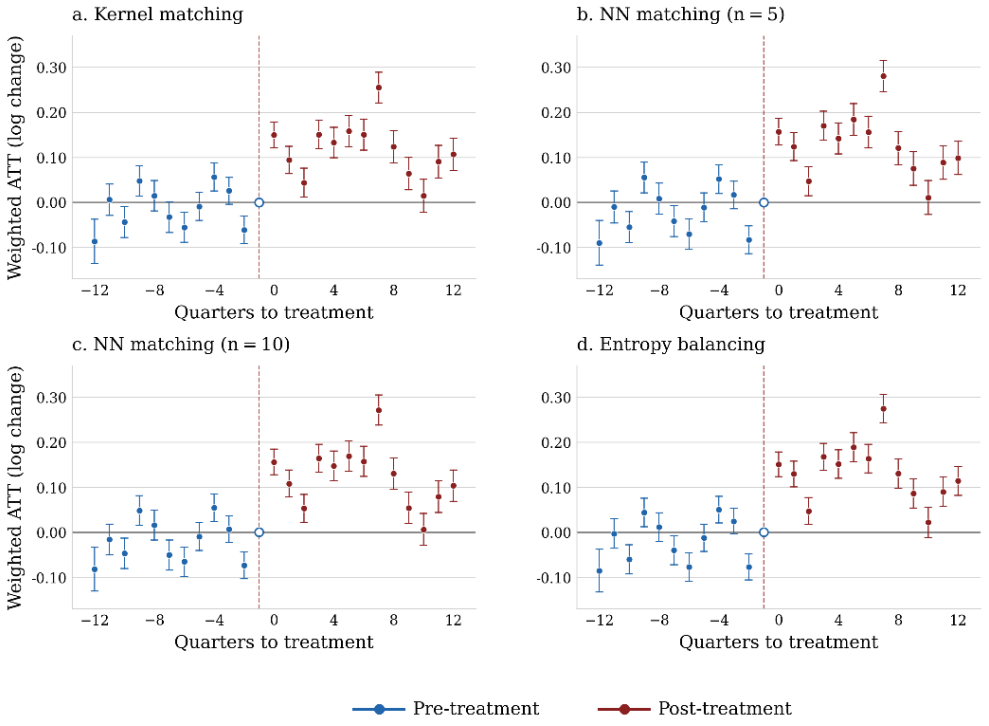
Table 11 addresses the concern that the similarity between the value-based and count-based measures of export reorientation may be driven by firm-product pairs that export exclusively to either trusted or non-trusted destinations throughout the sample period, for which both measures coincide by construction. The ATT based on the count share of trusted destinations is 16.6% in the full sample (column 1) and 15.4% when restricting to units whose count share varies over the sample period (column 2). The stability of the estimate confirms that the reorientation reflects genuine changes in firms' trading relationships rather than a mechanical consequence of the underlying distribution.

Table 10: Effect of dual-use listing on the export share to trusted countries

	(1)	**(2)**	**(3)**	**(4)**	**(5)**
	Baseline	*Kernel matching*	*NN matching (n = 5)*	*NN matching (n = 10)*	*Entropy balancing*
Overall ATT	0.168***	0.088***	0.097***	0.091***	0.101***
	(0.012)	(0.013)	(0.014)	(0.013)	(0.012)
Observations	47,112,716	17,702,888	7,677,096	10,539,312	40,951,904
Units	1,682,597	632,246	274,182	376,404	1,462,568
Estimator: PPML in all specifications.					

Notes: The dependent variable is the share of firm–product exports directed toward trusted destinations relative to total exports (excluding sanctioned destinations). Aggregated ATTs are computed as weighted averages of cohort–time ATTs via estat simple (post-jwdid STATA command), with unconditional standard errors. Column (1) uses the full estimation sample. Columns (2)–(4) use propensity score matching weights (kernel Epanechnikov, NN n = 5, NN n = 10). Column (5) uses entropy balancing weights. Standard errors in parentheses are clustered at the unit level. * p<0.10, ** p<0.05, *** p<0.01.

Figure 18: Effect on export shares to trusted countries - event-study estimates (robustness)



Notes: The dependent variable is the ratio of export values to trusted destinations to total export values (excluding sanctioned destinations), computed at the firm–product level. ATTs are aggregated as weighted averages of cohort–time specific estimates, with weights proportional to cohort size and number of post-treatment periods observed (see Section 5.1). Panel a. uses kernel matching weights (Epanechnikov kernel); panel b. nearest-neighbour matching with n = 5 neighbours; panel c. nearest-neighbour matching with n = 10 neighbours; panel d. entropy balancing weights. Vertical bars are 95% confidence intervals. The dashed vertical line marks the omitted reference period (t−g = −1). Standard errors clustered at the unit level.

Table 11: Effect of dual-use listing on the share of trusted destinations

	(1)	(2)
	All units	*Excl. constant units*
Overall ATT	0.166***	0.154***
	(0.012)	(0.012)
Observations	47,112,716	27,961,108
Units	1,682,597	998,611
Estimator: PPML for all specifications.		

Notes: The dependent variable is the share of a firm's active export destinations for a given product that are classified as trusted (the count share), computed at the firm–product–quarter level. Aggregated ATTs are computed as weighted averages of cohort–time ATTs via estat simple (post-jwdid STATA command), with unconditional standard errors. Column (1) includes all units; column (2) excludes units whose count share is constant over the whole 2017–2023 window (always 0 or always 1). Trusted destinations are Argentina, Australia, Canada, Iceland, Japan, Liechtenstein, New Zealand, Norway, Korea, Switzerland, Turkey, and the United States. Standard errors in parentheses are clustered at the unit level. * p<0.10, ** p<0.05, *** p<0.01.