BAFFI RESEARCH CENTER

Working Paper Master Theses Series

The Effect of Protectionism on Innovation: Evidence from Brexit

Niki Barbas | 2026

Università Bocconi
BAFFI
Centre on Economics, Finance and Regulation

The Effect of Protectionism on Innovation: Evidence from Brexit

Niki Barbas

Bocconi University

April 2026

Abstract

Trade integration and innovation have defined the past decade, yet Brexit — a paradigmatic episode of economic disintegration — occurred at a time when innovation-led growth was widely framed as central to long-run competitiveness. This thesis examines this tension by analysing the extent to which a protectionist shock halts a country's innovative pace. The thesis makes two contributions. First, drawing on a series of expert interviews, it develops a conceptual taxonomy of mechanisms linking protectionism to innovation, including trade barriers, migration, knowledge spillovers, uncertainty, and regulatory change. Second, it provides an empirical analysis using Brexit as a quasi-natural experiment. Product-level exposure to Brexit (UN COMTRADE) is matched to UK patent publications (WIPO), and recent advances in difference-in-differences with continuous treatments are used to recover heterogeneous effects along the exposure distribution. At the aggregate level, the 2016 Referendum is associated with a negative but statistically insignificant average effect on patent growth. However, non-parametric estimates reveal significant declines concentrated among the most highly exposed products. Sectoral evidence from chemicals — directly affected by the UK's exit from the EU REACH framework — shows economically large and statistically significant declines even at moderate exposures. Overall, the findings suggest that a protectionist shock does not automatically halt a country's innovative pace: the strong persistence of innovation acts as a buffer at low and moderate exposure levels. However, this buffer weakens as exposure intensifies. Once sector-specific vulnerability thresholds are crossed, negative effects materialise. More broadly, the innovation effects of deep disintegration shocks do not surface as abrupt breaks, but rather hide beneath aggregate stability, emerging only when examined with the appropriate lens.

Acknowledgements

I would like to sincerely thank Professor Gianmarco Ottaviano for supervising this thesis and for his continued guidance and support throughout the process. I am also deeply grateful to Thomas Sampson, John Van Reenen, Anna Valero, and Ruveyda Gozen for generously participating in my interview series, which allowed this thesis to incorporate a range of valuable insights and reflections beyond the empirical analysis. Finally, I thank Thomas Le Barbanchon for directing me to helpful resources that proved instrumental in shaping and refining my methodology.

Contents

1 Introduction

Trade and innovation have emerged as two of the characterizing macro-trends of the present decade. On the one hand, turns to protectionism and episodes of economic disintegration are increasingly common. On the other, rapidly evolving technological change and innovation have gained renewed importance in debates on long-run growth, competitiveness and global trade patterns, in a newfound Ricardian spirit. This thesis studies this interaction in the context of Brexit, and is motivated by what appears as a fundamental puzzle. Over the past years, the role of innovation as a driver of sustained economic growth has acquired exceptional prominence. In such an environment, one would expect countries to prioritise policies that foster innovation, preserve knowledge flows, and support integration into global value chains. Yet Brexit - a paradigmatic episode of economic disintegration - occurred precisely at this historical juncture. At a time when innovation-led growth featured prominently in economic and political discourse, the United Kingdom supported a political project that may undermine some of the structural conditions that foster innovation. Fundamentally, this points to a disconnect between the rhetoric of innovation and the politics of integration, reflecting the coexistence of two highly salient but potentially conflicting narratives.

This tension motivates the research question addressed by this thesis. **To what extent does a protectionist move halt a country's innovative pace?** Why does political momentum favour disintegration at a time where technological competitiveness is widely framed as central? More broadly, can trade integration and innovation - and the policies governing them - be thought about separately? In an attempt to answer these questions, this thesis revolves around two main contributions.

The first contribution is conceptual in nature. Through a series of interviews with cross-disciplinary experts on Brexit and UK innovation, this thesis tries to bring the innovation dimension of Brexit to the forefront. While most existing research on Brexit focuses on short-run effects on economic growth, analysing innovation is a first glimpse into structural determinants of long-run growth. Insights from the interviews are combined into a coherent conceptual framework that tries to construct a **taxonomy of mechanisms linking Brexit to innovation**. These channels span multiple dimensions, including trade barriers, migration, knowledge spillovers, uncertainty, expectations and the role of policy. This contribution serves a dual purpose. First, this classification is a useful tool that can be used to formulate theoretically grounded predictions on the effects of Brexit on innovation - predictions that can be extended to protectionist shocks more broadly. Second, this conceptualisation helps acknowledging the intrinsic complexity of a "shock" like Brexit; trying to reduce effects to a single channel or shock-type may risk losing fundamental grasp on the bigger picture and the structural adjustments at play.

The second contribution of this thesis is empirical. **Brexit is used as a quasi-natural experiment to study how protectionist shocks affect innovation**. Operating at the granular product level to mirror the nature of non-tariff-barriers, a measure of pre-referendum exposure to Brexit is constructed (UN COMTRADE) and matched to UK-origin patent publications (WIPO) using the Algorithmic Links with Probabilities methodology. Differently from existing studies on Brexit, this thesis builds explicitly on recent advances in Difference-in-Differences with continuous treatments, allowing treatment effects to vary flexibly along the Brexit exposure distribution. Intuitively, identification relies on within-sector variation in products' exposure to Brexit, which allows to compare the innovation pace of relatively more exposed products to their minimally exposed counterparts. Exploiting treatment intensity enables to uncover subtle features of the innovation response to disintegration that would otherwise remain hidden in conventional summary parameters. The main results are the following.

At the aggregate level, the **2016 Referendum is associated with a negative but statistically insignificant average effect of exposure on patent growth**. However, non-parametric estimates reveal that **significant negative effects are concentrated among the most highly exposed sectors**, while intermediate exposures display little or no decline. These results will be interpreted in light of an uncertainty channel operating strongly at high-exposures. The evidence for the 2021 Trade and Cooperation Agreement is more cautious - partly due to its overlap with COVID - and suggests a qualitatively opposite pattern, potentially consistent with delayed patent disclosure following earlier uncertainty. This naturally speaks to the core puzzle of the thesis: does a protectionist move halt a country's innovative pace? The aggregate results suggest a nuanced answer. Overall, they do not point to a broad-based collapse in UK innovation. **Innovation and patenting are highly persistent phenomena, and this persistence appears to act as a buffer against large protectionist shocks**. **At**

low and moderate levels of exposure, this structural inertia "protects" sectors from abrupt innovation slowdowns. **However, this buffer weakens as exposure intensifies. Once exposure crosses a certain vulnerability threshold, negative effects begin to materialise**. In this sense, Brexit does not uniformly derail innovative activity, but it does slow it down where integration is deepest. Even in aggregate terms, therefore, the innovation costs of protectionism emerge at the upper tail of exposure, rather than as an immediate economy-wide break in trajectory.

Interestingly, aggregate results turn out to conceal substantial sectoral heterogeneity, consistent with the sector-specific nature of Brexit-induced non-tariff barriers. The analysis therefore zooms into the salient case of chemicals, a highly innovative sector where contextual and institutional knowledge would strongly predict negative effects. In this context, Brexit signed the UK's exit from the EU REACH framework for the chemical industry, inducing substantial regulatory divergence in the sector. The empirical evidence confirms this narrative, building confidence towards this thesis' methodology. **In the chemical sector, Brexit exposure is associated with economically large and statistically significant declines in patent growth.** Strikingly, negative effects emerge already at relatively low levels of exposure, intensifying and stabilising with exposure — rather than appearing only at the extreme upper tail, as in the pooled specifications. This pattern is consistent with a story in which **Brexit materialises through multiple, overlapping sector-specific channels, that strengthen with the innovation intensity of the industry and its exposure to Brexit-induced non-tariff-barriers.** Ultimately, the sectoral evidence invites a reconsideration of the aggregate results. In fact, **pooled regressions may reflect an averaging over sectorally heterogeneous vulnerability thresholds**: in some industries, only very high exposure triggers innovation slowdowns, while in others — such as chemicals — even moderate exposure is sufficient. Aggregating across these heterogeneous dose–response functions mechanically concentrates statistical significance at the extreme of the exposure distribution, which is precisely what pooled regressions report.

Taken together, the findings offer a nuanced answer to the thesis' core puzzle. **A protectionist shock does not automatically halt a country's innovative pace: the strong persistence of innovation acts as a buffer at low and moderate exposure levels. How-**

ever, this buffer weakens as exposure intensifies. Once sector-specific vulnerability thresholds are crossed, negative effects materialise. Ultimately, this thesis suggests that analysing Brexit through aggregate averages risks overinterpreting an inherently incomplete picture. Instead, insightful findings on the innovation effects of protectionism are more likely to emerge when sectoral heterogeneity and exposure intensity are incorporated into the analysis.

This thesis is organised as follows. The remainder of **Chapter 1** gives some preliminary notions on Brexit and summarises existing literature. **Chapter 2** reports the four expert opinions and combines them into a taxonomy of mechanisms that link Brexit to innovation in the UK. **Chapter 3** explains the construction of *Exposure to Brexit* and *Patents* variables, discussing the relevant data and descriptive analysis. This chapter also features a glimpse into the Algorithmic Links with Probabilities matching strategy. **Chapter 4** dives into the methodology, starting by defining the causal parameters of interest with the relevant identification assumptions, and then outlining the TWFE estimating equation, the binarised version and the two non-parametric alternatives. **Chapter 5** reports the estimation results, distinguishing between aggregate and sectoral effects for the Referendum and the Trade and Cooperation Agreement, respectively. The chapter concludes with a discussion of limitations, suggesting directions of future research. Finally, **Chapter 6** concludes the thesis.

1.1 Understanding Brexit

This thesis studies the relationship between protectionism and innovation in the context of Brexit. Understanding the Brexit phenomenon and its institutional details is crucial in informing the correct design of any study wishing to identify its effects. This section defines Brexit, the referendum, the Trade and Cooperation Agreement, and the notion of non-tariff-barriers that will be used throughout this project.

This thesis defines **Brexit** as the decision of the UK to leave the European Union (EU). In the empirical part of this thesis, Brexit will be formalised as the union of two distinct events: the 2016 Brexit referendum, and the 2021 enactment of the Trade and Cooperation Agreement (TCA). A **timeline of events** is constructed to report all relevant events leading up to Brexit, integrating simultaneous developments in innovation policy. The dark blue events are highlighted as the most significant. The light blue comments describe the "economics" of the phenomenon, borrowing from Dhingra and Sampson [2022] to distinguish first- and second-moment shocks to expectations. The timeline was constructed drawing from a detailed description by Dhingra and Sampson [2022] and other sources cited directly in the figure. The following paragraphs describe the dynamics of the Brexit vote and the TCA.

The **Brexit referendum** happened on June 23rd, 2016, and the UK unexpectedly voted to leave the EU. An important aspect is that the outcome was largely **unanticipated**. For instance, betting markets put 30% odds on the UK voting "leave" [Bell, 2016]. **The fact that the referendum outcome was unanticipated is crucial for the credibility of identification**. A second essential aspect is that after the Brexit vote, there was no guidance on how Brexit would occur, and on whether it would occur at all. This initiated a pervasive wave of Brexit-induced uncertainty Dhingra and Sampson [2022]. On one hand, it induced a first-moment shock to expectations leaning towards reduced future openness. On the other hand, it created a second-moment shock about the many forms that Brexit could take, ranging from soft- to hard-Brexit scenarios being envisioned. This uncertainty lasted 4 years, until a Withdrawal Agreement passed through parliament. Following this idea, this thesis will refer to the "referendum effects" as those occurring in the period 2016-2020.

Brexit formally took place on January 31st, 2020, with the UK leaving the EU and entering a 1-year transition period. The actual Brexit-induced policy changes occurred on January 1st, 2021, when the TCA entered into force. The **Trade and Cooperation Agreement (TCA)** is the UK-EU agreement governing trade between the EU and the UK in the post-Brexit era. Crucially, the TCA retained tariff-free and quota-free goods trade. However, the TCA signed the exit of the UK from the EU Single Market and the Customs Union. The former implies a sacrifice of free movement of people, EU regulatory policy and jurisdiction of EU Courts. The latter implies the unraveling of the deep free trade agreement that existed prior to Brexit. In practice, this triggered the introduction of a customs and regulatory border separating the UK from the EU, and implied the implementation of non-tariff barriers. In this spirit, the TCA can be interpreted as the effective realisation of the Brexit trade shock. This thesis will therefore refer to the "TCA effects" as those materialising in the 2021-2025 period.

WHAT IS BREXIT? A TIMELINE OF EVENTS

1990s
Birth of political movement **against deepening** of European Integration through creation of Single Market and establishment of the European Union

1993
United Kingdom Independence Party (**UKIP**) established to advocate for Brexit

2000s
Wave of **immigration to the UK** from Eastern Europe leads movement to gain popularity (Sobolewska, Ford 2020)

2010-2015
Weak productivity growth and regional divergence; real wage stagnation **following the GFC** (Dorling 2018, Major and Machin 2018)

May 2010
PM David **Cameron** (conservative party) begins a fiscal austerity programme, reducing welfare payments with cuts concentrating in poorer areas. (Beatty and Fothergill 2014, Dhingra 2016)

2012-2015
UKIP gained popularity in areas more exposed to benefit cuts

2014
Start of Horizon 2020 as an EU R&D programme, with UK as one of largest beneficiaries

23 January 2014
In a speech at Bloomberg, PM David Cameron declares support for an **in or out referendum on EU membership,** conditional on conservatives winning the 2015 general election

22 May 2014
UKIP wins **European Parliament** election, winning 27.5% of the vote

7 May 2015
PM David **Cameron wins the general elections**, pledging to do a referendum on EU membership. Cameron and opposition Labor Party **opposed Brexit**, and opinion polls predicted a close vote but with a belief that voters would opt for the status quo

23 June 2016
Referendum on EU membership is held and **UK votes to leave the European Union (51.9% vote Leave)**: a major, unanticipated political and economic shock

24 June 2016
Cameron **resigns** as Prime Minister; **sharp depreciation of sterling value by 10-12% overnight** against the US dollar, and continued in the next month; **fall in the stock market** by 8%

Economics:
First-moment shock to expectations: Expected future openness declined, with the prospect of higher future trade barriers

Economics
Second-moment shock to expectations: policy uncertainty increased, as the referendum was a popular mandate without guidance over when and in what form Brexit would occur

2016-2020
Four years of debates on the **meaning of the Leave Vote**; No evidence of trade diversion from the EU : **EU's share of UK imports and exports stable**

13 July 2016
Theresa May elected PM, pledged to take UK out of Single Market, yet **retaining close relations** with the EU, avoiding regulatory barriers to goods trade, allowing an independent trade policy

2 October 2016
In her Conservative Party Conference speech, **May confirms the exit process (Article 50) will be triggered** before the end of March 2017

17 January 2017
May delivers her Lancaster House speech and **hardens her ideas about Brexit,** including an end to the jurisdiction of the European Court of Justice, exit from the Single Market and control over immigration from the EU

29 March 2017
The **UK triggers Article 50** of the Treaty on European Union **starting a two year process for negotiating** the terms of the UK's exit

8 June 2017
May's gamble on holding a general election backfires. She **loses her parliamentary majority** and forms a minority government with the support of the Northern Irish Democratic Unionist Party

November 2018
Draft **Withdrawal Agreement** between the EU and UK

2019
Horizon Europe (successor of Horizon 2020) negotiations begin, and UK seeks Associated Country status

15 January 2019
Parliament rejects May's Withdrawal Agreement; it will be rejected two further times in March, leading to an **extension of the Article 50 negotiations** and **May's resignation**

March 2019
Initially scheduled Brexit date, then postponed twice

24 July 2019
Boris Johnson becomes PM, in favor of Brexit giving the UK more extensive control over economic and trade policy; he **renegotiates the Withdrawal Agreement;** During May's mandate he led the Conservative Opposition.

October 2019
Johnson negotiates a revised Withdrawal Agreement, but **lacks parliamentary support to pass his deal**

12 December 2019
Johnson holds a general election and **wins** a large majority on the pledge to 'get Brexit done'; the **revised Withdrawal Agreement subsequently passes through parliament**

January 31, 2020
Brexit takes place, and the UK leaves the EU. The UK leaves the EU at 11pm UK time and **enters a standstill** transition period

Until December 2020
Transition period: trade and economic relations with EU unaltered

March 2020
Covid-19 pandemic.

24 December 2020
The UK and EU reach a deal on the **Trade and Cooperation Agreement (TCA)** setting the terms of their future relationship

Until here...
Brexit conceptualised as
- a **shock to expectations:** the first moment shift in expectations (decline in expected future openness of the UK to trade) and second moment shift in expectations (increase in uncertainty over future UK and EU economic policies). Firms take forward looking economic decisions.
- **a large depreciation of the pound sterling**

From here...
Brexit conceptualised as a **reversal of deep international economic integration** through the **introduction of NTBs**

1 January 2021
Trade and Cooperation Agreement (TCA):
- **tariff-free and quota-free trade** between EU and UK in products
- UK no longer a member of **EU Single Market**: no free movement of labor, no EU regulatory policy and no jurisdiction of European Courts
- UK no longer a member of **EU Customs Union**: independent trade policy for goods, no deep free trade agreement
- There are some **level playing field provisions** on labor and environemntal standards

After January 2021
Reduction of UK imports from the EU, differently from pre-TCA

September 2023
UK rejoins **Horizon Europe** after disputes in 2021-2022

Now, to understand the sector-level approach of this thesis, a discussion of the notion of non-tariff barriers is due. Some confusion could arise because the literature often uses the terms "non-tariff barriers" and "non-tariff measures" interchangeably. According to the WTO, the former is defined as a subset of the latter [World Trade Organization, 2012]. **Non-Tariff Measures (NTMs)** are defined as any policy measure other than ordinary customs tariffs that can affect trade. Instead, **Non-tariff Barriers (NTBs)** are a type of NTM that restricts trade protectionistically or unnecessarily. In other words, NTMs are intentionally neutral tools that serve legitimate policy goals, such as food safety or health standards. In contrast, NTBs have a normative connotation, being flagged as protectionist, discriminatory or unnecessarily trade-restrictive. Brexit largely concerned the latter, as most of the introduced trade frictions are given by regulatory divergence or customs formalities. To understand NTMs and their protectionist NTB counterparts, a useful classification is given by United Nations Conference on Trade and Development [2019]. Broadly, NTMs can be applied to imports or exports. On the import side, a further distinction is between technical measures and non-technical measures. Among **technical measures on imports**, there are (A) sanitary and phytosanitary measures, (B) technical barriers to trade (C) pre-shipment inspection and other formalities. Among the **non-technical measures on imports**, there are (D) contingent trade-protective measures, (E) non-automatic import licensing, quotas, prohibitions, quality-control measures and other non-technical restrictions, (F) price-control measures, including additional taxes and charges, (G) finance measures (H) measures affecting competition (I) trade-related investment measures, (J) distribution restrictions, (K) restrictions on post-sales services, (L) subsidies and other support, (M) government procurement restrictions, (N) intellectual property, (O) rules of origin. Finally, on the **export side**, there are (P) export-related measures.

Each of the above can contain a number of measures describing necessary certification, testing, inspection, quarantine, packaging, licensing, quality controls or marketing rules. Three examples can help understand the type of "shock" that Brexit represents and what the correct level of analysis is [United Nations Conference on Trade and Development, 2019]. **Rules of origin** are conditions under which a product can be considered as "economically originating" in a country, which in turn determines eligibility for preferential tariff treatment. If firms prove that a sufficient share of value-added originates from the UK or EU, then tariffs are null. It is intuitive to understand that this type of NTB affects disproportionately the manufacturing sectors with complex value-chains. **Sanitary and Phytosanitary (SPS) measures** are technical measures protecting human, animal and plant health. UK-EU trade flows entailing meat, dairy and plants need to provide health certificates and pass border inspections. This NTB type affects disproportionately the agriculture and food processing sector. Lastly, **passporting rights in financial services** are a subset of NTBs affecting trade in services. After Brexit, UK-based financial firms cannot "passport" services across the EU, but need a separate authorisation to operate and must deal with regulatory duplication. This last category disproportionately affected banking, insurance and asset management. From this discussion, it is evident that **by construction, NTBs operate through distinct regulatory channels that are intrinsically sector-specific**. As a result, Brexit effects can be expected to **materialise through a series of overlapping sector and product-specific effects**. **In absence of firm-level data, analysing variation at the sector level aligns with the institutional granularity at which NTBs operate. This thesis will develop following this intuition in developing a sectoral-approach to the study of the effect of Brexit on innovation.**

1.2 Literature Review

This thesis sits at the intersection of roughly three literatures: the theory of trade and innovation, the causal effect of trade liberalisation or protectionism on innovation, and the economic consequences of Brexit.

The first relevant stream of literature is that of **theories of trade and innovation**, where fundamental contributions started emerging in the 1990s. A central piece of this literature is represented by Grossman and Helpman [1991]. In their book, the authors present a series of endogenous growth models in which trade openness accelerates (long-run) innovation by expanding the effective market size and encouraging international spillovers. Underlying the former result is the intuition that integration with foreign markets allows firms to sell to a larger customer base, raising the returns to innovation and incentivising firms to invest in R&D seeking monopoly profits. Instead, the international spillovers channel refers to the benefits of cross-border flows of ideas for innovation, facilitated by trade openness. Romer [1990] endogenises technological change, viewing it as an outcome of firms choosing optimal R&D investment to maximise profits. This allows to micro-found the market-size rationale: larger economies -with more resources channeled in research- create innovations that allow innovators to appropriate returns. Thus, market size has a positive relationship with the equilibrium rate of innovation. Coe and Helpman [1995] provide the first empirical evidence for the role of international knowledge spillovers for innovation. In practice, using a panel of 22 OECD countries, they show that a country's total factor productivity is determined by its own R&D stock, as well as the R&D stock of its trading partners. Eaton and Kortum [2002] propose a Ricardian trade model where countries differ in their productivity draws. In their model, countries import goods from whoever has the best technology of production, implying that trade flows reveal the international distribution of innovation and that trade is a determinant of access to technology. Melitz [2003] crucially introduces firm heterogeneity in trade theory. The author shows how trade liberalisation triggers a reallocation of resources from low-productivity firms (that exit or shrink) to high-productivity exporting firms, ultimately raising aggregate industry productivity through a selection mechanism. An important implication is that barriers disproportionately harm the most productive, export-oriented firms. Keller [2004] provides a useful survey on the empirical literature relating to the channels through which technology spreads internationally. The author underscores the role of trade, FDI, patents, migration and geographical distance. Furthermore, the author reinforces the Coe and Helpman [1995] result whereby import-intensive sectors benefit the most from foreign R&D. Aghion and Howitt [1992] propose a growth model driven by Schumpeterian creative destruction. In the model, firms invest in R&D and innovate to avoid being displaced by future innovators. They also demonstrate that policies that lower the cost of R&D or increase returns to innovation accelerate growth. Rivera-Batiz and Romer [1991] provide an interesting contribution in terms of deep integration. The authors contend that goods market integration may have ambiguous growth effects, while *deep integration* - including knowledge integration - generate positive effects on long-run growth. Finally, Atkeson and Burstein [2010] build a model combining firm-level innovation decisions with trade, finding that trade liberalisation boosts aggregate innovation by reallocating resources towards high-productivity exporting firms. Crucially, however, their framework predicts that the aggregate welfare effects of trade costs are small, as changes in trade barriers trigger offsetting responses in product innovation that largely neutralise the impact on aggregate outcomes.

While the theoretical literature generates clear predictions, empirical literature is needed to discipline the sign and magnitude of these effects. This leads to a second related stream of relevant literature, that concerns the **empirics of the causal effects linking trade and innovation**. Two recent papers are particularly relevant. Aghion et al. [2018] use french firm-level data and find that firms whose export markets expand unexpectedly increase patenting, with effects concentrated in frontier firms. They show how the export channel operates through a scale effect rather than learning-by-exporting or competition mechanisms. Coelli et al. [2022] show that tariff cuts from multilateral trade agreements consistently increase patenting. They jointly attribute this to export-market size (firms innovate more as they can reach a large foreign market) and import-competition (domestic firms face foreign competition and innovate to escape it). Bloom et al. [2016] use the rise in Chinese import competition after WTO accession to identify the effects on EU firm innovation. Industries facing more competition display higher patenting, IT adoption and TFP growth, consistently with the escape-competition mechanism. The authors also show that lower productivity firms exit the market, coherently with the reallocation logic. Aghion et al. [2005] employ a panel of UK firms and find an inverted-U relationship between product market competition and innovation. They contend that at lower levels of competition, the escape-competition effect dominates, yielding increasing innovation. Instead, as competition intensifies further, the Schumpeterian effect dominates as innovation rents are eroded and disincentivise undertaking innovation in the first place. Other relevant empirical papers are Impullitti and Licandro [2018] and Bustos [2011]. Shu and Steinwender [2019] provide a thorough literature review of the impact of four different forms of trade liberalisation on firm productivity and innovation. They find that import competition has mixed effects on innovation (depending on whether "business stealing" or "escape competition" dominate), while export opportunities have generally positive effects. The authors also point to underexplored margin of trade disintegration rather than liberalisation. Other two valuable surveys of the broader literature are Akcigit and Melitz [2022] and Akcigit and Van Reenen [2023].

Finally, a third stream of related literature is that of the **economic consequences of Brexit**. Since Brexit is a relatively recent phenomenon, the papers relating to its consequences can be classified in two distinct approaches: some early papers and reports took an ex-ante approach to forecast consequences of Brexit under different scenarios; some very recent papers try to estimate the effects that have actually materialised ex-post. Dhingra et al. [2016] is an example of the former, where authors forecast the cost of Brexit on UK trade and living standards under two alternative Brexit scenarios. Dhingra et al. [2023] discuss the role of non-tariff barriers for trade and welfare, and assess the Brexit case-study, concluding that welfare losses from departure from the EU cannot be offset by new deals with key non-EU partners. Some ex-ante studies delve into specific sectoral effects. For instance, Gasiorek et al. [2019] anticipate that manufacturing could be especially vulnerable to Brexit NTBs. Among the very recent ex-post studies, Bloom et al. [2019] use the Decision Maker Panel to show the effects of the referendum. The paper finds evidence of large increases in business uncertainty, and a reduction of investment and productivity compared to a no-Brexit scenario. A similar and extremely recent paper from the same authors, Bloom et al. [2025] features a short extension which is the closest comparator to this thesis - operating at the firm level rather than the product level. Using PATSTAT, the paper provides tentative evidence showing that firms more exposed to the EU are less likely to apply for patents after 2016. They acknowledge that since innovation effects take time to materialise, it is plausible to think that effects will intensify over time. Dhingra and Sampson [2022] find that the referendum triggered higher import and consumer prices, lower investment, and slower real wage and GDP growth. The authors also show that in the aggregate, no trade diversion from EU happened before prior to the TCA. Also, Arnold et al. [2025] develop a granular measure of UK's diverging regulation with the EU, and find that Brexit-induced NTBs introduced significant trade frictions that have not been compensated by competitiveness improvements in non-EU markets. Freeman et al. [2024] use firm-level customs data to show the effect of the 2021 TCA on international trade. They find immediate declines in exports and imports from the EU, smaller firms being hardest hit, exports to non-EU countries remained stable, and imports with non-EU countries increased, signaling trade diversion. Moving to sectoral studies, a salient case concerned the chemistry sector, discussed in Jones and Burns [2024] to describe the developments and effects of the diverging regulation. Keiller [2024] shows that manufacturing firms exposed to EU trade saw a negative impact on investments post-referendum, with effects particularly pronounced for import- rather than export-exposure to the EU. Finally, Bakker et al. [2022] zoom into the food sector to provide evidence of the role of Brexit in rising food prices and inflation. Remarkably, this paper uses product-level pre-referendum EU import shares to measure exposure to Brexit - an approach mirrored by this thesis.

In terms of positioning, this thesis could be described as follows. It contributes to the broad empirical literature that tries to underpin the magnitude and direction of effects of integration shocks on innovation. However, it contributes to the latter in two relatively unexplored directions. First, while most empirical papers use liberalisation shocks, this thesis operates on the side of disintegration. Second, rather than studying a more traditional form of protectionism through tariff barriers, it focuses on the unraveling of *deep* economic integration through non-tariff barriers. The latter is a more subtle yet extremely relevant feature of our understanding of integration in a globalised world.

This project also contributes to the analysis of the economic consequences of Brexit, by focusing on the innovation outcome that, to date, is only marginally studied in Bloom et al. [2025], and in an unpublished paper by Boler, Gotzen and Martin (2026). Also, while existing studies on Brexit typically exploit pre-referendum EU import shares as a continuous measure of exposure, they often rely on conventional two-way fixed effects specifications that summarise treatment effects into aggregate parameters. In contrast, this thesis explicitly builds on recent advances in difference-in-differences with continuous treatments [Callaway et al., 2024], allowing treatment effects to vary along the exposure distribution. By moving beyond summary coefficients and potentially fragile TWFE estimators, the analysis recovers richer heterogeneity in the response of innovation to Brexit intensity. In this sense, the paper provides a methodological illustration of how frontier econometric tools can be applied to fully exploit the informational content of continuous treatment designs in the Brexit context.

Indirectly, this thesis also contributes to the first stream of theoretical literature. Through four interviews with experts in international trade and innovation, the Brexit case-study is thoroughly analysed. This provides two contributions to this theoretical stream of literature. First, reasoning through an applied context reveals additional channels that often remain implicit in more stylised theoretical models. For instance, an innovation economist's perspective suggested that disintegration can affect the longer-term direction of technical change - an aspect that is often overlooked but is of primary importance. Second, stepping back and examining the mechanisms jointly in a real-world setting helps intuitive understanding of how these forces interact.

2 Expert Opinions

This thesis explores the broad question of the effect of Brexit on innovation in the UK. Innovation is generally not considered a first-order effect of Brexit, as primary effects would concern trade or GDP. As such, it is not immediate to think about the channels that link these two phenomena. At the same time, thinking about Brexit shock propagations into innovation allows to gain some insights on the non-trivial question of what kind of economic shock Brexit really represents.

These considerations motivate this chapter, that leverages on expert opinions to form an integrated perspective about the relationship between Brexit and innovation. These interviews explore the topic by discussing **theoretical mechanisms, dimensions of effect heterogeneity and policy prescriptions**. A nice feature of this interview series is that it integrates cross-disciplinary perspectives from trade economists and innovation economists. This was done with the intention of observing how different fields can reach distinct yet complementary conceptualisations of the same phenomenon, while stressing the importance of particular dimensions and angles of the issue. Working as a pre-doctoral research assistant at the Centre for Economic Performance (CEP) of the London School of Economics gave me the opportunity to arrange interviews with **Thomas Sampson, John Van Reenen, Anna Valero and Ruveyda Gozen**. Being in the UK and working in the innovation or Brexit fields for some time, the researchers provided stimulating insights on the topic. **Sections 2.1 to 2.4** summarise each of the interviews in essay form. **Section 2.5** extrapolates joint conclusions from the four talks. Crucially, it also features what is considered one of the contributions of this thesis, namely a conceptual map constructing a **taxonomy of mechanisms** linking Brexit to innovation.

2.1 Some thoughts from Thomas Sampson

Thomas Sampson is an Associate Professor at the LSE, and an expert in international trade and economic growth. Among the four interviewees, Thomas could bring a trade economist's angle to the discussion, as he extensively published about Brexit and related effects. He also worked extensively on what links trade and innovation. [1]

[1] Interview with Thomas Sampson, London School of Economics, 10.02.2025, conducted by the author.

When thinking about the effect of Brexit on innovation, the first step is to inspect more generally the **role of trade and economic integration for innovation**. Thomas worked extensively on this link and summarised what the theory has to say about how trade integration matters for innovation. He explains that there are two main mechanisms at play. The first is a **market size effect**; viewing innovation in the form of investment as a fixed cost, any increase in the size of the potential market generates increased incentive to invest on innovation. From this perspective, a **positive** empirical relationship between trade openness and innovation emerges as a natural theoretical prediction. The second mechanism is an **import competition effect**; when a country integrates, firms gain export opportunities but also face more foreign import competition domestically. The literature thinks of this channel as being a bit more nuanced, as import competition can change innovation incentives in ambiguous ways. For instance, certain models predict a negative effect on innovation, that can potentially offset the positive market size effect. The underlying idea is that as firms face import competition, they will sell less, effectively reducing the market size (as in, profits and market share) and, ultimately, innovation incentives. On the other hand, Schumpeterian models in the Aghion-Howitt spirit, predict escape competition type of effects that move positively with innovation. The rationale of these models is that firms innovate to escape competition. It follows that as competition increases, firms invest in innovation to differentiate, meaning that competition positively stimulates innovation. Thus, while the market size channel is clear and positive, the competition channel remains theoretically ambiguous - yet likely important.

Having clarified how the theory thinks about the link between trade and innovation, Thomas suggests to apply this framework to rationalise the **effect of Brexit on innovation in the UK**. Theoretically, Brexit can be seen as a reduction in trade integration. Through the market size effect, this should have a negative effect on UK innovation, while the reduced import competition effects remain ambiguous. While highlighting that any consideration on the net effect of Brexit remains an empirical question, Thomas **expects that the net effect of Brexit on UK innovation would be a negative one**. He explains that, under certain assumptions, say a Constant Elasticity of Substitution world, the market size and import competition effect in a single sector economy exactly offset each other - according to an old Grossman Helpman result. Yet, when allowing for Schumpetarian escape competition effects, the negative effects that could come from import competition are mitigated or even overturned, so the net effect of more integration turns positive. Symmetrically for the Brexit case, this would mean a negative effect on innovation stemming from reduced integration. The combination of these two facts lead him to expect a negative effect on innovation from Brexit.

However, is this what we've seen so far and what we expect to see the in the future? Thomas suggests that we still **do not have a careful empirical study** of whether that is indeed the case. He explains that at aggregate level, UK GDP growth and trade growth have been hit negatively by Brexit. Thus, it would be natural to think that part of this slowdown in GDP growth might be coming from low innovation. However, to conclude this definitively, he highlights that we would like to go into the data and this has not been done yet. In terms of the evolution of effects over time, Thomas expects **effects to cumulate slightly over time**, rather than fully unfolding in the first couple of years. This is due to the forward looking nature of innovation and, since current innovation is an input for future innovation, it does take some time to emerge.

Thomas proceeds to underline that everything discussed above has been holding knowledge spillovers from abroad fixed. Yet, reduced integration may also **reduce the flow of knowledge from abroad**. If this is the case, then we have an additional channel that would be expected to reduce UK innovation. In practice, we have less of an empirical handle on knowledge spillovers, namely, whether and how this channel operates. However, one mechanism that people appointed to in recent work is that when trading, people interact with foreign buyers and sellers, learning something from them. In this view, trading less implies fewer interactions and less learning. Ultimately, less trade means lower knowledge spillovers and, if we think of knowledge spillovers as an input into innovation, then that reduces the incentives to innovate.

Having explored and applied a theoretical framework to understand the effect of Brexit on innovation, an interesting aspect is to consider potential **dimensions of heterogeneity** of these effects. Thomas starts by explaining that an obvious dimension of heterogeneity is between firms that are more or less exposed to trade with the EU; if this mechanism is operating through changes in trade with the EU, then firms that trade with the EU will be most exposed, on either the export or import front. Another relevant dimension could be the **supply**

of innovators to the UK and **changes in migration policy**; if UK firms rely on being able to draw talent from the EU to take part in R&D and now that becomes harder through Brexit, that could also be an important dimension of the problem. In short, we should pay attention to firms' ability to hire cross-border as that can determine which firms are most affected by Brexit in terms of their innovative capacity. When prompted to reflect about whether the effects on innovation would be **level effects** -less innovation overall- or **compositional effects**-different types of innovation-, Thomas suggests that answering is rather complex. He explains that there is a lot of uncertainty around this question, as our understanding of what drives innovation is weaker than our understanding of what drives, for instance, trade. Yet, he guesses that both level and compositional effects will be relevant. He then reinforces that industries exposed to EU trade will be most affected, yet also featuring an overall net negative effect, as previously discussed.

Having elaborated on the trade-innovation link in the Brexit context, several natural questions arise in terms of policy implications. For instance, unfolding the **role of innovation policy** and making a snapshot of the current state of innovation policy in the UK can help put things into perspective. Thomas explains that, broadly, innovation policy changes the incentives to innovate, either positively or negatively, depending on what the policy is. The UK seems keen to encourage innovation by providing various R&D tax credits, but at the same time it does not provide particularly high levels of funding for R&D. For instance, historically, UK levels of R&D has not been as high as other advanced countries.

Reasoning about the link of patent policy and trade, Thomas makes reference to his paper *Trade, Innovation, Optimal Patent Protection*. As countries become more integrated through trade, the more cross-border spillovers of patent policy, so that countries will tend to underprovide patent protection (a free rider problem). They will do that generically, but the effect will be stronger the more integrated the countries are - hence, the stronger the case becomes for cooperating in patent policy across countries. Now, reading Brexit as a form of disintegration, **incentives to cooperate over patent policy have decreased**. However, this will be relatively small and second order, compared to the other things (market access, trade costs, uncertainty...) that were shocked in UK trade at that time. In short, Thomas suggests that Brexit slightly changes patent cooperation incentives, but the dominant innovation effect of

Brexit operates through market size and integration, not patent coordination.

This thesis and discussion prompts to think about what Brexit reveals about the **extent to which trade and innovation policy can be thought about separately**. Thomas assessed that, to answer this, we need someone to do the study of what happens to innovation in the UK after Brexit. If it is the case that Brexit has had a big negative effect on UK innovation, it would clearly suggest that innovation outcomes depend on trade policy. If instead studies suggested that innovation was not much affected, it would suggest they are more independent. Hence, the question remains unanswered pending the empirical work to see what has actually happened.

To conclude our interview, I asked Thomas about what questions in this field are still open and what he would be curious to know. His first suggestion concerned the type of research that Nick Bloom and his team have conducted, by looking extensively at the uncertainty dimension around Brexit. In this view, there is a question of to what extent **Brexit induced uncertainty has had dampening effects on innovation**, separate from the trade channels described previously. To the extent that R&D is a forward looking investment, just as capital investment, it is likely to be dampened by uncertainty - and Brexit did create a lot of uncertainty for about 5 years. Hence, the uncertainty channel would be operating independently and might be worth exploring. His second suggestion is about the **empirical question** of the effect of Brexit on Innovation that we discussed in this interview and that is central to this thesis. Finally, Thomas highlights a further open empirical question relating to **regulatory autonomy**. Pro-Brexit arguments frequently emphasised that leaving the EU would grant the UK greater regulatory freedom, potentially allowing it to to regulate in a way that promotes entrepreneurial startups, ultimately encouraging innovation. This raises the empirical question of whether post-Brexit regulatory changes have materially encouraged entrepreneurial activity or innovation, thereby offsetting some of the negative trade-related effects. However, whether such a channel has occurred and the extent of its quantitative importance remains an open question.

2.2 Some thoughts from John Van Reenen

John Van Reenen is the Director of POID, and senior economic advisor to the UK Chancellor of the Exchequer, Rachel Reeves. John is probably one of the biggest UK innovation experts, having extensively published in the field, co-authoring papers on innovation with - among others- Philippe Aghion, David Autor, Daron Acemoglu and Nick Bloom. Hence, his opinion would combine innovation expertise and a policy drive [2].

When prompted with the question of how Brexit affected UK innovation, John starts by taking a step back. Brexit has a lot of first order economic impacts, and, even if there was no effect on innovation at all, the **overall impact on welfare would remain negative**. Ex-ante, the expectation was that of significant negative effects from Brexit, the more so, the harder the Brexit deal. The ex-post evaluations of Brexit have proved to be consistent with that. For instance, several papers coauthored with Agnes Norris Keiller demonstrated significant declines in investment post-Brexit. With the premise in mind of not needing negative innovation effects to observe overall welfare losses from Brexit, John proceeds to discuss prospective innovation effects. The likely effect of Brexit would be a negative one on innovation, and this can be understood through three main mechanisms. This should be thought about primarily from a **trade perspective**; namely, trade stimulates innovation through the market access effect coherent with the endogenous innovation trade models, or through import competition. A useful survey about the effects of trade on innovation can be found in his book, *The Economics of Creative Destruction*, in a chapter written by Melitz and Redding. A second channel could be that of **Foreign Direct Investment**. Namely, to the extent that Brexit drives a reduction in FDI, and the latter is a booster of innovation, innovation should be expected to decline. A third mechanism would be through **immigration**. Brexit ended free movement with the EU. Since EU migrants tended to be higher skill, more educated and younger, reductions of EU migrants could be driver of a negative effect. One of the ironies of Brexit is that post-Brexit, the government - including Boris Johnson - actually made the migration system much more *liberal*. This led to an increase in total immigration, with the 2023 value of net migration was the highest registered in recent history - with a large portion being lower skill individuals. As a result, the overall immigration channel may be weaker than initially anticipated, with the net innovation impact depending on the skill composition of post-Brexit migrants. For instance,

[2] Interview with John Van Reenen, London School of Economics, 16.02.2025, conducted by the author.

this immigration channel should be rationalised as a change in composition of migrants away from the EU, and more towards China, India and Nigeria.

Another channel somehow related to the trade ones outlined above is that of **creative destruction and firm dynamism**, which mostly operates through **weaker import competition**. Before Brexit, UK firms competed directly with EU firms, who could export freely into the UK; low trade costs triggered by the customs union implied a high competitive pressure. Brexit caused a sudden increase in Non Tariff Barriers, leading EU firms to export less to the UK and releasing some of the competitive pressure borne by UK firms. Creative destruction crucially relies on inefficient firms exiting, efficient firms growing and entering, with competitive pressure forcing innovation. As escape competition weakens through Brexit, incumbent firms face less pressure, low-productivity firms survive longer, markups increase, and incentives to innovate fall. Some business dynamism effects could also materialise through **market size effects**. Before Brexit, UK firms had frictionless access to EU markets which gave scaling opportunities through export. Large markets reward high-productivity firms, and exporting firms tend to occupy the right tail of productivity. This implies that exposure to foreign markets encourages upgrading. With Brexit, exporting became more costly, effectively shrinking effective market size. This means that high-productivity UK firms may scale less, entry incentives will fall, and reallocation towards top firms slows. Thus, creative destruction could also weaken through reduced export opportunities. However, while theory suggests that weaker competition could reduce dynamism, the magnitude of such firm turnover and reallocation effect remains an empirical question.

In terms of empirically identifying the reallocation effects, John suggests to **differentiate between two margins of creative destruction**. The first would be the between-firm margin; this would require checking if lower productivity firms are exiting less, and if the rate of firm turnover is declining, which would be indicative of weaker creative destruction. The second margin is within firms; that is, incumbent firms adjust internally by changing their innovation intensity, adjusting productivity without altering their entry-exit behaviour. Another obvious identification problem is given by the **occurrence of COVID**. One way of reasoning about this is to leverage on the fact that the industries most exposed to the Brexit shock are different from the ones most exposed to the COVID shock. For instance, COVID industries would include the face-to-face sectors of retail, hospitality etc., while Brexit industries would be the most exposed to goods trade. Thus, if the innovation decline is stronger in goods-trade sectors (not hospitality), that supports a Brexit channel.

This reflection naturally leads to a consideration of **heterogeneity in Brexit's effects on innovation**. While there may be broader general equilibrium effects operating at the aggregate level, the first-order mechanisms are likely to operate through trade exposure. In this sense, sectors that are more integrated into EU trade routes are expected to be disproportionately affected. Importantly, Brexit did not primarily introduce new tariffs under the TCA, but rather increased **non-tariff barriers**. These include rules of origin requirements, customs frictions, regulatory divergence, and constraints in services trade such as the loss of financial passporting. Such **NTBs generate sector-specific disruptions that are more subtle but potentially pervasive**. For example, sectors heavily reliant on EU migrants may face innovation constraints if EU and non-EU workers are not perfect substitutes. Similarly, financial services may be affected through the loss of passporting rights, which could alter the scale and integration of UK-based innovation in that sector. Additionally, during the 2016–2021 period of uncertainty, industries that anticipated potential tariff exposure — such as the automotive sector exporting to the EU — may have experienced a "chilling effect," whereby firms delayed or reduced innovation and investment in anticipation of future trade barriers. Overall, Brexit's impact is therefore unlikely to be uniform, but instead to manifest through **multiple overlapping and sector-specific channels** linked to trade intensity, labour composition, regulatory exposure, and expectations.

In discussing the policy response, John emphasises that the most direct way of addressing any negative innovation effects from Brexit is through **reducing the trade frictions** that were introduced. From this perspective, the most far-reaching option would be rejoining the European Union, while intermediate options include re-entering the Customs Union or the Single Market. Short of these, any measure that lowers trade costs would help mitigate the underlying impact of Brexit. An alternative strategy has been to pursue new trade agreements with non-EU partners, such as the recent India trade agreement, accession to the CPTPP, and bilateral agreements with Japan and South Korea. However, in his view, operating on the trade-reintegration margin addresses the issue at its root, whereas **innovation policy op-**

erates more as a compensatory mechanism and is therefore of second-order importance relative to restoring trade integration.

That said, he acknowledges that the UK has undertaken substantial **innovation policy activity** since 2016. Initiatives such as the introduction of ARIA, the Help to Grow Management and ICT programmes, the broad innovation agenda outlined under Rishi Sunak, and the decision to rejoin Horizon Europe represent meaningful efforts to sustain research funding and international collaboration, particularly in science and higher education, one of the UK's strongest export sectors. The central question, however, is whether these measures have effectively compensated sectors most negatively affected by Brexit. While targeted support has been provided in certain industries, such as motor vehicles, it remains unclear whether comparable compensating mechanisms have been implemented across all of the sectors that experienced the largest trade-related shocks.

When asked about **innovation policy priorities**, John begins by noting that, while the innovation angle is important, it may not be the most central margin through which Brexit operates. Even if Brexit had no effect on innovation at all, overall welfare losses would still arise through other economic channels. Turning specifically to innovation policy, he identifies several key areas. One major priority concerns **artificial intelligence** and how the UK can retain as much of the productivity gains from this general-purpose technology as possible. The government has taken steps in this direction, including the establishment of AI growth zones aimed at supporting the development and scaling of AI-related activities. A second priority relates to the **overall level of R&D spending**. At the start of the most recent government, there were plans to reduce public investment, including R&D expenditure, as part of broader fiscal tightening. John emphasises that maintaining, and indeed increasing, public investment in R&D is crucial, and he encouraged the reversal of these proposed cuts. This led to a substantial increase in public investment (on the order of £110 billion) despite the usual fiscal pressures that often result in reductions in investment during periods of financial constraint. Finally, he refers to developments in **industrial strategy**, particularly the decision to focus on eight key sectors seen as potential future growth areas of the UK economy. These include life sciences, artificial intelligence, and finance, supported through regulatory changes, targeted policy measures, and skills-oriented frameworks. Together, these initiatives reflect an attempt to strengthen the UK's innovation capacity, even if they do not directly address the underlying trade frictions introduced by Brexit.

2.3 Some thoughts from Anna Valero

Anna Valero is Director of CEP's Growth Programme, and deputy Director of POID. She is an expert in innovation, and currently serves as an Industrial Strategy Adviser to Chancellor Rachel Reeves. I was interested in Anna's opinion because it could combine an innovation economists' expertise with a strong awareness and presence in UK policy [3].

An interesting way of thinking about the effect of Brexit on innovation in the UK, is to think of **how innovation actually happens within a firm**. While with other interviewees the focus was more on aggregate effects, Anna develops some thoughts from the point of view of a firm deciding to innovate. She identifies three main mechanisms that could be at play. The first mechanism concerns **trade and trade barriers**, and their implications on market access. For instance, a firm deciding to innovate will consider the market they will be able to reach with that innovation. If the market for the innovative product is large and easily accessible through aligned regulations, the firm will be encouraged to innovate. However, if non-tariff barriers are introduced, the implied reduced market access and regulatory misalignment will naturally act as a disincentive to innovate in the first place. This becomes particularly important for the development of new technologies, where expectations on the target market is crucial. This channel may instead be less important for firms whose innovations exclusively target the UK market. A second firm-level mechanism concerns **skills and human capital** as a crucial input to innovation. For many years, the UK has been a magnet for talented people from Europe, who would then became entrepreneurs or set up innovative businesses. When Brexit happened, many of these people felt less welcome, and quite a few left the UK and decided to go back to their countries. This is a problem for innovation output, both directly and indirectly. Directly, through entrepreneurs who own firms, or who would have set up innovative firms absent Brexit. Indirectly, through staff that could have enabled innovation within existing firms. A third mechanism would be the ability of firms and innovators to access **funding and innovation grants**. In the case of Brexit, this concerned European shared innovation grants, and Horizon Europe. Anna highlights that rejoining Hori-

[3] Interview with Anna Valero, London School of Economics, 27.02.2025, conducted by the author.

zon Europe was extremely important for the UK, and this was very much in the public debate for years. The fourth and last channel is that of **collaborations**. At the origin, this largely concerns UK researchers in universities, engaging in international collaborations towards the development of innovations. However, this type of research has significant spillovers and implications for firms and their ability to innovate. In this view, Brexit resulted in lower access to collaborations with Europeans, which is damaging to the UK's innovation system more broadly.

Interestingly, Anna explained the trade-related channels by making reference to the role of **regulatory alignment**, which seems to be the margin of NTBs that is most relevant to innovation. Prompted to expand on this aspect, Anna explains that historically, the UK always had some regulatory matters from which it was opting out. Nevertheless, it was largely aligned with EU regulation. When thinking of a protectionist shock, it is not only tariffs that matter for goods passing through borders, but also regulatory alignment. The latter becomes particularly important when it comes to **innovation happening in new areas, where regulation is still being created**. This is because in that case, regulation acts as a **key enabler of whether the innovation will be able to access the market**. In this view, for a firm deciding to innovate in the UK, it is important to have the certainty that the UK and EU markets will be aligned on the regulation relevant to that specific innovation. Alignment means that the UK can expect to access the large market at its doorstep, which is obviously a favourable environment for innovation to occur. The alternative to alignment is to opt for divergence from EU regulation, mostly in the form of deregulation. On one hand, this could potentially benefit some particular products through enhanced regulatory freedom. On the other hand, these firms would face the additional burden of navigating diverging regulation to (eventually) be able to reach European markets with their innovations. Ultimately, if there is incompatible regulation, innovations may end up being "trapped" in the UK market. Overall, it is safe to say that Brexit NTBs trigger a new burden of complexity for innovating firms.

A salient case of regulatory misalignment after the referendum was that of the **chemicals industry** - with relevant implications on pharmaceuticals too. Prior to Brexit, the UK chemicals industry was subject to the the EU regulatory framework REACH - Registration, Evaluation, Authorisation and Restriction of Chemicals. Under EU REACH, firms registered substances once, and this allowed them to sell across the entire EU market. This was managed by a centralised regulatory authority, the European Chemicals Agency (ECHA). After Brexit, the UK could no longer rely on EU Reach, and had to created a parallel system called UK REACH to duplicate all relevant regulations. Under this scheme, firms wanting to sell chemicals in both UK and EU markets need separate registrations, compliance processes, data submissions and testing - all of which is costly. This is especially important for innovation overall, as chemistry is a highly innovative sector in the UK. This is a key example of how **regulatory divergence increases innovation complexity** and, in turn, **regulatory alignment is an innovation enabler**.

Having dived into the mechanisms, it is natural to think about possible dimensions of firm heterogeneity. Anna explains that she would expect **smaller firms to be the most impacted**. Small firms are already less likely to export and face larger frictions to market access - for instance due to lack of expertise. Moreover, reasoning about the first channel, such firms are known to be less able to navigate the complexities of diverging regulation. In particular, small high-growth firms that are relevant to innovation and growth may have suffered the most, eventually deciding to relocate to Europe, or deciding not to start up at all. On the other hand, larger firms usually find ways to navigate divergence, for instance by moving headquarters or specific functions to Europe, in order to retain market access. Finally, Anna suggests that heterogeneity would be best understood by looking at the relevant firm-level data.

An angle of the Brexit-innovation question that could be tackled with Anna is about what types of innovation would be most impacted. Innovation can be of multiple types: frontier innovation, diffusion of existing technologies, adoption, product or process innovation. It is reasonable to expect that certain types of innovation would be more vulnerable to Brexit than others. Anna explains that she would imagine **frontier innovation to be the most impacted**. In fact, frontier innovation heavily depends on international talent, cross-border collaboration, research grants, access to large markets, regulatory alignment and the openness of the innovation ecosystem overall - all of which are part of the channels that were earlier discussed on the Brexit impacts. Thus, the weakening of international openness would make frontier R&D suffer first - and the consequences of a closed innovation system should be a concern. On the other hand, diffusion and adoption are less vulnerable to Brexit. Diffusion and adoption barriers typically concern management quality, organisational capability, information frictions, financing constraints and internal skills - all of which are mostly domestic structural issues, largely unrelated to EU integration. For instance, one could think of the adoption of digital technologies. Most of these are invented in the US and diffuse to other countries, depending on the factors cited above. Realistically, we would not think of Brexit as a constraint to digital diffusion. A similar argument applies to argue that product innovation would be more impacted then process innovation.

The second part of Anna's interview leveraged on her role of Advisor for UK Industrial Strategy to gain intuition on the policy dimension. Anna explains that, historically, the UK has reasonably emphasised the role of innovation; even under the Boris Johnson years, there were commitments to increase public investments in R&D. The present government also increased **real investment in R&D**, something that Anna personally contributed to. Another relevant dimension of innovation policy is about making sure that **grants are going to firms that have high-growth potential**. Furthermore, a set of policies is in place to encourage local spillovers, for instance through the **Local Innovation Partnerships Funds**, currently managed by the Department of Science.

Despite the many policy changes in recent years, Anna explains that the UK's underlying innovation strengths have not fundamentally altered, and this gives the country some positive underlying continuity. For example, life sciences are historically a field where the UK retains a comparative advantage from an innovation perspective, and in continuous to do so despite Brexit. Similarly, many recent efforts were concentrated in the development of an AI Action Plan. However, previous industrial strategies had already concentrated on AI, robotics and the digital. Anna's take on the current state of innovation policy is that at the **granular level there is some structural continuity on areas of focus, while most issues are related to policy churn**. Policy churn - namely frequent changes in policy directions, priorities or institutional programmes - create uncertainty for businesses when it comes to innovation. It is therefore essential for innovation policy to **keep building on the valuable underlying frameworks** in place, operating from the bottom-up, while seeking more stability to manage business expectations to favour innovation.

Finally, Anna suggested a dimension of the Brexit-innovation question that she believes is under-discussed in the public debate and that would benefit from more attention. She explains that government policies are not only about increasing the amount of innovation, but also about influencing the direction of innovation. The latter is called **directed technical change**. The UK being close to the EU meant being part of shared agendas and discussions on direction and scope. Leaving the EU implied not being part of such discussions in recent years in the same way. A salient example would be about wanting to shape AI to have positive social purposes, "good use", for instance trying to enhance rather than replace workers. Together with "Safe AI", other objectives in European agendas are shared security and resilience objectives, or environmental objectives. Anna has substantial experience on the latter, having largely worked in Net Zero. She explains that **Net Zero** is an area where directed technical change is extremely important. This is also closely related to resilient growth, as it is more resilient to have home-grown renewable energy. The UK was the first advanced economy to put Net Zero by 2050 into law, being a leader on climate policy. This consensus started breaking down in recent years, even if the current government committed to clean energy missions. There was an era where the US was also aligned on this, now this is not the case anymore. Conversely, the EU is still much emphasising the green transition. Thus, **being out of the EU during this period has prevented the UK from being aligned on certain directions of innovation with countries that see these issues broadly in the same way**.

2.4 Some thoughts from Ruveyda Gozen

Ruveyda Gozen is a Research associate at the LSE for POID and hosts the Innovation and Diffusion podcast with John Van Reenen. Ruveyda published in the field of innovation for several years, and is currently working on the paper *Brexit and the Falling Innovation Dynamism*, jointly with Ralf Martin and Esther Boler. Ruveyda brings specific knowledge on the topic and hands-on experience on this type of data [4].

Ruveyda and her coauthors provide what is currently the first empirical working paper about the effect of Brexit on innovation. Their team started working on this topic because they viewed Brexit as a historically significant event that is likely to trigger substantial **structural change**. Both before and during Brexit, economists raised concerns on prospective negative effects. In the past years, the literature provided substantial evidence about this; Brexit has triggered

[4] Interview with Ruveyda Gozen, London School of Economics, 13.02.2025, conducted by the author.

several short run effects including declines in business investment, productivity and trade. Ruveyda explains that their paper was interested in inspecting the unexplored innovation angle of Brexit, because innovation and knowledge production directly feed into a country's growth path and can inform on what to expect in the longer run.

In particular, *Brexit and the Falling Innovation Dynamism* is an empirical investigation on the **effect of the 2016 Brexit vote on innovation collaborations and knowledge production**. Their paper offers two distinct contributions. The first contribution is mostly descriptive in nature, and leverages PATSTAT patent data to observe the evolution of patent collaborations. The hypothesis is that UK-EU inventor collaborations would fall after the referendum, compared to the UK-UK and UK-Rest Of the World (RoW) counterparts. In particular, they maintain that all inventor collaborations would fall due to Brexit, but the UK-EU would be hardest hit. The second and main exercise they perform is to merge PATSTAT with ORBIS data to conduct a firm level analysis. Their objective is to analyse whether firms who are more exposed to UK-EU collaborations prior to the referendum, see a subsequent reduction in their innovation intensity.

In the **patent level analysis**, they observe that patent collaborations between the UK and the EU declined significantly after the Brexit vote. In terms of magnitude of the effect, they find that the 2019 patenting level is even *lower* than the one registered in the year 2000. They interpret this as Brexit having a more disruptive effect on innovation compared to the financial crisis. However, since the time series ends in 2019, there could be a later recovery that remains unobserved in the dataset. In the **firm level analysis**, their research finds that firms who were engaging in UK-EU innovation collaborations, were subject to a greater innovation decline compared to firms who were only domestically involved in innovations, or engaged in collaborations with the rest of the world. This result is particularly interesting in light of inspecting a potential compositional change mechanism: in principle, it could have been that while UK-EU collaborations were declining, UK-ROW were increasing as a form of compensation mechanism. Yet, this is not the case. The takeaway is that Brexit affected the UK's innovation activity overall, which Ruveyda views as even more concerning.

While these effects remain true in the aggregate, an interesting extension is done in terms of heterogeneity analysis. For instance, Ruveyda explains that an important angle is that of **green innovation**. She explains that working with green patents dramatically reduces the sample size. Nonetheless, the negative effects explained above seem to remain true on the sample of green patents. In particular, they estimate that one patent out of six vanishes after the Brexit vote. Another exercise they perform is to distinguish **intensive and extensive margin effects**. The former refers to how much a firm innovates, while the latter reflects whether a firm innovates at all. In particular, they looked at patent intensity, conditional on the firm being an innovator pre-referendum. Their research finds that among firms that were already patenting, the number of patents declines after the Brexit vote. This suggests that the slowdown is not only about firms stopping innovation, it is also about innovators slowing down their pace. One last dimension of heterogeneity is that firms collaborating with the EU are different from firms collaborating with the RoW. Ruveyda argues that precisely the fact that firms collaborating with the EU are hardest hit, builds confidence on the patenting decline being driven by Brexit. She specifies that their findings are robust to conditioning on a series of firm observables.

Overall, this research was conducted in the spirit of extracting the sign of the relationship between Brexit and innovation, without specifically looking into a mechanism to isolate to claim causality. However, the channel that is implicitly captured by their investigation is that of the **uncertainty shock** triggered by Brexit; innovation is by construction a risky investment, due to uncertainty over its outcome. This leads innovation to be under-supplied in general, which explains the importance of innovation policy and subsidies to compensate for this. In this view, the Brexit vote further enhanced the uncertainty that is already intrinsically associated with innovation. In turn, this can result in falling innovation efforts. Another more practical reason for focusing on such mechanism is related to data availability. The patent data they use is the 2023 PATSTAT release, whose last credible year of data is 2019-2020, because of the lag in patent data. It follows that, since in 2019 the TCA was not in place, the observed effects must be related to uncertainty and expectations, rather than market access or alternative channels.

Ruveyda also signals that in future work, she envisions an investigation of the potential **brain drain mechanism** underlying the effect of Brexit on innovation. More precisely, inventor names and locations could be leveraged to study if researchers and scientists left the UK after the vote - for instance, for VISA complications and associated uncertainty; if they did, it would be interesting to track their new destinations. The significance of the brain drain mechanisms can be grasped through a **historical evaluation of the changing nature of innovation**. In the 19th century US or UK, "garage inventors" were the drive of technological change. Namely, individuals developing impactful innovations. Over time, with globalisation and the emergence of corporations, **innovation became a team effort**. This is where thinking of innovation as a trade and flow of ideas becomes particularly important. In this view, Brexit brings structural change that complicates the process of bringing together these teams to facilitate innovation through international collaboration.

This last reflection underlines the importance of being able to attract innovators to come and stay in the UK. More generally, Ruveyda articulates that **innovation policy** and **migration policy** need to be jointly discussed. In the 19th and 20th century, we learned that the innovation-migration ecosystem rests on the interaction of three pillars: universities (knowledge production and talent formation), firms (commercialization, patenting and scaling) and government policy (shaping innovation incentives and mobility). In the UK, the ability of universities to attract high-skill students and researchers has been a historical contributor to innovation. After Brexit, universities not only started struggling on the financial front, but also faced difficulties in attracting international students. The latter weakens the innovation ecosystem more broadly and suggests that innovation policy should not be thought about separately from migration policy. Ruveyda cites, among others, a paper by Moser, Voena and Waldinger about the innovation effect of Jewish migrants fleeing to the US from Nazi Germany, demonstrating that talent mobility shapes innovation outcomes. Ruveyda adds that, being the UK an education exporter with a conglomerate of good universities, it is in a particularly good position to utilize international talent. The priority should therefore be to attract these students and keep them in the country. Once in the country, facilitating the creation of networks is also essential. For instance, Cambridge and Oxford tried to link their universities through improved transportation, and take advantage of the "bumping into each other" network effects. More broadly, Ruveyda argues that it is essential to promote a global innovation network, not only with the EU but also with the rest of the world.

Overall, Ruveyda's reflections highlight the importance of examining the innovation consequences of Brexit not merely as short-run fluctuations, but as forces that may shape the UK's longer-term growth trajectory. Moreover, her emphasis on the close interdependence between innovation policy and migration policy underscores that the UK's innovative capacity is deeply tied to its ability to attract and retain global talent. In this view, Brexit's impact cannot be assessed solely through trade metrics, but must also be understood through its implications for the broader innovation ecosystem.

2.5 Extrapolating Conclusions from Expert Opinions

This interview series has touched upon different theoretical or empirical aspects of the Brexit effects on innovation. Thomas proposed a trade-theoretic lense on the effects of economic disintegration. John integrated the discussion with insights from firm dynamism and creative destruction, and provided a policy hierarchy and sectoral-effect considerations. Anna contributed interesting thoughts about the types of innovation involved, the role of regulatory alignment for innovation and directed technical change. Ruveyda provided first evidence of Brexit-induced uncertainty effects on innovation collaborations and stressed the relevance of migration policy and brain drain risks. Overall, the theory seems to predict **negative net innovation effects of Brexit**. The remainder of this thesis will try to establish if this is verified in the data.

Insights from these interviews can be combined to construct a **taxonomy of mechanisms** linking Brexit with innovation. **Figures 1 and 2** jointly provide a systematic classification of all channels that emerged from the interviews. The mechanisms can be classified into four distinct yet intertwined areas: "Trade and Trade Barriers", "Migration and Knowledge Spillovers", "Uncertainty and Expectations", and the "Role of Policy". These charts serve a dual contribution. First, having such classification allows researchers and policymakers to **formulate theoretically grounded and mechanism-specific predictions on the effect of Brexit on innovation**. In turn, this is useful to guide future research by formulating a coherent set of testable hypotheses, or to locate existing literature. Second, what emerges from this classification is an objective **acknowledgment of the multidimensionality implied by protectionist shocks**. It follows that Brexit cannot be boxed into a single economic shock type - say, a trade shock. Doing this could in fact result in overly-narrow framings and reductionist perspectives. Instead, it is important for research to inspect individual channels in the awareness that each of them is part of a broader multifaceted protectionist phenomenon where many effects are jointly at play. While the literature acknowledges the intuitive multi-dimensionality of Brexit, it is striking that even through the narrow lens of innovation alone, the disintegration shock reveals considerable embedded complexity.

A last remark for this section is about another channel that I believe could be relevant for evaluating NTB effects, but is not explicitly highlighted in theories and did not come up in any of the expert interviews. The intuition builds from a paper by Bakker et al. [2022] that finds evidence of an imperfect pass-through of NTB costs into consumer prices, making the case for Brexit-induced sectoral inflation. Taking this reasoning a step further, higher costs for firms are *partially* passed on to consumers in the form of higher prices, while part of the cost rise will be incurred by firms themselves. If this cost rise cuts into firm margins, companies will find themselves with fewer funds to re-invest into R&D. In turn, lower inputs to innovation may lead to a decreased pace in innovation output, as tracked by lower patent counts. This argument implicitly assumes imperfect capital markets and the presence of financial frictions, whereby firms are partially constrained to finance R&D through internal funds. This **financial frictions mechanism** is likely to operate for smaller, younger firms without established credit access. In turn, this could be particularly important for small innovative startups. With firm-level data on R&D expenditure, this channel could be tested following an empirical approach that mirrors that of this thesis.

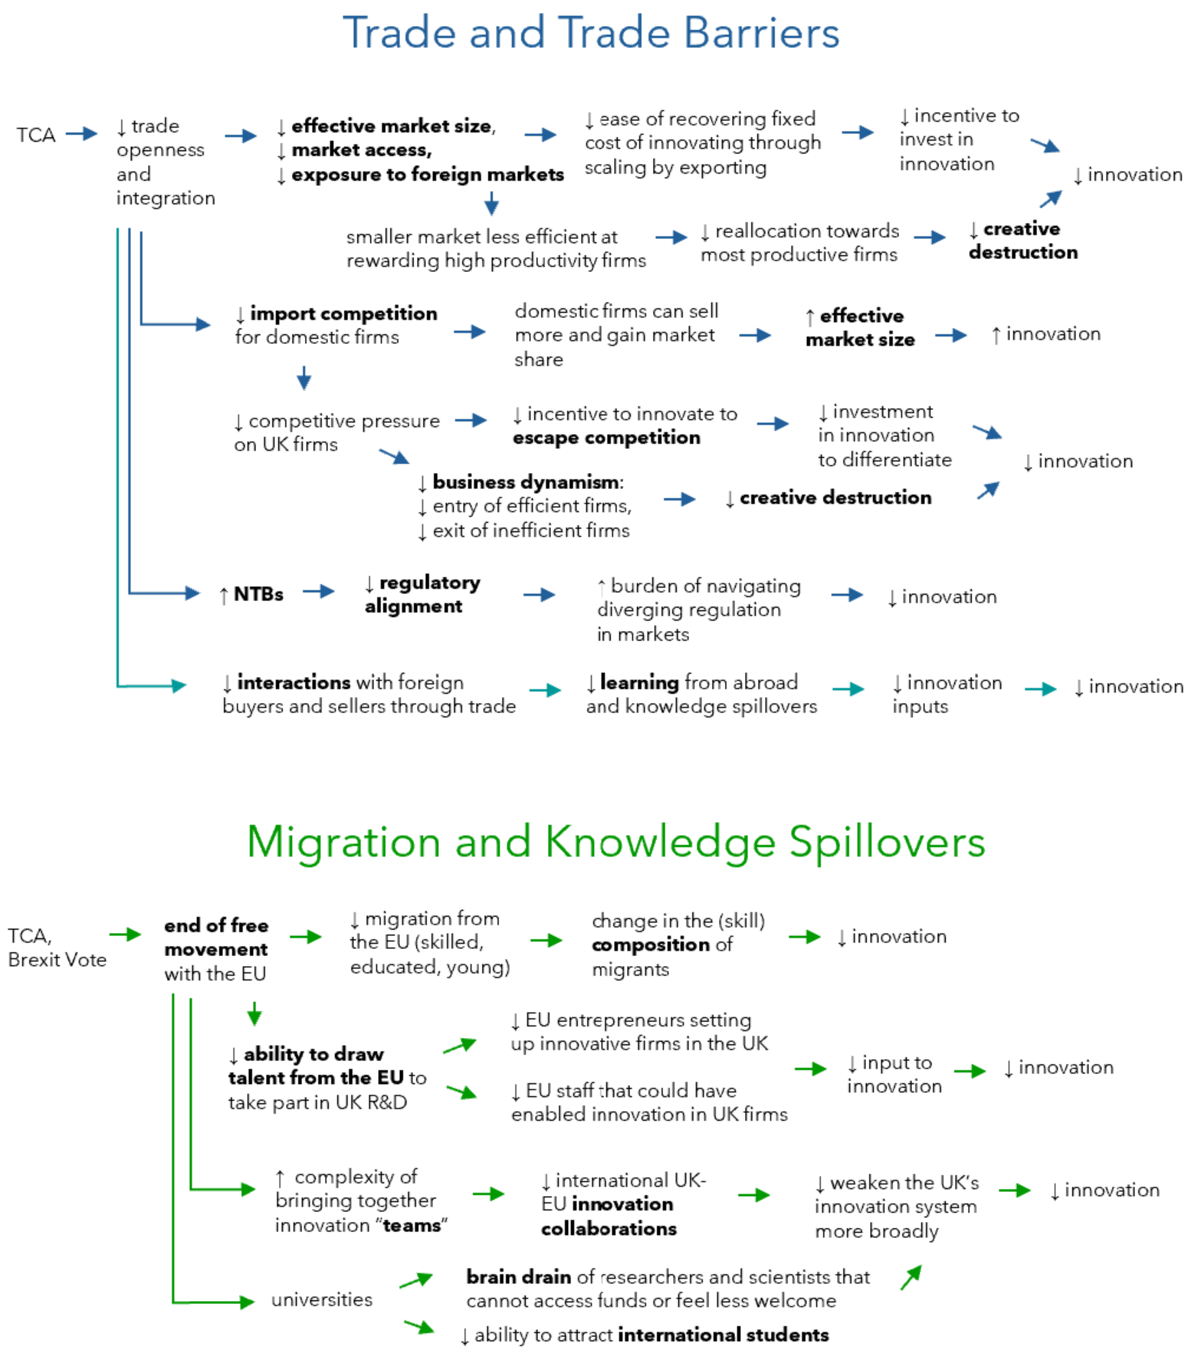

Figure 1: Mechanisms Linking Brexit to Innovation

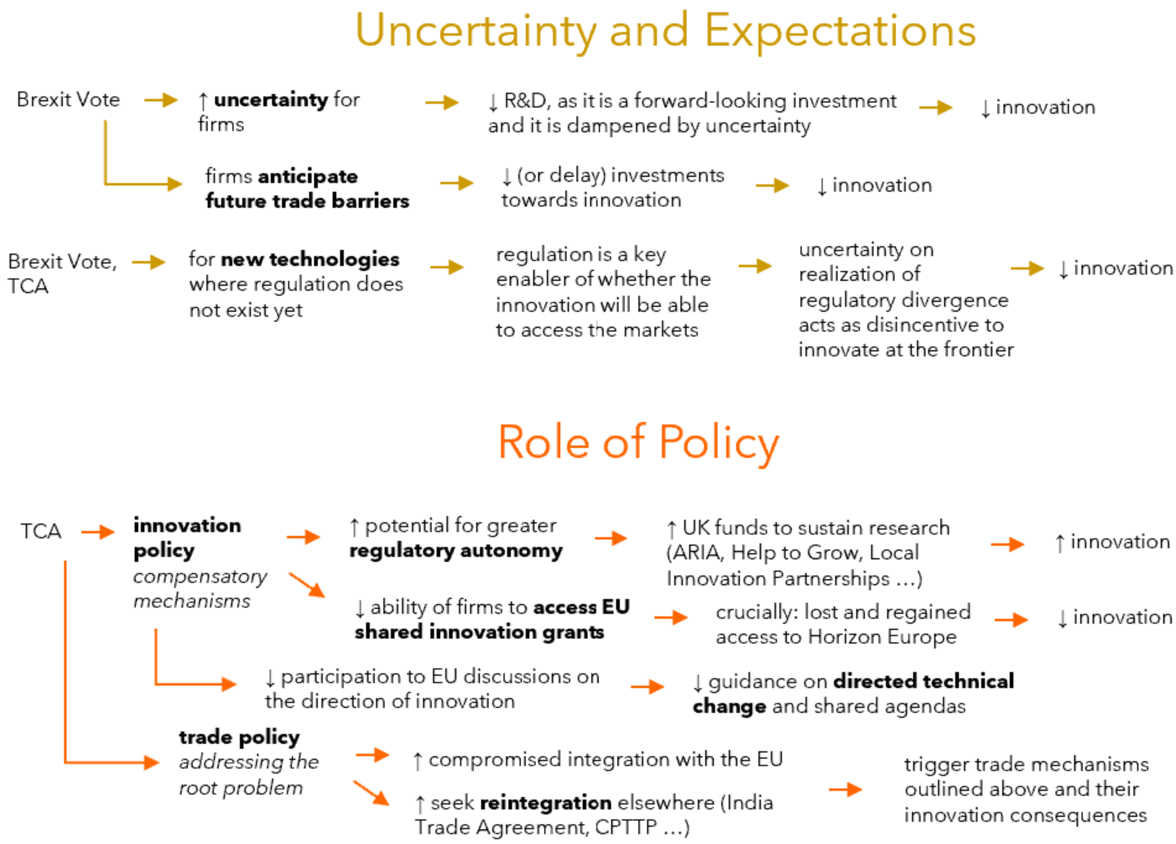

Figure 2: Mechanisms Linking Brexit to Innovation

3 Data and Descriptive Analysis

This chapter turns the research question into an empirical exercise. It describes the data sources, the construction of the datasets and relevant variables, and provides a descriptive analysis throughout. This chapter is organised as follows. **Sections 3.1 and 3.2** describe the building blocks of the dataset in their native classification spaces: innovation in the IPC space and Brexit exposure in the SITC space. Since identification will require both variables to be defined on the same sectoral unit, **Section 3.3** expresses innovation in the SITC space using the Algorithmic Links with Probabilities concordance. Finally, **Section 3.4** considers the merged innovation-trade dataset and provides some first descriptive observations. Throughout the sections that follow, the descriptive analysis will highlight the identifying variation that will be crucial for the later identification exercise.

3.1 Measuring Innovation through Patents

The dependent variable of this study seeks to capture UK-origin, internationally-oriented innovation by technology over time. To do this, this research employs data from the **World Intellectual Property Organization (WIPO)** to obtain **PCT publications by IPC class for the period 2005-2025**[5]. Some remarks are due to understand the meaning and choice of this data.

WIPO's Patent Cooperation Treaty (PCT). Following World Intellectual Property Organization [2025], the Patent Cooperation Treaty is an international intellectual property agreement governed by the World Intellectual Property Organization (WIPO). It was introduced in the 1970s to allow for parallel patent protection in PCT contracting states (now composed of 150 signatory nations). In practice, filing a **PCT application through WIPO acts as a temporary global placeholder** for 30 or 31 months, securing an international filing date (or "priority date") and ensuring rights to seek later protection. During this time, applicants can choose to enter a national or regional phase to seek legal protection for their invention. For instance, applicants can seek european protection at the European Patent Office (EPO), or US protection at USPTO, or Japan protection in JPO. These offices conduct the actual examination of the patent application and decide whether to grant protection. In

[5] IP Statistics Data Center, PCT, Indicator: 6 - PCT publications by IPC class; Report type: yearly statistics; from 2005 to 2025; Origin: United Kingdom; IPC cass: 4 digit, all.

particular, EPO is the regional patent office of the European Patent Organisation. It includes all EU member states, as well as UK, Norway, Switzerland and Turkey. Once an application is filed at the EPO, an examination is conducted and, if granted, the patent becomes a bundle of national patents enforceable in the designated states. In this view, datasets from the two patent authorities capture different things. **WIPO captures inventive activity with global intent**, with no protection granted but with global coverage. This is UK's inventor's willingness to patent internationally regardless of region. Instead, EPO captures inventions actively seeking European Protection. The use of **PCT publications** means that patents with international intent are selected, which targets internationally-oriented innovation. The disadvantage of this measure is that it potentially excludes purely domestic patents - which would be directly filed at the UK's national patent office. However, those patents are unlikely to be affected by trade disintegration and are not of interest for this thesis: it is the firms that were most exposed to EU trade that are likely to seek protection internationally and therefore be impacted by Brexit in their innovation efforts.

PCT Publications. Moreover, the publication stage is the the one that follows the application stage, but precedes the examination and the eventual decision of granting or rejecting the paper. In particular, publication refers to the public disclosure of the patent application, that happens 1 to 2 years after the application [World Intellectual Property Organization, 2025]. This will be pivotal in **accounting for an appropriate lag in the DiD regressions.** Finally, focusing on publications is convenient because it allows enough time to avoid counting "bad application filings", but without getting into grant lag distortions. Ideally, this should capture innovation efforts.

Geography. WIPO collects patent flows data, which is recorded along two main dimensions. The **origin** refers to where the applicant (inventor or firm) is from, while **office or filing route** refers to where protection is sought [World Intellectual Property Organization, 2025]. This thesis selects patents filed by UK residents regardless of where protection is sought.

Period. The data is selected for the years 2005-2025. Recalling the Brexit timeline reported in Section 1.1., the two fundamental events happen in 2016 and 2021 - the Brexit vote and the TCA, respectively. In this view, the chosen time window allows for 11 years of pre-period before the Brexit vote, 9 years post-vote, and 4 years after the TCA. This should allow for enough observations to conduct a credible Difference in Differences analysis around both target years.

4-digit IPC Subclass. The International Patent Classification (IPC) is the common classification for patent documentation that entered into force in 1975 [World Intellectual Property Organization, 2023]. It functions through a hierarchical structure of sections, classes, subclasses, groups, and complete classification symbols. **Table 1** summarises the structure of the classification. This thesis operates at the subclass level (4 digit IPC code), which is the finest granularity for which a credible link with trade classifications can be carried out. A 1-digit IPC analysis will be provided in the descriptive statistics section to gain intuition on the evolution and composition of patents over time.

Level	**Digits / Symbol**	**Categories (approx.)**	**Example (code : description)**
Section	1 letter	8 (A–H)	H : Electricity
Class	3 characters	129	H01 : Basic electric elements
Subclass	4 characters	∼650	H01S : Devices using light amplification (LASERs)
Main group	5-8 characters /00	∼7,000	H01S 3/00 : Lasers
Subgroup	5-8 characters /xx	∼70,000	H01S 3/14 : Characterised by the material used as the active medium

Table 1: Structure of the International Patent Classification (IPC)

3.1.1 Patents in the innovation space: Descriptive Analysis

Having understood the measure of **patent counts by IPC subclass**, this section provides several descriptive statistics on the dependent variable.

Table 2 provides some descriptive statistics for the patent subsections and for the total yearly patent counts. It can be observed that there is a strong heterogeneity in patenting activity across IPC 4-digit classes, differently from the relatively smooth evolution of aggregate patenting over time. At the IPC4–year level, the mean number of patents is 13.7, but the standard deviation (37.9) is nearly three times as large, with observations ranging from zero to 712 patents. This indicates a highly skewed distribution, where many technology–year cells record little or no activity while a small number of fields account for very large patent counts. Instead, the annual totals are more stable; over 2005–2025, PCT publications average 8,678 per year with a standard deviation of 885, fluctuating between 7,374 and 10,045. The relatively moderate dispersion in the aggregate series suggest that any Brexit-related effects are unlikely to be visible in a simple time-series aggregate and instead motivate exploiting cross-technology heterogeneity in exposure and responses.

Variable	**Obs.**	**Mean**	**Std. Dev.**	**Min**	**Max**
Patents (IPC 4-digit $\times$ year)	13,272	13.73	37.93	0	712
Total patents (yearly)	21	8,678.10	885.22	7,374	10,045

Table 2: Descriptive statistics for PCT patent publications, 2005–2025

It is useful to understand the general trend of patents in the time series. **Figure 3** displays the yearly counts of UK-origin PCT publications between 2005 and 2025. PCT publications exhibit a clear **upward trend**, rising from roughly 7,400 to around 10,000 annual filings, indicating sustained growth in internationally oriented inventive activity. This upward trend will motivate the use of year fixed effects in the Difference in Differences regressions. Furthermore, two main dips are visible in the time series. A decline around from 2008 aligns with the global financial crisis, while a similar change in 2022 can be attributed to the COVID pandemic. Both occurrences are coherent with the cyclical sensitivity of innovation to macroeconomic shocks. The dashed vertical lines indicate the 2016 referendum and the 2021 TCA - as the magnitude of the 2021 dip is too large to be attributed to Brexit effects, while it is coherent with the pandemic's disruptions. This graph suggests that any Brexit-related effects are unlikely to appear as an immediate aggregate collapse, reinforcing the **need to focus on sectoral heterogeneity** rather than aggregate trends alone.

In view of understanding the composition of patents by different fields of technology, **Figure 4** provides the evolution of patent counts by the 8 IPC sections (A-H). The section-level breakdown reveals that UK PCT activity is concentrated in Human Necessities (A) and Physics (G), followed by Chemistry (C) and Performing Operations (B), while Textiles (D) and Fixed Constructions (E) remain comparatively small throughout the sample. **Figure 5** plots contributions of each IPC section to each year's total patent counts. The compositional shares of patenting across broad IPC sections remain relatively stable over the sample period, with Human Necessities (A) persistently accounting for roughly one quarter of total PCT publications.

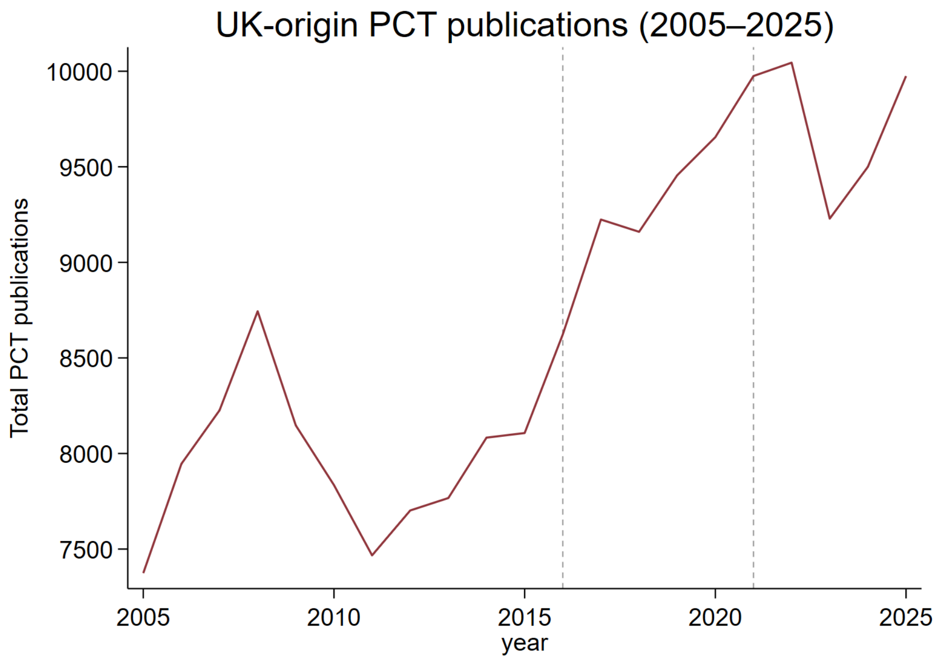

Figure 3: Time series of total patent counts

It is important to clarify how the patent data should be interpreted at this stage. The descriptive evidence presented above is organised along technological classifications (IPC), which group inventions by their scientific and technical domain. However, Brexit exposure is defined along trade classifications (SITC), capturing differential dependence on EU trade across sectors. As a result, because of the current "cut" of the patent data, these preliminary descriptive figures are not meant to show any meaningful dips at Brexit-relevant event dates. Rather, they provide contextual evidence on rising aggregate innovation trends and a rather stable technological composition. Meaningful interpretation of Brexit effects requires re-aggregating patent data to the SITC level, where variation in Brexit exposure is defined. Finally, all the above suggest that the Brexit story is worth exploring at the sector-specific level, rather than the aggregate.

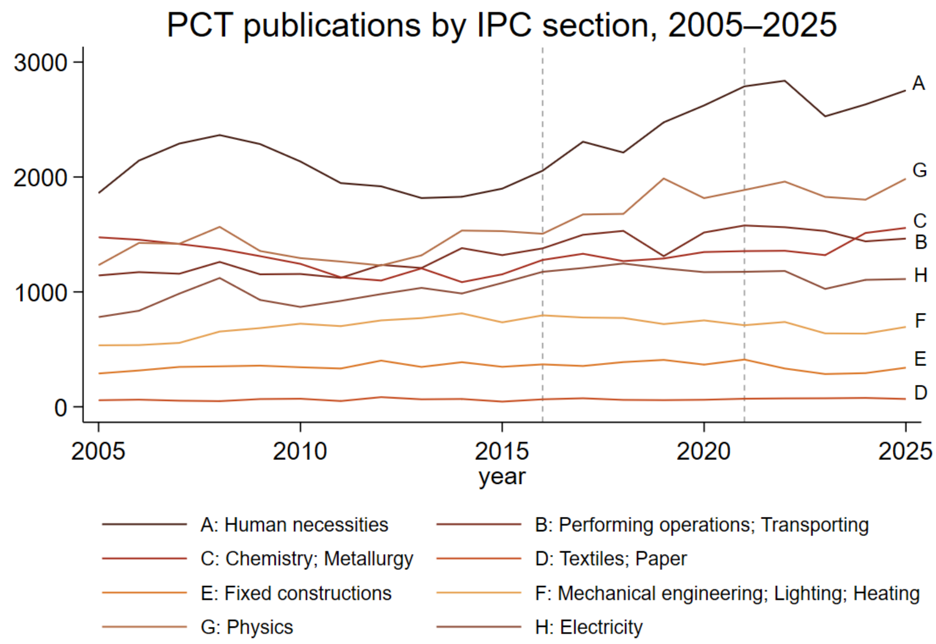

Figure 4: Time Series of Patent Counts by Sector of Technology

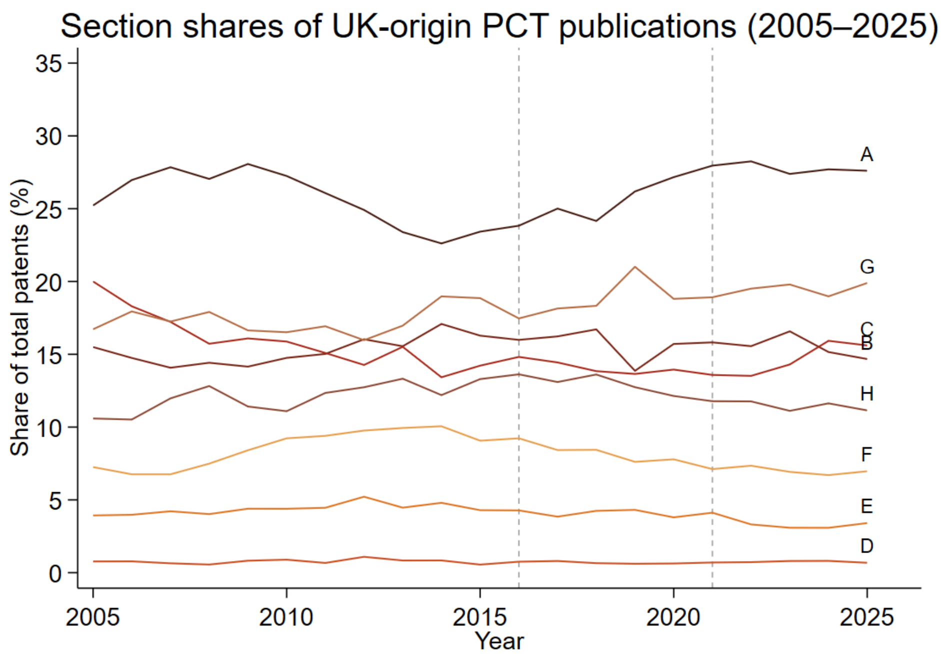

Figure 5: Evolution of Sector Shares of Patents

3.2 Measuring Exposure to Brexit

The independent variable of this study is a measure of sectoral exposure to Brexit. Following and adapting from Bakker et al. [2022], this thesis constructs a sector-level indicator based on the **UK's share of imports sourced from the EU in 2015**, i.e. the pre-referendum EU trade share. In particular, the exposure to Brexit of sector s, reads:

$$\text{BrExposure}_s = \frac{\text{Imports}_{s,\text{EU},2015}}{\text{Imports}_{s,\text{World},2015}}$$

There are two crucial features to note about the use of this variable.

First, using exposure to EU trade **a year prior to the referendum** is crucial to ensure the variable is predetermined, rather than mechanically affected by post-Brexit adjustments in trade patterns. In turn, the latter is central to the credibility of identification in the DiD design. It is also important to notice that Exposure to Brexit is a variable that is by construction continuous and bounded in the interval [0,1]. Having such **continuous treatment variable** motivates the methodology drawing from the continuous DiD literature.

A second remark concerns the correct interpretation of this variable. The rationale for focusing on import exposure — rather than exports — is that innovation is often strongly input-dependent. Firms rely on imported intermediate goods, components, and embedded technologies to sustain their innovative activity. Brexit primarily raised border and regulatory frictions affecting cross-border supply chains; as such, sectors that were more reliant on EU imports were plausibly more exposed to disruptions in the inputs underlying production and innovation. Ultimately, this measure captures something economically meaningful: it reflects **trade dependence at the product-sector level**, which is closely correlated with value-chain integration. Sectors that import a substantial share of final goods from the EU are often the same sectors that rely on EU suppliers for intermediate inputs. The measure should therefore be interpreted as a **proxy for how deeply a UK sector was "entangled" with EU trade prior to Brexit**.

The following sections describe the data cleaning process for the construction of the variable, and provide some descriptive statistics on the final measure of Exposure to Brexit.

3.2.1 Data and Data Cleaning

To measure Brexit, this thesis uses the UK's share of imports from the EU in 2015. Data on UK imports from the EU and from the World is taken from **UN Comtrade** [6]. The paragraphs that follow will motivate several choices stemming from the degrees of freedom of this data.

Working with import values rather than quantities. Exposure to Brexit wishes to capture how dependent each UK sector was on imports from the EU before the referendum. This dependence relates to economic significance for Brexit, namely how much each sector relied on EU inputs. This dimension is captured by the total monetary value of imports, it is comparable across products of very different physical units, and it matches how firms feel trade shocks - that is, in monetary terms rather than in kilos. The only disadvantage of working with values is that the measure is sensitive to price changes (e.g. exchange rate, inflation, commodity spikes). However, this is mitigated by the fact that this thesis employs this value as a pre-Brexit baseline rather than a time series.

Working with CIF value rather than primary value. UN Comtrade reports import values in two forms. The CIF value is the value of imported goods, including (i) the cost of the goods themselves (ii) the insurance cost during transport and (iii) the freight/shipping cost to the importing country's border. It measures how much the importer actually pays to get the goods delivered to the cost of entry. This is the standard for valuing imports in international trade statistics. On the other hand, the Primary value (or, FOB value) is defined as the base customs value of the goods, excluding insurance and freight. It coincides with the cost of the goods up to loading onto the transport vessel in the exporting country. Since this research takes a UK perspective, we use the CIF value because it represents what the importer pays when a good is delivered to its border.

Currency choice. Import values on UN Comtrade are reported in US dollars, rather than in GB Pounds. However, this is irrelevant in view of taking a ratio between EU and World imports to compute Exposure to Brexit, where the unit of measurement will simplify nonetheless.

[6]UN Comtrade data was seleted as follows: goods; annual; SITC S4 classification; AG4 All SITC rev 4, 4 digit codes; reporter: UK; Trade flows: imports; Period: 2015; Partners: World, EU countries individually selected excluding the UK; Modes of transport: total; 2nd Partner: World; Custums codes: total; breakdown: classic; aggregate by: none.

European Union aggregate. One issue arising when working with UN Comtrade is that a European Union aggregate is unavailable. Hence, the 27 countries of the EU (excluding the UK) in 2015 were manually selected and aggregated to form an EU aggregate. Other sources, such as EUROSTAT, have such aggregate but have some caveats on sectoral classification codes and descriptions that made EUROSTAT data less suited for this project.

Sectoral classification. Having described the various dimensions of this data, it is worth having a small digression about the sectoral classification that will be used throughout this thesis. The **Standard International Trade Classification (SITC)** is a system developed by the United Nations to classify internationally traded goods [7]. The current version is SITC Revision 4, published in 2006. The classification follows a hierarchical structure and SITC codes can be expressed through different levels, each corresponding to a 1 to 5 digits classification [United Nations Statistics Division, 2006]. **Table 3** summarises the structure, and provides an intuitive example. The fact that there can be code ranges such as 00–98 but only 67 active categories implies that SITC codes act as structured placeholders, rather than a continuous numeric sequence.

Level	Digits	Categories (approx.)	Example (code : description)
Section	1	10 (0–9)	7 : Machinery and transport equipment
Division	2	67 (00–98)	71 : Power-generating machinery and equipment
Group	3	262 (001–984)	711 : Steam or other vapour generating boilers
Subgroup	4	1,012 (0010–9840)	7112 : Steam and other vapour generating boilers
Item	5	3,993 (00100–98400)	71121 : Boilers for power generation

Table 3: Structure of the Standard International Trade Classification (SITC Rev. 4)

[7] A useful guide to **SITC rev4** : https://unstats.un.org/unsd/publication/SeriesM/SeriesM_34rev4E.pdf

SITC data cleaning. This research utilises the **subgroup level (SITC 4 digit)**, while keeping track of the more aggregate categories for consistency and for sectorial classifications. An aspect that may seem confusing is that not all 4-digit SITC codes are composed of 4 digits. For instance, a code made of 2 digits can be marked 4-digit because SITC has no finer aggregation below it, for that specific product. For example, code 11 (Bovine animals, Live) appears in the SITC 4 digit aggregation, as it is effectively its final detail. It is part of Section 0 (Food and Live Animals), Division 00 (Live Animals Other Than Animals of Division 03), Group 001 (Live Animals Other Than Animals of Division 03), Subgroup 0011 (Bovine animals, live). However, UN COMTRADE returns it as "11" with aggregation "4". Hence, some cleaning must be done to homogenise the 4-digit codes to be actually composed of 4 digits. Otherwise, aggregating upwards would become problematic because it would sum across categories based on a potentially wrong parent. In the above example, it would sum across the the codes beginning with "11", while it should sum across the codes beginning with "001". Thus, the **cleaning homogenised the codes** to match the number of digits implied by the aggregation level reported.

As explained, the data was downloaded at 1- 2- 3- and 4-digit SITC levels. While the unit of analysis is 4-digit codes due to granularity, the more aggregate levels are useful for interpretation and graphical purposes. However, the **additive consistency** between these categories is imperfect. Additivity is maintained at fine granularities (4 to 3, 3 to 2) but tends to get lost from the 2-digit to the 1-digit aggregation. This could be due to confidentiality suppression, rounding, currency conversions, direct SITC data submissions, or adjustments for consistency with national accounts. To decide how to approach this inconsistency in aggregating, it is useful to understand if Comtrade's SITC data is built bottom-up or top-down. National authorities measure trade data at the Harmonised System (HS) level, and report it to UN Statistics Division. UN Comtrade then converts HS to SITC 4-digit, and then aggregates it up to 3-digit, 2-digit and, ultimately, 1-digit SITC. Since the **aggregation is bottom-up**, this research will prefer to maintain internal consistency by operating at the 4-digit level and aggregating it up when needed for interpretation or graphical intuition.

3.2.2 Exposure to Brexit: Descriptive Analysis

Having constructed the Brexit Exposure measure at the SITC 4-digit level, this section provides several descriptive statistics on the variable.

Table 4 provides a univariate analysis of the Brexit Exposure variable at the SITC 4-digit level. Panel A shows that, after the cleaning process, we are left with a sample of 1005 subgroups (i.e., products) for which we have an associated Brexit exposure measure. The average exposure across subgroups is 57%, meaning that on average, 57% of pre-referendum UK imports in a given SITC 4-digit product category came from the EU. This signals a high integration of the UK with EU supply chains. Moreover, both values of 0 and 1 are attained, implying that some products were not imported at all from the EU, while others were fully imported from the EU. The substantial variability of the measure is also confirmed by the 0.27 standard deviation. Panel B shows the distribution of the Exposure measure by 4 intensity categories. We observe that products are substantially heterogeneous in their exposure to Brexit, although the distribution seems slightly skewed towards the higher end of exposure intensities. Again, this points to substantial pre-Brexit integration with EU supply chains.

Table 4: Distribution of Sector-Level Brexit Exposure (SITC 4-digit subgroups)

Panel A: Summary Statistics		**Panel B: Exposure Categories**		
Statistic	Value	Category	Count	Percent
Observations (N)	1,005	Low (< 0.25)	155	15.42
Mean	0.571	Mid-low (0.25–0.50)	249	24.78
Std. Dev.	0.272	Mid-high (0.50–0.75)	295	29.35
Min	0.000	High (≥ 0.75)	306	30.45
Max	1.000			

Figure 6 provides a graphical version of the statistics reported in Panel B, again suggesting substantial variability and a tendency towards higher Brexit exposures. **Figure 7** provides an alternative frequency histogram with a finer partition of bins, and a fitted Kernel density.

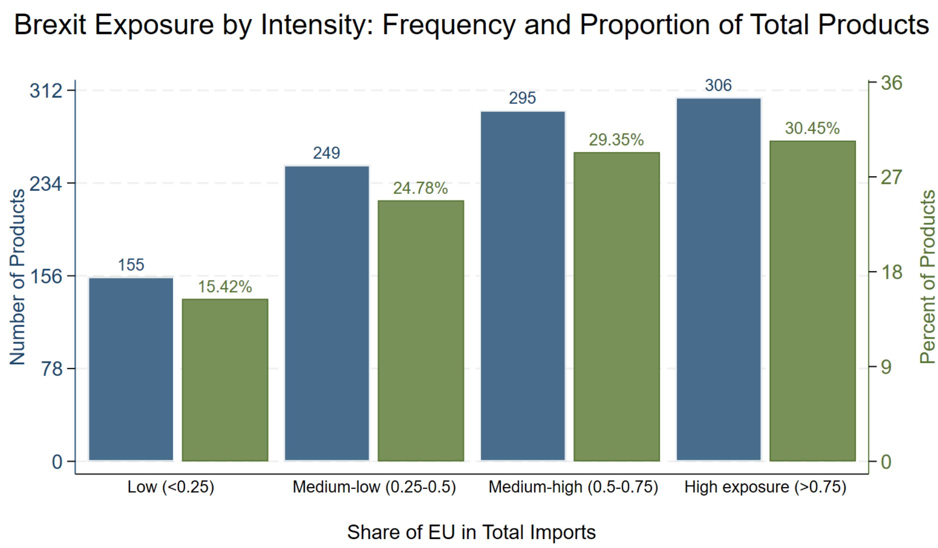

Figure 6: Brexit Exposure by Intensity: Frequency and Proportion of Total products

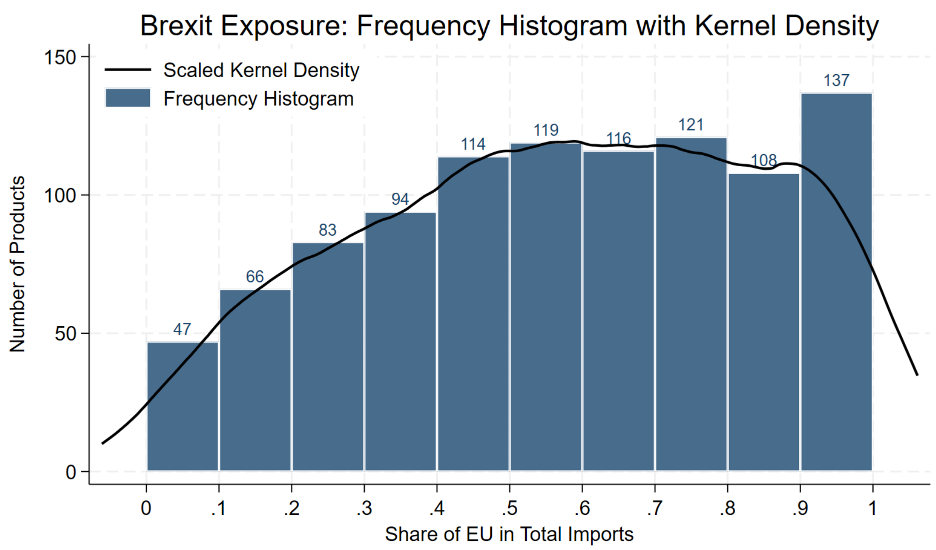

Figure 7: Brexit Exposure: Frequency Histogram with Kernel Density

While the above discussion suggests an **encouraging variability of the independent variable at the 4-digit product level**, a different exercise must be carried out to evaluate its suitability for identification purposes. The key is to understand that through this analysis, we will use low exposed products as counterfactuals for highly exposed products. Yet, maintaining a pooled vision of the 1005 4-digit products risks comparing a bicycle with a bovine. As will be explained in the methodology section, continuous DiD relies on parallel trends (in their standard version *and* strong version). To maintain parallel trends, this thesis has to rely on within sector identification, i.e. comparing the behavior of different products within the same sectors. Intuitively, rather than comparing a bicycle with a bovine, we would like to compare a bovine with a sheep. In this view, **Figure 8** separates the 1005 4-digit products into their respective 10 1-digit sections, namely the coarsest partition of the SITC classification. We will assume that (strong) parallel trends holds *within* each of the 10 sections. We can conclude two things from observing the histograms by sector - which is to be inspected in conjunction with **Table 5**. First, we observe that some SITC sections have insufficient data for inferential purposes and will likely be dropped (for instance, section 9: commodities and transactions n.e.c.; potentially section 1: beverages and tobacco). The second thing we can observe is that different sectors display different distributions of Brexit exposure of their products. For instance, sections 5,6,7 and 8 are closer to normal distributions centered around different means. Others are skewed towards higher exposures, namely sections 0 and 4, while the rest are closer to being uniformly distributed, such as section 2. These shapes will become important to consider in light of identifying Level Treatment Effects and Causal Responses, which are later discussed in the methodology section.

SITC 1-digit	Description	N	% of Total
0	**Food and live animals**	131	13.03
1	Beverages and tobacco	11	1.09
2	**Crude materials, inedible, except fuels**	107	10.65
3	Mineral fuels, lubricants and related materials	22	2.19
4	Animal and vegetable oils, fats and waxes	21	2.09
5	**Chemicals and related products, n.e.s.**	129	12.84
6	**Manufactured goods classified chiefly by material**	229	22.79
7	**Machinery and transport equipment**	213	21.19
8	**Miscellaneous manufactured articles**	139	13.83
9	Commodities and transactions n.e.c.	3	0.30

Table 5: SITC 1-digit sectors: descriptions, number of products, and share of total (N = 1005).

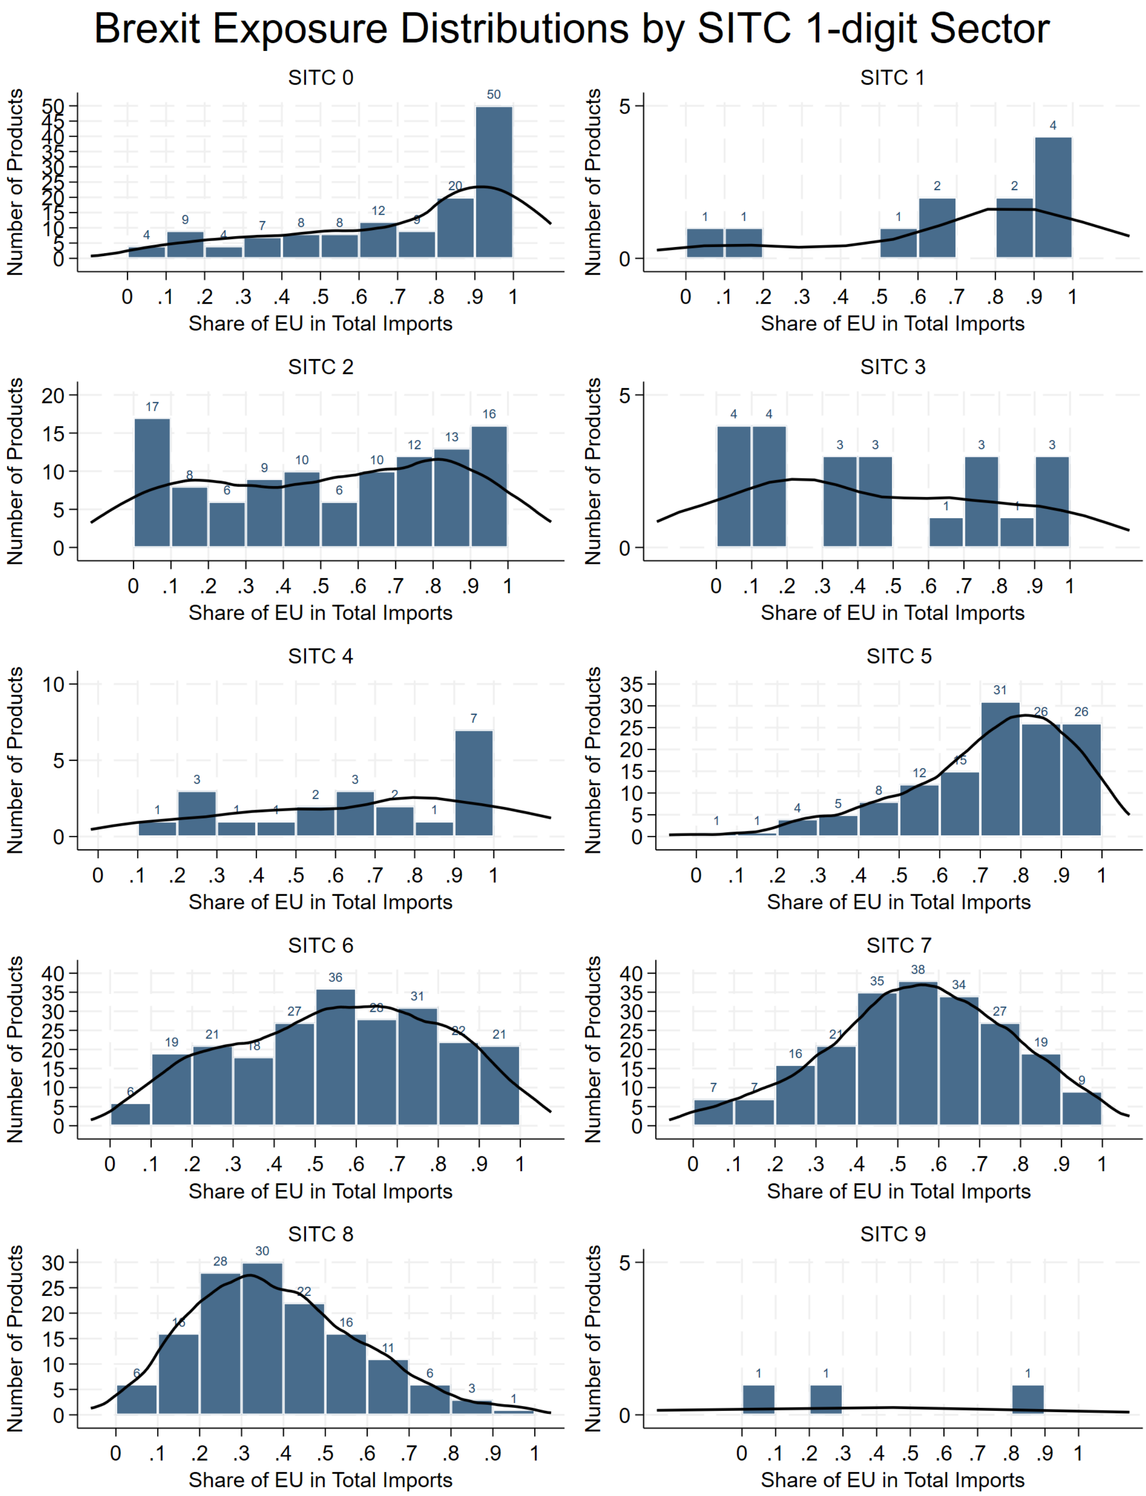

Figure 8: Brexit Exposure by SITC1 Sector

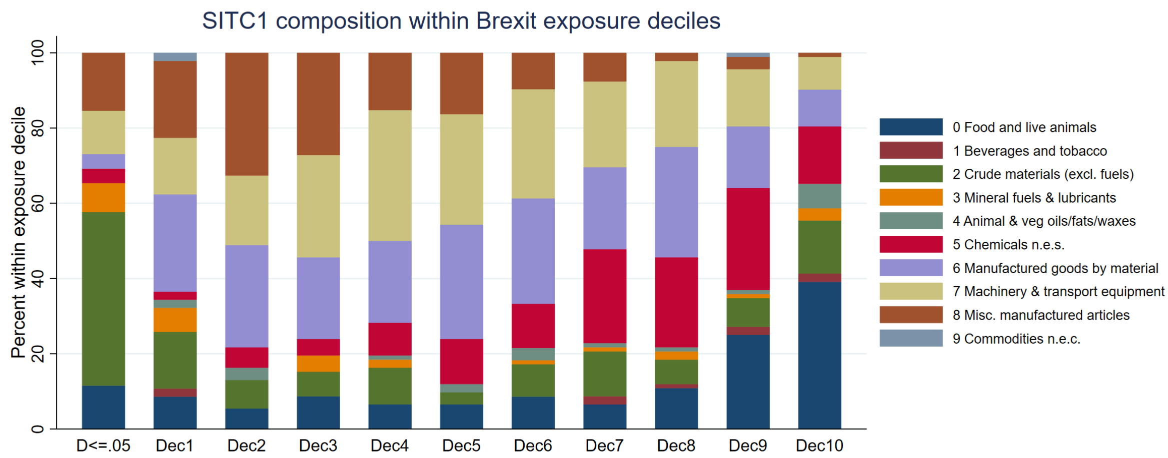

Figure 9: SITC 1 Composition of Brexit Exposure Deciles

Figure 9 shows that the sectoral composition of products varies markedly across Brexit exposure deciles. Low-exposure products are heavily concentrated in primary and resource-based sectors, particularly crude materials and related categories, while higher exposure deciles increasingly comprise manufacturing activities. In particular, machinery and transport equipment, manufactured goods by material, and miscellaneous manufactures expand steadily as exposure rises, with machinery becoming especially dominant in the upper deciles. By contrast, food and live animals and other low-tech sectors are disproportionately represented at the bottom of the exposure distribution. This systematic re-sorting of SITC1 sectors across exposure deciles indicates that exposure intensity is strongly correlated with broad sectoral composition. As a consequence, identification strategies that rely purely on cross-decile comparisons risk conflating exposure effects with underlying sectoral differences. This pattern strongly suggests that doing **identification within SITC1 sectors** — or controlling flexibly for SITC1×time effects — is important to ensure that estimated treatment effects are not driven by structural differences in sectoral composition rather than by Brexit exposure itself.

3.3 Linking trade and innovation data

This section explains the process of linking trade and innovation data, by operating on patent data to change their classification space. In practice, it is an exercise of translating the IPC technological space into the SITC trade space. This is necessary in order to subsequently merge the patent dataset with the exposure to Brexit dataset. To do this, this thesis follows the Lybbert and Zolas [2014] approach of Algorithmic Links with Probabilities, which the next section outlines.

3.3.1 Algorithmic Links with Probabilities

Linking trade data with innovation data is non-trivial. The literature uses several different concordances, which can be broadly categorised into two separate categories. On one hand, the Yale Technology Concordance (YTC) provides a mapping between IPC and cSIC [Kortum and Putnam, 1997]. The positive aspect of YTC is that it is based on theoretical underpinnings. However, YTC is widely criticised for being outdated, as it is built on a narrow set of Canadian patents frozen in time, space, and technology. The same category contains concordances that are derived from YTC. For instance, the OECD Concordance [Johnson, 2002] emerged to link IPC and ISIC, yet relied on YTC and simply added additional layers of concordance. Thus, such derived concordances are subject to the same limitations of YTC, and the additional concordance layer causes the strength of the technology-industry linkage to atrophy. A second category of concordances emerges from an empirical approach based on Bayesian inference. This project follows this route, by using the **Algorithmic Links with Probabilities** methodology by Lybbert and Zolas [2014]. ALP works as follows: (i) extracts keywords from SITC official industry descriptions (ii) uses data mining on PATSTAT to collect all patents that match the SITC keywords, and produce an SITCxIPC frequency matrix (iii) converts frequencies into probabilities using Bayes' rule, while defining three alternative weighting schemes for the probability of IPC_j conditional on $SITC_i$ - this thesis employs probability weights (iv) defines a cutoff probability to drop low probability links, and normalises the remaining weights.

An example can aid intuition on this methodology. Say a patent abstract mentions "Protective helmet with improved shock absorption for cyclists", and the patent is classified in IPC subclass A42B (Hats, head coverings); ALP text-mines SITC descriptions using keywords such as "helmet, headgear", and it finds SITC subgroup 8484 (Headgear and fitting thereof);

mining is applied to all patens by repeating the above process; a frequency table for SITC 8484 is built, reporting all the IPC codes it is associated with (say, IPC A42B, A62B, F41H) and the number of matching patents by IPC (say, 800, 100 and 50, respectively); retrieve a weight for each of the three IPCs (say, 0.75, 0.15, 0.1, respectively); this allows to claim that 75% of technology in SITC 8484 comes from IPC A42B; repeating this process for all SITC codes, a concordance is built for SITC-IPC. This summarises the ALP logic.

3.3.2 Patents in the Trade Space: Descriptive Analysis

This section wishes move the focus from the question "how many patents are in technology subclass i?" to "how many patents are associated with trade sector s?". In particular, ALP considers an IPC_j and associates it to a set of tuples made by pairs of $SITC_i$ and a corresponding weight w_i that sums to 1 within each IPC class. The idea is that one patent gets fractionally allocated across trade sectors. The mapping therefore reads: $IPC_j \to \{(SITC_1, w_1), (SITC_2, w_2), \ldots\}$. Taking the concordance as given, the following cleaning steps were performed on the WIPO dataset described in **Section 3.1**.

1. **Many to many expansion.** Merge the patent dataset (IPC subclass/year/patent count) with the ALP concordance for IPC4 to SITC4; For example, if IPC A01B maps to 2 distinct sectors, one patent will become two rows.

2. **Apply fractional allocation.** Multiply the patent count by the weight corresponding to the associated SITC sectors; For example, if IPC A01B has 10 patents in year 2005, and it maps into SITC 7112 and SITC 7113 with weights 0.4 and 0.5 respectively, then SITC 7112 will be assigned 4 patents while SITC 7113 will be assigned 6.

3. **Aggregate to sector-year.** Sum patent counts associated with the same SITC sector, obtaining a clean SITC-year panel with patent counts (SITC subgroup/year/patent count).

4. **Patent counts variable.** Two distinct patent count variables are created. The raw measure is that of patent counts by SITC sector, by year. The second measure proposed is a classic log transformation of the first measure, namely, $log(1 + patents)$. A remark on the rationale of this transformation will follow shortly.

The above steps allow to obtain a measure of innovation associated with trade sectors. Note that patent counts become a **measure of patent intensity**, a fractional measure obtained by the IPC mass being distributed across many SITC sectors. In what follows, some descriptive statistics on patent counts by trade sectors are reported.

Table 6 report some descriptive statistics for the two outcome variables, namely patents and log-patents by SITC and year. On average, a SITC4 sector-year cell is associated with about 9 patents, which is quite large considering the narrow definition of a 4-digit SITC. The key, however, is that the standard deviation exceeds the mean by more than 3 times, suggesting the presence of a strong right skew. This means that some sectors will have very high patent mass, yet the majority will have low patent mass - which is typical for innovation data. This intuition is also confirmed by the sample range: while the minimum patent count attained per sector is null, the maximum is 590 patents, which is extremely high relative to the mean. Note that these extreme observations may be problematic in view of a regression, as it means that some sectors will mechanically have a larger influence in level regressions and in OLS. This motivates the employment of the log-transformed variable to compress the range and reduce the mechanical influence of very large counts. In fact, log-patents retain a standard deviation that exceeds the mean, without altering the right-skew of the data, but reducing the extent of it by stabilising the variance. This can be also noticed by the reduced sample range, that is now between 0 and 6. **Figure 10** is useful in visualising the result of this transformation. The left panel reports a histogram of the raw patent counts, while the right panel reports an analogous graph for the log variable. In both cases, patent activity is highly skewed across SITC4 sectors, with a small number of sectors accounting for a disproportionately large share of innovation. However, the figure in logs stabilises the variance and thickens the tail of the distribution, hopefully reducing the mechanical influence of large realisations. The remainder of the analysis of this section will use log-patents.

Variable	**Obs.**	**Mean**	**Std. Dev.**	**Min**	**Max**
Patents (SITC4 $\times$ year, weighted)	20,265	8.92	29.55	0	589.57
Log patents ($\log(1 + \text{patents})$)	20,265	1.16	1.28	0	6.38

Table 6: Descriptive statistics for sector-level patent counts (SITC4 $\times$ year), 2005–2025

In view of using patent counts as a dependent variable in a continuous difference-in-

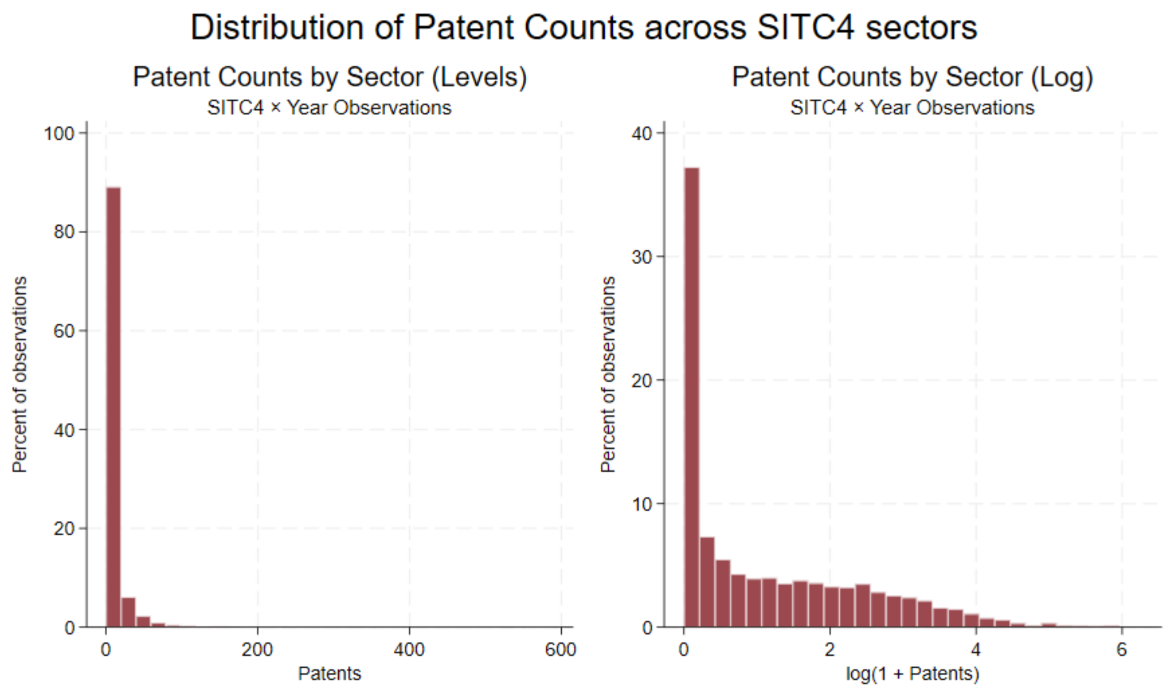

Figure 10: Distribution of patent counts by SITC4

differences regression, it is useful to zoom into the **variability of the outcome**. Namely, disentangling the overall variation of patent counts in between- and within-sector variation. In particular, as will be explained in the methodology section, this thesis will regress patent counts of sector s at time t on sector fixed effects, time fixed effects, and the interaction of Brexit exposure of sector s with a "post" variable. In a panel like the one constructed, there are two sources of variability in the outcome.

First, there is **cross-sectional (CS)** variation, that, fixing a year, compares the patent counts of different sectors. Until now, descriptives and histograms in this section have relied on the pooled patent dataset, where one observation referred to a SITC-year pair. To isolate cross-sectional variation, the left panel of **Figure 11** reports the distribution of patent counts for three representative years, namely 2005 (beginning of sample and pre-financial crisis), 2015 (before the Brexit vote) and 2022 (after the TCA). The kernel densities indicate persistent cross-sector heterogeneity in innovation intensity, with a reasonably skewed distribution coherent with the prior analysis. Another interesting fact is that there is a gradual rightward shift of the distribution and a thickening of the right tail, which reflects an overall growth in patenting over in the last 20 years. To conclude, cross-sectional variation is substantial as expected in innovation data, but this will be absorbed by sector fixed effects.

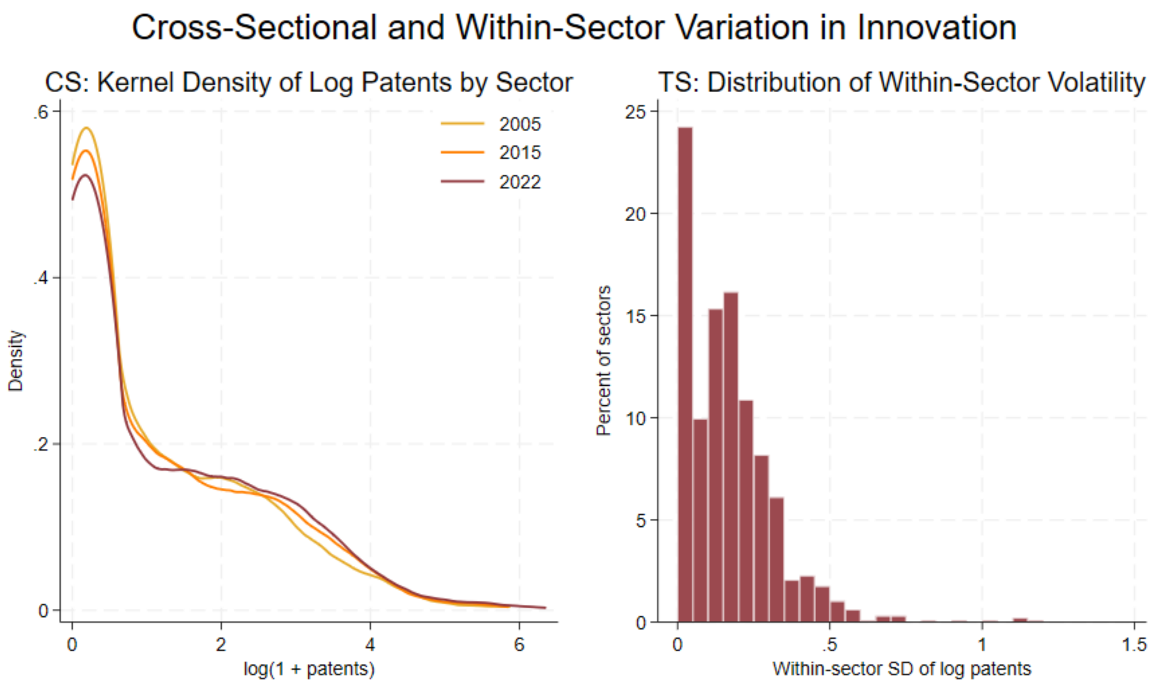

Figure 11: Cross-sectional and Time-series variation of innovation

The second source of variation is **time-series (TS)** variation within sectors. That is, fixing a sector, the variation of patent counts over time. This is precisely the variation that the regression will exploit, trying to inspect whether sectors highly exposed to Brexit change their patenting behavior after Brexit, differently from low exposed sectors. To capture this, the right panel of **Figure 11** plots the distribution of the standard deviation of within sector patent counts over time. While the variability is visibly in favour of lower values, sufficient variation is exhibited to credibly carry out a regression. Sectoral patent counts are expected to be mostly persistent in time, which makes the latter result unsurprising.

A last exercise can be performed to disentangle variability. **Table 7** reports a statistical decomposition [8] of the overall variation of patent counts. The between-variation measures how much sectors differ from each other on average (collapse each sector to its average over time, then compute variance across sectors). The within variation measures how much sectors move over time around their own average (for each sector, subtract its mean, and compute the variance of those deviations). DiD removes between variation through sector fixed effects, common time ariation through year fixed effects, and identify the causal coefficient of interest

[8] This can be done in STATA using the xtsum command.

using within-sector time variation that differs by Brexit exposure. This decomposition reveals that almost all variation comes from persistent differences across sectors, which is normal for innovation data. However, 17% of variation is within-sector over time, which is economically meaningful and hopefully sufficient for credible identification.

Measure	**Std. Dev.**	**Sectors (n)**	**Years (T)**
Overall variation	1.28	965	21
Between-sector variation	1.26		
Within-sector variation	0.22		

Table 7: Variance decomposition of log patent counts (SITC4 $\times$ year), 2005–2025

A last contextual result can be obtained by plotting patent counts by broad SITC 1-digit classification. This should virtually mirror the information contained in Figure 5, but breaking down the patents by trade classification rather than by technology classification. The idea is that identification is carried out within each of these lines: fixing a SITC 1-digit, we will compare the behavior of the SITC 4-digit nested inside it. **Figure 12** shows the time series. The results are intuitive except for one outlier. The most innovative trade sectors are Machinery and Transport Equipment, Chemicals and Related Products, and Miscellaneous Manufactured Articles. The least innovative sectors are Animal and Vegetable Oils and Fats, Food and Live Animals, and Mineral fuels and oils. One outlier is the exceptional growth in patenting for Beverages and Tobacco, which shows a surprising surge in innovation activity. This is not an artifact of the log-transformation, and a version in levels of this graph can be found in **Figure 21** in the Appendix. The spike around 2017 is due to an explosion in tobacco-related technology, for instance related to the use of e-cigarettes or related products. This is evident by plotting the time series for the Beverages and Tobacco sector, disaggregated by the underlying products. **Figure 22** in the Appendix clearly shows that the trend is uniquely driven by subgroup 1222, namely "Cigarettes containing tobacco". Notably, this includes IQOS that introduced the heated tobacco technology in 2016, and was quickly followed by competitors in subsequent years.

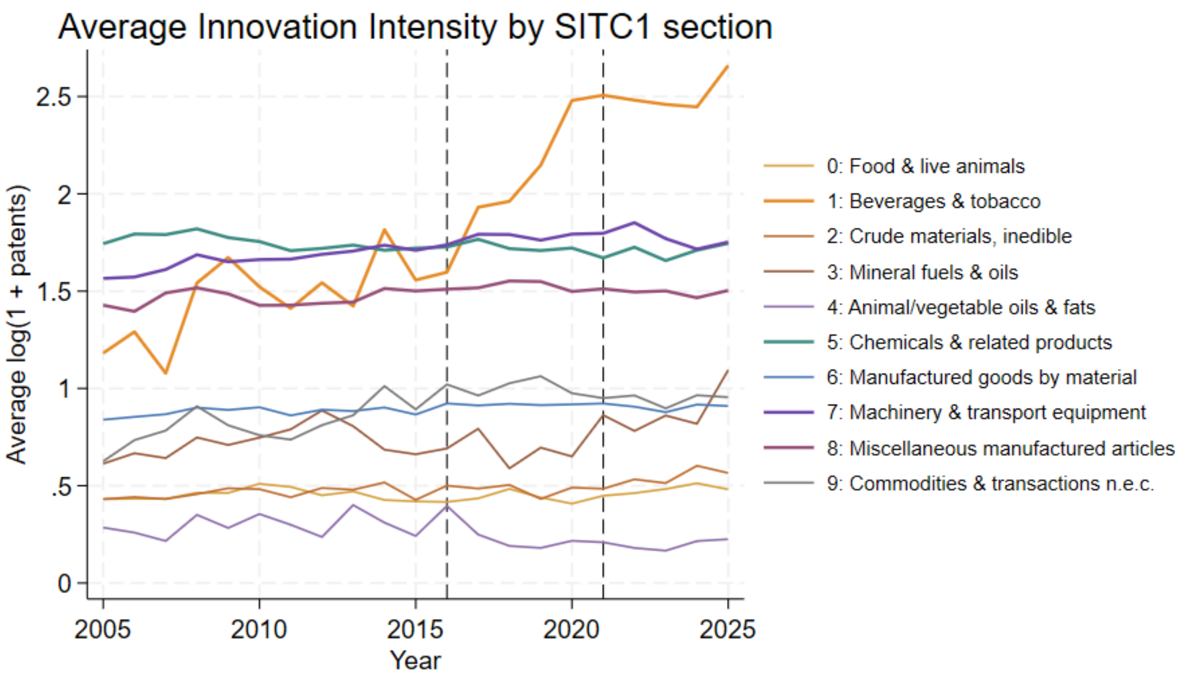

Figure 12: Evolution of patent intensity by SITC1 sector

3.4 Merging Brexit Exposure with Patents

This section finally merges the patent panel by SITC 4-digit sectors with the exposure to Brexit cross-section. The merge resulted into a 97% matching rate. The final dataset is made of 19,908 observations, namely 948 SITC 4-digit sectors over 21 years. The unmatched observations are mostly primary goods, raw materials and basic agricultural goods [9]. This is not a threat to the validity of this study, as these sectors are likely to have little impact on PCT patent mass. Lastly, note that baseline regressions will drop SITC1 sectors with insufficient observations for credible inference (namely, SITC 1, 3, 4, 9 - coherently from insights from **Table 5**), whenever identification is carried out within SITC1 sectors.

Having obtained the final dataset, some first correlational plots can be useful to get some intuition on the results that could come. For instance, this section tries to visualise the relationship between sectoral exposure to Brexit and patent intensity. A difficulty in doing this stems from the fact that exposure to Brexit is continuous variable. This section therefore distinguishes between above-median exposure (in blue, throughout) and below-median exposure (in red, throughout). Note that the median exposure to Brexit is 0.58 in the sample.

[9] A complete list of unmatched (hence, excluded) sectors can be found in the Appendix.

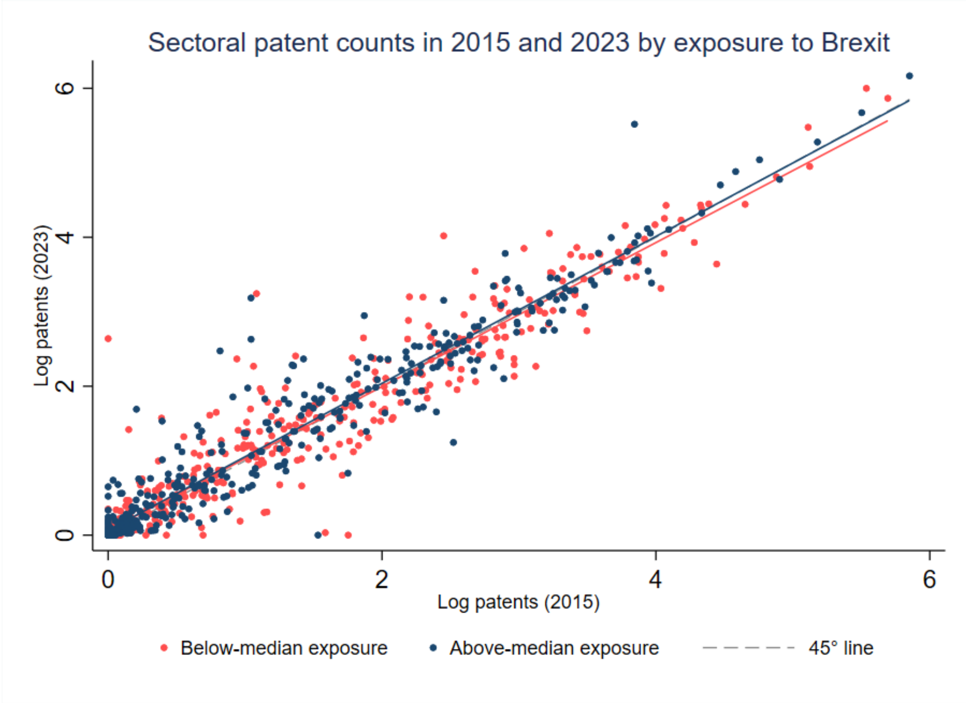

Figure 13: Scatterplot of SITC4 Sectors by Above- and Below-Median Exposure to Brexit

Figure 13 plots compares the patenting behaviour of sectors above- and below-median exposure to Brexit [10]. The x-axis reflects patent intensity in 2015 (pre-referendum), while the y-axis shows the same measure but for 2023 (after the TCA). Observations along the 45° line reflect sectors that display an unchanged patent intensity in 2015 and 2023. The blue dots and the respective linear fit refers to sectors that have above-median exposure to Brexit, while the red dots and its linear fit reflect the below-median counterparts. The first thing that can be noticed is that there is an extremely **strong persistence** in sectoral innovation intensity, as dots are highly clustered around the 45° line. This is coherent with previous insights on weak time-series variation, and the dominance of the cross-sectional one. A second consideration that arises from the graph is that the two clouds overlap heavily, and the two fitted lines are pretty close. If anything, the above-median exposure sectors seem to have maintained a stable innovation activity, while the below-median sectors have marginally diminished it. However, it is crucial to remember that this plot is still dominated by between-sector differences, which will be absorbed. The DiD effect will probably be of second order relative to sector fixed effects. Overall, however, there does not seem to be strong evidence of a differential break.

[10] Figure 24 in the Appendix reports the same exercise, performed by top and bottom exposure quartiles.

Figure 14 plots the time series of the mean patenting activity by above- and below-median exposure [11]. Interestingly, it is evident that the **most innovative sectors are those that were exposed less to Brexit**. Instead, the less innovative sectors were the most exposed to Brexit. The level gap seems to be persistent and slightly widening over time. Now, assuming that innovative sectors are economically stronger, this could signal that Brexit hit the most vulnerable sectors, eventually worsening inequality in sectoral economic outcomes. Another useful insight from this figure concerns a first assessment of the plausibility of **parallel trends**[12]. In the time series, looking at the pre-period 2005-2015, both groups show similar fluctuations and an overall growth pattern. Note, however, that continuous DiD does not compare a treatment and control group, but it compares across different doses of exposure. This implies that this plot is only indicative of parallel trends. A third contribution of this graph regards the **response to Brexit**. The solid vertical lines indicate the 2016 referendum and the 2021 TCA. Accounting for a lag in patent publication responses, the dashed vertical lines indicate a one- and two-year lag compared to the two Brexit events. It can be seen that after 2017, sectors that are relatively more exposed to Brexit see a decline in patenting, while sectors that are less exposed to Brexit maintain their positive innovation trajectory. While this is correlational, it is suggestive of a negative effect of Brexit on patenting activity. The evidence is weaker for the TCA, where from 2022 both sectoral groups see a decline in patenting during the COVID period.

[11] Figure 22 in the Appendix does the same exercise, comparing top and bottom deciles and quartiles.

[12] Figure 23 in the Appendix breaks down the plot by SITC1.

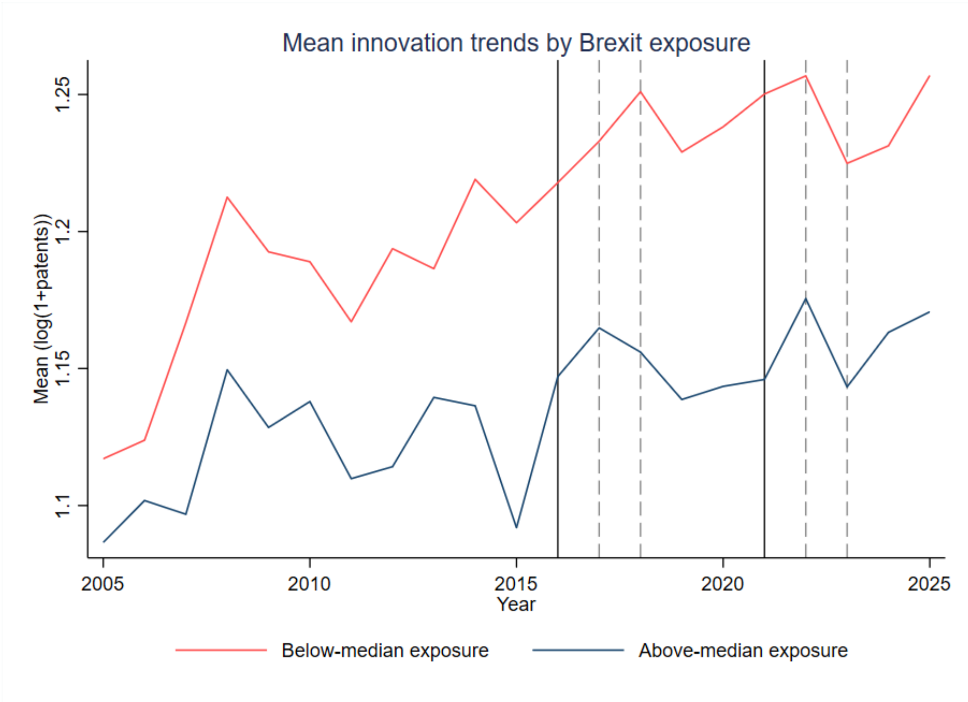

Figure 14: Innovation Trend by Above- and Below-Median Brexit Exposure

4 Methodology

The **identification strategy** used in this research is non-trivial and requires some attention, as it addresses several possible designs: difference-in-differences with continuous treatment, dose-response models and event studies. This section discusses the relevant econometric theory and explains the specifications that will be presented in the results section.

The goal of this research is to identify the causal effect of exposure to Brexit on innovation. Since exposure to Brexit is continuous and bounded between 0 and 1, the identification framework followed in this research is that of **Difference-in-Differences with continuous treatment**. The literature explains that this design is fundamentally different from the conventional binary treatment case. Canonical DiD designs compare outcomes of treated and untreated groups, before and after treatment, to identify average treatment effects under parallel trends assumptions. In the continuous Brexit exposure case, the treatment does not "turn on and off", but it rather has a **dose** d or treatment intensity that differs by sector. **Intuitively, identification relies on within-sector variation in products' exposure to Brexit, which allows to compare the innovation pace of relatively more exposed products to their minimally exposed counterparts.** An advantage of this configuration is that it allows to study Brexit in a more nuanced manner, allowing to ask comparative questions about the effects of being more or less exposed to Brexit, rather than the effects of the "occurrence" of Brexit, which is now a fact. The disadvantage of this continuous treatment variable is that identification involves several complications.

Recently, the literature has scrutinised such designs from an econometric perspective. **Callaway et al. [2024]** provide a state of the art summary of DiD with continuous treatment. They discuss the identification of treatment effects under parallel trends assumptions, show that TWFE estimators fail to have the desired causal interpretation, and propose alternative nonparametric estimators that tackle these issues.

4.1 Causal Parameters of Interest

This section explains the causal parameters of interest. The premise is that having a continuous treatment intensity implies that multiple counterfactual comparisons can be carried out. The authors explain several important distinctions that arise with continuous treatments, that are econometrically immaterial in conventional binary treatment designs. Following Callaway et al. [2024], we define three causal parameters to estimate. Note immediately that a crucial distinction is between causal parameters that relate to a specific dose *d* (the "building blocks"), and the summary causal parameters that, using a certain weighting scheme, try to extract a general causal response to the phenomenon of interest. Denote $\Delta Y = Y_{t=2} - Y_{t=1}$.

- **Level Treatment Effect ($LTE(d)$ or $ATT(d \mid d)$):** The level treatment effect at exposure level *d* is the average change in patenting for sectors with exposure level d, relative to the zero exposure counterfactual. This measures the effect of moving from no treatment to a particular level *d* of treatment. This can be computed $\forall d$. Note that this is a treatment effect that is "local" to units that experienced dose d.

$$LTE(d) = ATT(d \mid d) = \mathbb{E}[\Delta Y \mid D = d] - \mathbb{E}[\Delta Y \mid D = 0]$$

- **Average Treatment Effect on the Treated (ATT_0)**: To aggregate the parameters estimated $\forall d$ in order to get a summary measure, we average them using the distribution of the dose among treated units. By comparing all treated units (whetever the $d > 0$) with the zero counterfactual, the average Brexit effect across exposed sectors is obtained.

$$ATT_0 = \mathbb{E}[ATT(D \mid D) \mid D > 0]$$

- **Average Causal Response ($ACR(d)$):** Comparing adjacent dose groups, we obtain the marginal effect of a slight increase in exposure. This reflects how sensitive innovation is to incremental increases in exposure to Brexit. Note that this holds under strong parallel trends (SPT) as discussed in the next section.

$$ACR(d) = \frac{\partial}{\partial d}\mathbb{E}[\Delta Y \mid D = d]$$

Crucially, LTE and ATT capture **level effects**, by comparing outcomes for dose *d* to zero-exposure counterfactuals. Instead, ACR captures **marginal effects** of an increase in exposure.

4.2 Identification Assumptions

Identification with continuous treatment relies on a number of assumptions.
The first three assumptions are standard [Callaway et al., 2024]: (1) **Random Sampling**: observe two periods of independent and identically distributed panel data (a pre-outcome, a post-outcome, and a treatment dose); (2) **Continuous or multi-valued discrete treatment**: some observations are untreated, some are treated with a continuous or discrete treatment; (3) **No anticipation of future treatments**: we observe untreated potential outcomes in the pre-period, and potential outcomes for dose d experienced in the post-period.

The central assumption for DiD identification is given by parallel trends. Also in this case, the presence of a continuous treatment augments the number of possible comparisons, and with it, the number of identifying assumptions. In the following, the assumptions of Parallel Trends (PT) and Strong Parallel Trends (SPT) are defined for the present context.

- **Parallel Trends (PT)**: in the absence of Brexit, sectors with different exposure levels would have experienced parallel trends in their patenting activity; this ensures that the untreated group (zero exposure) is a valid counterfactual for any treated observation at intensity *d*, $\forall d$. In other words, the untreated potential outcomes of all sectors would have evolved similarly in terms of innovation trend. Under PT, each $ATT(d \mid d)$ is identified.

- **Strong Parallel Trends (SPT)**: the observed outcome changes for every dose group reflect what would have happened to all other groups, had they received that dose. For instance, a sector with exposure $D = d_1$ would have followed the patenting trend of $D = d_2$, had it been exposed to $D = d_2$. Similarly, the path of outcomes for low-dose units represents what the higher dose units would have experienced, had they received the lower dose. Essentially, SPT rules out selection-on-gains across exposure groups, ensuring that differences across dose groups reflect causal responses rather than heterogeneity bias. This allows to identify ACR.

Note that PT is an assumption that makes comparisons with the zero-dose counterfactual. Instead, SPT is an assumption that compares all possible doses d.

To assess the **plausibility of PT and SPT**, the authors suggest to examine event study-type aggregate effects in pre-treatment periods. Finally, note that there exists a version of SPT that hold after conditioning of covariates. In the thesis' context, **Conditional SPT** may hold after conditioning on SITC1 sectors.

4.3 Estimation

Having discussed the causal parameters of interest and identification assumptions, this section starts introducing estimation. The following points remain true across all specifications.

- **Brexit events.** This thesis tries to identify the causal effect of Brexit on innovation. In particular, it distinguishes between the effects of two main events: the 2016 Brexit Referendum and the 2021 TCA.
- **Patent lags.** As discussed in the descriptive section about patenting, the variable chosen is given by patent *publications*. Since these come with a lag compared to the actual innovation, all specifications are run considering alternative lags.
- **Individual and time dimension.** All regressions exploit the finest granularity of the sectoral classification. Thus, they are run at the SITC 4-digit level ("S4" in regressions) over years ("t" in regressions). In some specifications, the year fixed effects are replaced by SITC1xyear fixed effects. The rationale is to restrict identification to comparisons between SITC4 products that belong to the same SITC1 sector.

4.3.1 Two-Way-Fixed-Effects with Continuous Treatment

We can now turn to estimation. For several years, the most natural approach to deal with continuous treatment DiD has been to run conventional **two-way-fixed-effects (TWFE)** regressions. For instance, TWFE regresses the log-patent count on sector and year fixed effects, and the interaction of Brexit exposure with the relevant post dummy:

$$\text{lpatents}_{S4,t} = \alpha_{S4} + \gamma_t + \beta^{twfe} \left(\text{BrExposure}_{\text{S4}} \times \text{PostYYYY}_t\right) + \varepsilon_{S4,t} \tag{1}$$

Where: $\text{lpatents}_{S4,t} = \log(1 + \text{patents})$ for SITC4 product $S4$ in year t; α_{S4} = SITC4 fixed effect; γ_t = year t fixed effect; $BrExposure_{S4}$ = Brexit exposure (0 to 1) for SITC4 $S4$; $PostYYYY_t = 1$ for $t \geq YYYY$ and 0 otherwise; β^{twfe} = marginal effect of exposure in the post-YYYY period; $\varepsilon_{S4,t}$ = error term with standard errors clustered at the SITC4 level.

TWFE will be treated as a standard benchmark. However, it has a number of limitations that lead it to not have a causal interpretation. Following **Callaway et al. [2024]**:

1. **TWFE aggregates dose-specific level effects using non-transparent weights.** Summary measures have the goal of aggregating the dose-specific estimands ("building blocks") into a unique measure of the overall effect of a phenomenon. In fact, β^{twfe} can be expressed as a weighted integral of average level treatment effect parameters, yet weights integrate to zero. The latter implies the existence of negative weights. In particular, TWFE assigns negative weights to below-average dose units, and positive weights to above-average dose units. This means that TWFE essentially runs a weighted binary DiD, using as treatment group the "high dose" units and as control group the "low dose units" - where weights are the distance of observations from the mean dose.

2. **TWFE suffers of selection bias under PT.** Decomposing β^{twfe} by causal response, the weights integrate correctly to one and remain non-negative. However, the ACR is only identified under SPT assumption. In fact, under standard PT, β^{twfe} identifies the ACR plus treatment heterogeneity across doses. This is intuitive: comparisons across different dose groups include a causal response, but possibly also the *different treatment effects of the same dose for different dose groups.* For example, if units with higher doses had larger treatment effects at each dose, the coefficient will be biased upwards.

3. **TWFE imposes global linearity**. TWFE is biased if the true data generating process yields different treatment effects depending on the dose. In other terms, the coefficient is the slope of the best linear approximation to the dose-response function.

Thus, the coefficient of this regression will be treated with caution. In principle, under SPT and linearity, β^{twfe} can be interpreted as a linear approximation of the average causal response, i.e., how much patenting causally increases on average where Brexit exposure increases by one unit. In this case, since exposure runs from 0 to 1, it is the effect of moving from no exposure to full exposure. Picking a representative dose, say the sample median of 0.58, the product $0.58\beta^{twfe}$ is interpreted as the average treatment effect for the median exposure dose. However, the discussion above suggests that β^{twfe} does *not* have a clear causal interpretation due to weighting, bias, and linearity. The fundamental problem therefore lies in the aggregation of dose-specific estimands into more generic summary results. In this view, the best approach is to choose aggregation schemes explicitly or to rely on dose-specific estimates. This leads Callaway et al. [2024] to suggest the use of non-parametric techniques.

4.3.2 Nonparametric estimators

To address the limitations of TWFE outlined above, this section follows the path suggested by Callaway et al. [2024]. The main idea is to employ non-parametric estimators to do three distinct estimation exercises.

Multi-valued Discrete Treatment

The first non-parametric approach requires the treatment to be **multi-valued discrete**. This means the continuous Brexit exposure indicator should be discretised into bins, for instance based on 10 deciles. Then, one can regress outcome changes on the (saturated) dose indicators, choosing as omitted category the untreated units. The regression equation would read:

$$\Delta \text{lpatents}_{S4} = \beta_0 + \sum_{j=1}^{J} \mathbf{1}\{\text{BrExposure}_{S4} = d_j\}\beta_j + \varepsilon_{S4} \tag{2}$$

Equation (2) can then be estimated through OLS and the coefficients of interest are the set of $\hat{\beta}_j$ for each j. Under PT, these coefficients identify the $ATT(d \mid d)$, namely the average treatment effect for sectors in each of the J exposure bins. Taking the difference between consecutive coefficients and assuming SPT, $\hat{\beta}_j - \hat{\beta}_{j-1}$ identifies the $ACR(d_j)$.

Understanding this equation is very simple, by reasoning through estimating the previous TWFE equation in differences, and considering that exposure is binned into J categories. For instance, let $1BrExposure_s = d_j$ be the indicator that sector $S4$ lies in exposure bin d_j, $j = 1, \ldots, J$, and let $Post_t$ be an indicator equal to 1 in the post period and 0 in the pre period.

In the pre-period, t=1 and Post=0, and the TWFE equation reads:

$$\text{lpatents}_{S4,1} = \alpha_{S4} + \gamma_1 + \sum_{j=1}^{J} \beta_j^{\text{twfe}} (\mathbf{1}\{\text{BrExposure}_{S4} = d_j\} \times 0) + \varepsilon_{S4,1} \quad = \alpha_{S4} + \gamma_1 + \varepsilon_{S4,1}$$

In the post-period, t=2 and Post=1, and the TWFE equation reads:

$$\begin{aligned} \text{lpatents}_{S4,2} &= \alpha_{S4} + \gamma_2 + \sum_{j=1}^{J} \beta_j^{\text{twfe}} \big(\mathbf{1}\{\text{BrExposure}_{S4} = d_j\} \times 1\big) + \varepsilon_{S4,2} \\ &= \alpha_{S4} + \gamma_2 + \sum_{j=1}^{J} \beta_j^{\text{twfe}} \mathbf{1}\{\text{BrExposure}_{S4} = d_j\} + \varepsilon_{S4,2} \end{aligned}$$

Taking the difference between the post- and pre-period equation, the sector fixed effects simplify, and re-arranging yields:

$$\Delta\text{lpatents}_{S4} = (\gamma_2 - \gamma_1) + \sum_{j=1}^{J} \beta_j^{\text{twfe}} \mathbf{1}\{\text{BrExposure}_{S4} = d_j\} + (\varepsilon_{S4,2} - \varepsilon_{S4,1})$$

Finally, defining $\beta_0 \equiv (\gamma_2 - \gamma_1)$ and $\varepsilon_{S4} \equiv (\varepsilon_{S4,2} - \varepsilon_{S4,1})$, it is easy to recognise equation (2). Notice that the difference of time fixed effects collapse to the constant, in the regression. Finally, notice the role of using untreated units as omitted category. Omitting one category is necessary to avoid perfect multicollinearity. The specific omission of the zero-dose groups allows to interpret the coefficients as as level effects of does d with respect to the no exposure counterfactual. That is, each $\hat{\beta}_j$ is the difference in outcomes between bin j and zero-exposure sectors - which is in fact the definition of an $ATT(d_j)$. The constant will identify the average change in outcomes for zero-exposure sectors.

Continuous Treatment and Dose-Response Functions

A second non-parametric exercise is a continuous counterpart of the technique described above. Conceptually, we have change in sectoral patenting $\Delta lpatents_{S4}$ and sectoral exposure $BrExposure_{S4}$. We wish to estimate $\mathbb{E}[\Delta lpatents \mid BrExposure = d], \forall d$, to understand, say, the effect of Brexit for sectors with 30, 45, or 80% exposure. Altogether, this information forms a **dose-response function**, that maps each level of treatment intensity into an expected outcome change. As introduced in previous sections, TWFE assumes this function to be linear. This means that TWFE cannot capture a situation where effects vary by exposure. For instance, it may be that effects are null for exposures 0-20%, negative for 20-60%, and then null again. A fitted straight line is clearly incapable to capture this. This motivates the desire to use a tool that allows for a more flexible description of the relationship between treatment intensity and patenting activity.

In practice, the authors suggest to run the following regression:

$$\Delta \text{lpatents}_{S4} = \sum_{k=1}^{K} \psi_{Kk}(\text{BrExposure})\, \beta_{Kk} + \varepsilon_{S4}. \tag{3}$$

where $\psi_{Kk}(\cdot)$ are K known flexible transformations of the dose, and β_{Kk} are the k parameters that form vector β_K. Equation (3) can be estimated by OLS. Under PT, these estimate the $ATT(D \mid D)$; under SPT, they estimate $ATE(d)$, whose derivative is $ACR(d)$.

A standard choice for the $\psi_{Kk}(\cdot)$ is a **cubic B-spline** with K=4. A **spline** is a flexible curve built from small pieces of polynomials. The idea is that instead of fitting one straight line across the entire dose range, it is better to divide the dose range in K segments (say, K=4), fit a curve in each segment, and smoothly join the K curves. Each of the K functions, one per segment, is called a **basis function**. The junction points that divide the dose range into segments are called **knots**, and they are usually placed at quartiles. A spline is **cubic** if the curve fitted inside each segment is a cubic polynomial of the exposure variable. Finally, a cubic **B-spline** is simply a computationally convenient way of constructing this piece-wise cubic function. Overall, this method allows to estimate the dose-response curve without forcing it to be linear, allowing the data to determine the shape of the relationship between exposure and outcome.

Summary Measures of Treatment Effects

So far, the two proposed nonparametric estimators were tied to the specific dose d. However, one may be interested in deriving more generic estimators, to answer global questions on the overall effects of the Brexit Vote or the TCA, virtually collapsing the exposure dimension. As discussed, this is exactly where TWFE fails, as it aggregates the dose-specific estimands either using non-transparent weighting schemes, or generating a biased estimate. The author's suggestion is to run a simple binarised DiD regression, by comparing all treated units

($BrExposure > 0$, regardless of the specific dose) to untreated units ($BrExposure = 0$). The estimating equation reads:

$$\Delta \text{lpatents}_{S4} = \beta_0^{bin} + \mathbf{1}\{\text{BrExposure}_{S4} > 0\}\ \beta^{bin} + \varepsilon_{S4}. \tag{4}$$

In the equation above, the binarised DiD estimator β^{bin} identifies the ATT_0 under PT, and ATE_0 under SPT. Thus, this coefficient can be interpreted as the average change in patents for any exposed sector, compared to the zero-exposure counterfactual. This estimator is non-parametric in the dose dimension, it is a simple difference in means with no assumption on the shape of responses by exposure.

Testing Pre-Trends

In binary treatment DiD, one should check if the treated group trends differently from the control group, prior to the treatment. The continuous DiD counterpart, following Callaway et al. [2024], is to inspect whether the slope of Brexit exposure predicts innovation differences before treatment. Intuitively, testing parallel trends directly for all doses is complicated, so one can estimate the following event-study regression and look at the coefficients before treatment; if the latter are approximately zero, there is no differential pre-trend by exposure. Note that the clean pre-treatment period is limited to 2005-2015, namely relative to the pre-referendum period. Doing a similar exercise for the TCA would be confounded by the referendum effects and the transition period. Thus, defining event time relative to the referendum year 2016: e = t - 2016. Year 2015 (e = -1) is omitted, and taken as baseline.

$$\text{lpatents}_{S4} = \alpha_{S4} + \gamma_t + \sum_{e \neq -1} \beta_e \Big(\text{BrExposure}_{S4} \times \mathbf{1}\{t - 2016 = e\} \Big) + \varepsilon_{S4,t}. \tag{5}$$

In equation (5), β_e captures how strongly innovation differs with exposure, relative to 2015. Parallel trends corresponds to $\beta_e = 0, \forall e < 0$. Note that (5) can also be estimated with SITC1xyear fixed effects, to allow for SITC1-specific time shocks, and restraining comparisons to within-SITC1 sectors. This concludes the methodology section. The next section will show the regression results obtained by estimating equations 1, 2, 3 and 4.

5 Results

This Chapter reports the results of this thesis, by estimating the five equation types described in the methodology. **Section 5.1** carries out pre-trend tests to build confidence towards the plausibility of parallel trends. **Section 5.2.** identifies the aggregate effect of the 2016 Brexit Referendum. **Section 5.3.** reports the aggregate effect of the 2021 TCA. **Section 5.4.** proposes some estimations of sectoral heterogeneity. Finally, **Section 5.5.** discusses the limitations of this study and areas for future research.

5.1 Testing the Plausibility of Parallel Trends

To assess the plausibility of parallel trends, equation (5) was estimated setting 2015 as the baseline year. A joint test of the pre-2016 exposure-year coefficients fails to reject the null hypothesis of coefficients being equal to zero ($F(10, 947) = 1.09, p = 0.366$). This suggests that there is no evidence of systematic differential pre-trends across levels of Brexit exposure. **Figure 17** confirms that pre-treatment coefficients are small and statistically null. The figure and F-test refer to the specification with SITC1xYear fixed effects, but results are robust to the baseline version that only includes SITC4 and time fixed effects. Note that 2013 and 2014 coefficients in the SITCxYear specification are individually weakly significant, but the joint test makes the collective assessment of no *systematic* differences from zero. In turn, this allows us to build confidence towards the identification strategy of using differently exposed products within the same SITC1 category as counterfactuals.

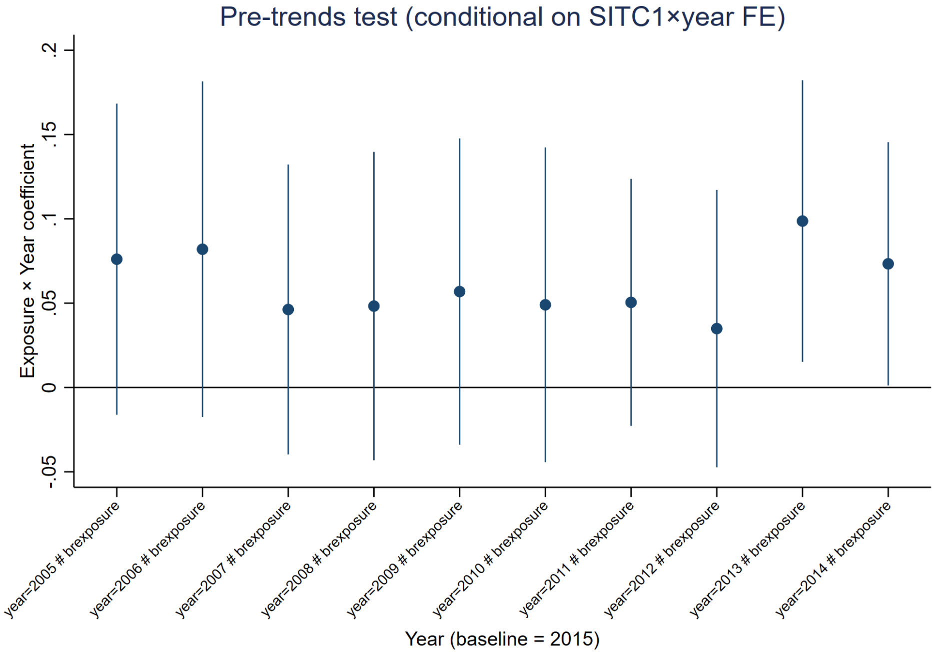

Figure 15: Pre-Trends Test for Plausibility of Parallel Trends

5.2 The Aggregate Effect of the Brexit Referendum on Innovation

Having inspected the credibility of identifying assumptions, this section leverages sectoral variation in exposure to Brexit to estimate the effects of the referendum on patenting activity.

Tables 8 and 9 report the results from the estimation of equations (1) and (4), namely TWFE and Binarised DiD. It is immediately visible that coefficients are negative across all specifications and insignificant. The TWFE coefficients identify the Average Causal Response under SPT and linearity (which is likely to hold in specifications with SITC1 fixed effects), while under PT it is a weighted average of heterogeneous Level Treatment Effects. The preferred specification is reported in column (3), and includes a two-year patenting lag and SITC1 fixed effects. The coefficient suggests that **one-unit increase in Brexit exposure (namely, moving from no exposure to full exposure), is associated with a 2.7 percentage point lower change in log patenting after the referendum.** However, the coefficient is small and statistically insignificant.

Moving to the binarised DiD that compares all the treated to the untreated (empirically defined as SITC4 sectors whose pre-referendum exposure to Brexit is < 0.05), the picture does not dramatically change. The coefficient, under PT, indentifies ATT_0, the average treatment effect across exposed sectors, against the zero-exposure counterfactual. The preferred specification in column (2) suggests that, **on average, exposed sectors experienced 5 percentage point lower patent growth relative to minimally exposed sectors**, two years after the referendum. The coefficient becomes smaller for effects registered three years after the referendum. Again, no estimates are statistically different from zero.

Jointly, both summary estimators point to the conclusion that, at least in the aggregate, there is **no statistically significant effect** of Brexit exposure on patenting. If any, however, these effects are **negative and relatively small** in magnitude.

Table 8: **TWFE – 2016 Referendum effects**

	(1) Post 2017	(2) Post 2017	(3) Post 2018	(4) Post 2018
Exposure $\times$ Post	-0.01724	-0.050	-0.02651	-0.055
	(0.02575)	(0.034)	(0.02696)	(0.035)
SITC4 FE	Yes	Yes	Yes	Yes
Year FE	No	Yes	No	Yes
SITC1$\times$Year FE	Yes	No	Yes	No
Observations	18,774	19,908	18,774	19,908
Clusters (SITC4)	894	948	894	948
Within R^2	0.0001	0.0009	0.0002	0.0011

Notes: Columns (1) and (3) report restricted-sample estimates (SITC1 $\neq$ 1,3,4,9). Dependent variable is $\log(1 + \text{patents})$ at the SITC4 level. Robust standard errors in parentheses, clustered at the SITC4 level in full-sample regressions.

Table 9: **Binarised DiD — 2016 Referendum effects**

	(1) 2018	(2) 2018	(3) 2019	(4) 2019
Treated (D $>$ 0.05)	-0.07267	-0.05440	-0.00477	-0.00909
	(0.05674)	(0.05454)	(0.03891)	(0.03055)
SITC1 controls	No	Yes	No	Yes
Years used	2016,2018	2016,2018	2016,2019	2016,2019
Observations	948	894	948	894
Clusters (SITC4)	948	894	948	894
R^2	0.0023	0.0162	0.0000	0.0172

Notes: Dependent variable is $\Delta Y = \log(1 + \text{patents})_t - \log(1 + \text{patents})_{2016}$ at the SITC4 level. Each column reports the coefficient on the dummy `treated` (brexposure $>$ 0.05). Standard errors clustered at the SITC4 level in parentheses. Columns with SITC1 controls correspond to the restricted sample (SITC1 = 1,3,4,9 dropped).

Nonparametric Estimators by Exposure to Brexit

Nonparametric estimates give additional insight on where these average effects could be coming from. **Table 10** reports the coefficients associated with Brexit exposure deciles. Under parallel trends, this coefficient represents β_j, for $j = 1, \ldots, 10$, namely $ATT(d_j \mid d_j)$, the level treatment effect for the each of the exposure deciles relative to the (virtually) zero exposure control group. Interestingly, the highest-decile coefficients become negative and significant at the 5% level, providing an estimate of $ATT(d_{10} \mid d_{10})$. After the referendum, **highly exposed sectors exhibited approximately 11 percentage points lower patent growth relative to their minimally exposed counterparts**. The effect refers to the specification that uses a two-year lag and includes SITC1 fixed effects.

The left panel of **Figure 16** reports a predictive margins plot by exposure deciles, corresponding to the discussed specification [13]. The right panel reports the continuous treatment counterpart, namely the estimation of the spline dose-response function. Importantly, notice that while the coefficients reported in **Table 10** correspond to level treatment effects relative to the control group, the margins plotted in **Figure 16** represent the predicted mean change in patenting for each exposure bin, without differencing relative to the control. That is, they trace the conditional expectation $E[\Delta Y \mid D = d]$ across exposure bins and show the control mean. The spline traces the conditional expectation of patent growth as a smooth function of exposure intensity. It immediately emerges that there is no strictly monotonic dose-response pattern across most of the exposure distribution. **Significant negative effects appear concentrated in the extreme upper tail of exposure**. **Middle-exposure deciles show positive and statistically significant growth in patenting**, despite this growth still being weaker than that of the zero-exposure control. The spline dose-response curve confirms and smoothens this pattern.

This pattern can be interpreted in light of the **uncertainty and expectations channel**. Following the referendum, sectors that relied heavily on EU trade faced uncertainty regarding future market access and regulatory conditions. In response to this, they decreased or delayed R&D forward-looking investments, resulting in the observed decline in patenting activity.

[13] The regression table of predictive margins for the Referendum can be found in the Appendix, Table 18.

However, this mechanism seems to have operated only at very high-exposure margin, at least in the aggregate. By contrast, sectors with intermediate levels of EU integration may have exhibited greater resilience. A moderate degree of integration may have allowed these sectors to adjust along existing commercial and collaborative networks rather than undertaking costly and uncertain restructuring. Reallocating activity within established supply chains and innovation partnerships is plausibly less disruptive than forming entirely new trade and research relationships. As a result, these sectors may have been able to maintain continuity in their patenting activity despite the referendum shock.

Table 10: **Discretised Treatment by Decile — 2016 Referendum Effects**

	(1) 2018	(2) 2018	(3) 2019	(4) 2019
Decile 1	-0.0400	-0.0728	-0.0242	-0.0045
	(0.0624)	(0.0660)	(0.0421)	(0.0491)
Decile 2	-0.0291	-0.0629	0.0187	0.0103
	(0.0624)	(0.0619)	(0.0410)	(0.0471)
Decile 3	-0.1018*	-0.0950	-0.0323	-0.0133
	(0.0594)	(0.0608)	(0.0402)	(0.0469)
Decile 4	-0.0020	-0.0261	0.0138	0.0116
	(0.0600)	(0.0615)	(0.0389)	(0.0456)
Decile 5	-0.0321	-0.0526	0.0219	0.0225
	(0.0607)	(0.0616)	(0.0394)	(0.0459)
Decile 6	-0.0580	-0.0948	0.0033	-0.0096
	(0.0613)	(0.0626)	(0.0386)	(0.0457)
Decile 7	-0.0338	-0.0574	0.0046	-0.0107
	(0.0597)	(0.0625)	(0.0398)	(0.0465)
Decile 8	-0.0445	-0.0706	-0.0140	-0.0141
	(0.0601)	(0.0611)	(0.0428)	(0.0475)
Decile 9	-0.0658	-0.0699	-0.0097	0.0006
	(0.0589)	(0.0608)	(0.0406)	(0.0471)
Decile 10	-0.1165**	-0.1243**	-0.0467	-0.0405
	(0.0584)	(0.0618)	(0.0367)	(0.0466)
SITC1 FEs	Yes	No	Yes	No
Years used	2016, 2018	2016, 2018	2016, 2019	2016, 2019
Observations	894	948	894	948
R^2	0.0336	0.0124	0.0242	0.0040

Notes: Dependent variable is $\Delta Y = \log(1 + \text{patents})_{post} - \log(1 + \text{patents})_{2016}$ at the SITC4 level. Each coefficient reports the difference in mean ΔY between exposure decile j (among treated observations) and the control group (bin 0), where treated is defined as BrExposure > 0.05; Columns with SITC1 controls correspond to the restricted sample (SITC1 = 1,3,4,9 dropped); Robust standard errors in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

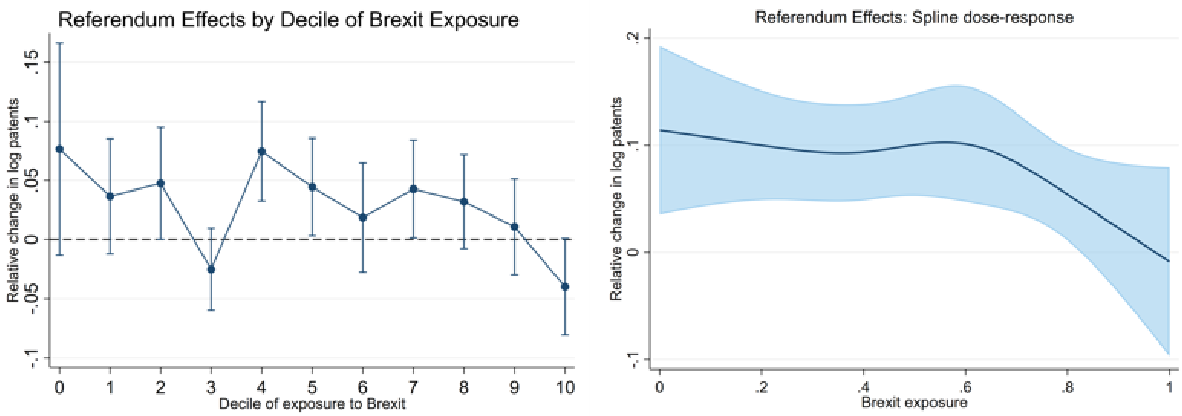

Figure 16: Nonparametric Effects of the 2016 Referendum

5.3 The Aggregate Effect of the TCA on Innovation

Moving to the TCA, the effects are more ambiguous and might have to be taken with a grain of salt in light of the COVID occurring in the same period. **Tables 11 and 12** report estimates of TWFE and binarised DiD. All coefficients are small, mostly negative, and insignificant. Starting from TWFEs, the preferred specifications that include SITC1 fixed effects report that there is no statistically detectable linear relationship between Brexit exposure and patent growth following the implementation of the TCA. Namely, a virtually null or very weakly negative average causal response. **Binarised DiD** mostly confirms this finding, suggesting absent average level effects.

Table 11: **TWFE – 2021 TCA effects**

	(1) Post 2022	(2) Post 2023	(3) Post 2022	(4) Post 2023
Exposure × Post	-0.00139	0.00061	-0.023	-0.019
	(0.02873)	(0.02962)	(0.035)	(0.035)
Year FE	No	No	Yes	Yes
SITC1×Year FE	Yes	Yes	No	No
Observations	18,774	18,774	19,908	19,908
Clusters (SITC4)	894	894	948	948
Within R^2	0.0000	0.0000	0.0001	0.0000

Notes: Columns (1)–(2) report the restricted-sample estimates (SITC1 $\neq$ 1,3,4,9); Dependent variable is $\log(1+\text{patents})$ at the SITC4 level; Robust SE in parentheses, clustered SE at the SITC4 level in the full-sample regressions.

Table 12: **Binarised DiD — 2021 TCA effects**

	(1) 2024	(2) 2024	(3) 2023	(4) 2023
Treated (D $>$ 0.05)	-0.05519	-0.00461	-0.01308	0.03108
	(0.05238)	(0.04955)	(0.03425)	(0.03429)
SITC1 controls	No	Yes	No	Yes
Years used	2021,2024	2021,2024	2021,2023	2021,2023
Observations	948	894	948	894
Clusters (SITC4)	948	894	948	894
R^2	0.0010	0.0481	0.0001	0.0118

Notes: dependent variable is $\Delta Y = \log(1+\text{patents})_t - \log(1+\text{patents})_{2021}$ at SITC4; each column reports the coefficient on the dummy `treated` (brexposure $>$ 0.05); SE clustered at SITC4 in parentheses; restricted columns drop SITC1 $=$ 1,3,4,9.

Nonparametric Estimators by Exposure to Brexit

Nonparametric estimations suggest that effects differ by the degree of exposure, although this affects the sign but not the statistical significance of the coefficients. **Table 13** shows the estimation results by decile of exposure to Brexit. Despite absence of statistical significance, an interesting pattern can be found in the signs of the coefficients across deciles, although they are quite sensitive to the choice of lag. The preferred specification is in column (3). Across specifications, a surprising pattern is that the highest-decile coefficients are positive. **Figure 17** aids intuition by providing the associated predictive margins plot, together with the spline dose-response. The graph suggests that the patenting growth of the tenth-decile is an outlier compared to the other degrees of exposure [14]. For instance, sectors with median exposures to Brexit are suffering small patenting declines, even compared to the null control mean. This pattern is confirmed by the spline dose-response function.

Given the sensitivity of the estimates to the assumed patent lag, as well as the confounding effects of the COVID-19 shock during the TCA period, the results should be interpreted with caution. Nonetheless -taking these effects at face value- it is striking that the nonparametric estimates for the **TCA display a pattern that is qualitatively opposite to that observed after the referendum**. In the referendum period, it was precisely the highest exposure decile that exhibited a relative decline in patenting activity. The subsequent rise in patenting for this same group after the TCA is therefore quite surprising, although it is not statistically significant. This result could be consistent with an a **delayed disclosure mechanism**. Namely, firms in highly exposed sectors may have postponed or slowed the formal filing of patent applications during the heightened uncertainty surrounding the referendum and throughout the pandemic. Innovation projects may have continued in the background, but their registration could have been deferred. **Once both policy uncertainty and macroeconomic disruptions receded, these cumulated innovation efforts may have materialised in a concentrated wave of patent applications**. Under this interpretation, the positive tail effect observed during the TCA period reflects the timing of innovation disclosure rather than a structural increase in innovation incentives. However, if present, this effect is not quantitatively substantial, as elicited by the statistical insignificance of coefficients.

[14] The regression table of predictive margins for the TCA can be found in the Appendix, Table 19.

Table 13: Discretised Treatment by Decile — 2021 TCA Effects

	(1) 2023	(2) 2023	(3) 2024	(4) 2024
Decile 1	0.0645	0.0151	-0.0378	-0.0864
	(0.0475)	(0.0473)	(0.0615)	(0.0635)
Decile 2	0.0249	-0.0242	0.0230	-0.0522
	(0.0443)	(0.0429)	(0.0581)	(0.0603)
Decile 3	0.0109	-0.0379	-0.0215	-0.0739
	(0.0442)	(0.0438)	(0.0597)	(0.0618)
Decile 4	-0.0051	-0.0286	-0.0061	-0.0629
	(0.0468)	(0.0459)	(0.0572)	(0.0597)
Decile 5	0.0074	-0.0392	-0.0334	-0.0975
	(0.0446)	(0.0425)	(0.0576)	(0.0589)
Decile 6	0.0004	-0.0412	-0.0441	-0.0887
	(0.0404)	(0.0391)	(0.0538)	(0.0565)
Decile 7	0.0544	0.0135	-0.0013	-0.0429
	(0.0450)	(0.0442)	(0.0560)	(0.0584)
Decile 8	0.0052	-0.0383	-0.0344	-0.0798
	(0.0434)	(0.0417)	(0.0561)	(0.0585)
Decile 9	0.0450	0.0276	-0.0070	-0.0184
	(0.0414)	(0.0396)	(0.0567)	(0.0587)
Decile 10	0.0649	0.0224	0.0849	0.0514
	(0.0448)	(0.0431)	(0.0605)	(0.0607)
SITC1 FEs	Yes	No	Yes	No
Years used	2021, 2023	2021, 2023	2021, 2024	2021, 2024
Observations	894	948	894	948
R^2	0.0206	0.0108	0.0623	0.0219

Notes: Dependent variable is $\Delta Y = \log(1 + \text{patents})_{post} - \log(1 + \text{patents})_{2021}$ at the SITC4 level. Each coefficient reports the difference in mean ΔY between exposure decile j (among treated observations) and the control group (bin 0), where treated is defined as BrExposure > 0.05; Columns with SITC1 controls correspond to the restricted sample (SITC1 = 1,3,4,9 dropped); Robust standard errors in parentheses; ${}^{*}p < 0.10$, ${}^{**}p < 0.05$, ${}^{***}p < 0.01$.

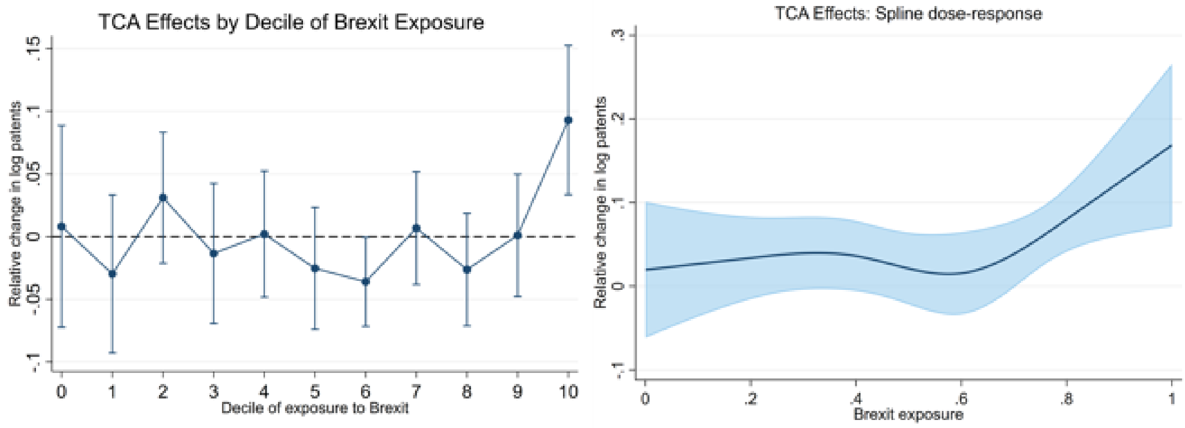

Figure 17: Nonparametric Effects of the 2021 TCA

5.4 The Sectoral Effects of Brexit on Innovation

The previous sections of this chapter have analysed the aggregate effects of Brexit. In the analysis, summary parameters from TWFE and Binarised DiD proved to consistently be small and insignificant. This is not extremely surprising, for two reasons. First, as shown in the descriptive analysis, most of the variation in patenting activity is cross-sectional rather than time-series, which was suggestive that -if any- effects would be very modest in magnitude. Second, as explained in **Section 1.1**, a correct and more credible interpretation of Brexit effects is best obtained by operating at the **sector level**, rather than relying on pooled aggregates. Intuitively, this is due to the fact that the TCA ended up introducing non-tariff-barriers, that by construction are highly sector-specific. This section has two purposes. First, it shows that sectoral heterogeneity indeed exists. Second, it zooms into the evidence regarding the chemistry sector, which is the salient case of Brexit-induced regulatory divergence.

To show the relevance of sectoral heterogeneity, **Table 14** shows the result of a simple TWFE estimation mirroring equation (1), yet augmented with SITC1 interactions - only considering sectors with sufficient observations for inferential purposes. That is, the estimating equation now reads:

$$\begin{aligned}\mathrm{lpatents}_{S4,t} &= \alpha_{S4} + \gamma_t + \beta_0 \left(BrExposure_{S4} \times PostYYYY_t\right) \\ &\quad + \sum_{S1 \neq S1^{\text{base}}} \beta_{S1} \left(BrExposure_{S4} \times PostYYYY_t \times \mathbf{1}\{S1\}\right) + \varepsilon_{S4,t}\end{aligned}$$

Where: $\mathrm{lpatents}_{S4,t} = \log(1 + \text{patents})$ for SITC4 product $S4$ in year t, α_{S4} = SITC4 fixed effect, γ_t = year fixed effect, $BrExposure_{S4}$ = Brexit exposure (0–1) for SITC4 $S4$, $Post2022_t = 1$ for $t \geq 2022$ and 0 otherwise, $\mathbf{1}\{S1\}$ = indicator for SITC1 sector $S1$ (with one base sector omitted), β_0 = effect for the base SITC1 sector, β_{S1} = differential effect relative to the base sector, $\varepsilon_{S4,t}$ = error term with standard errors clustered at the SITC4 level.

Table 14: Sectoral marginal effects of Brexit exposure (by event)

	(1) Referendum (post-2017)	(2) TCA (post-2022)
Food (SITC 0, base)	-0.06088*	-0.00616
	(0.03409)	(0.03673)
Crude materials (SITC 2)	-0.00755	0.06841**
	(0.02944)	(0.03471)
Chemicals (SITC 5)	-0.12054***	-0.08664***
	(0.02811)	(0.03096)
Manufactured goods (SITC 6)	-0.03880	-0.04141
	(0.03187)	(0.03484)
Machinery (SITC 7)	0.09388**	0.06299
	(0.03667)	(0.04322)
Misc. manufacturing (SITC 8)	0.01468	-0.02770
	(0.04776)	(0.04913)
Observations	18,774	18,774
Clusters (SITC4)	894	894
Within R^2	0.0120	0.0043

Notes: Each cell reports the marginal effect of a one-unit increase in *brexposure* ($0 \to 1$) after the indicated event. Coefficients are obtained from sector-specific `lincom` estimates based on models with SITC4 and year fixed effects; standard errors clustered at the SITC4 level in parentheses. Sample restricted as in the main specification (SITC1 = 1, 3, 4, 9 dropped). Significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Several observations emerge from **Table 14**. First, exposure to Brexit features highly heterogeneous sectoral effects, coherently with the sector-specific nature of non-tariff-barriers. In line with the pooled results in the previous section, most coefficients are small in magnitude and statistically insignificant. Interestingly, effects also present opposite directions depending on the sector. Overall, **this heterogeneity exercise confirms the hypothesis that Brexit effects are likely to manifest through multiple, partially overlapping sectoral effects.**

When it comes to interpreting single sectoral coefficients, a first-best would be to be able to link each of the SITC1 sectors to contextual information on the NTBs that impacted them. However, being SITC1 a trade classification rather than an NTB classification, this exercise should be done by aggregating across products by the NTB type affecting them. Following this idea, categories such as "Manufactured goods" are too heterogeneous to be mapped credibly into a single regulatory mechanism. Because non-tariff barriers operate at a more granular product and regulatory level, attributing structural interpretations to each SITC1 coefficient would risk over-interpretation driven by aggregation.

An important exception to this argument is the **Chemicals sector** (SITC 5). This sector stands out for three reasons. First, recalling **Figure 12**, Chemistry is the SITC1 sector that displays the **highest patenting intensity** throughout the entire time series. Thus, analysing chemistry is a way of isolating the effect of Brexit on patenting, in sectors where innovation is known to be economically meaningful. In turn, this cleans for the eventuality that results are driven by low-innovation-activity sectors maintaining a low and constant patenting activity before and after Brexit. Second, as emerged in Anna Valero's interview in **Section 2.3**, the chemistry sector represents the **most salient case for Brexit-induced NTBs**. This is a case in which there exists a clear institutional channel linking Brexit to chemical-sector innovation, namely that of **regulatory divergence**; Brexit signed the exit of the UK from the EU REACH regulatory framework, thus requiring firms to duplicate compliance procedures under the new UK Reach framework. In turn, this implied an increase in fixed costs and regulatory complexity for serving both EU and UK markets simultaneously. Third, as can be seen in **Table 14**, the chemistry sector is the only one that displays negative and highly significant effects in response to both Brexit events considered. For instance, the TWFE estimates suggest the following pattern. Concerning the referendum, **moving from zero to full exposure to Brexit is associated with a 12 percentage point lower change in log-patenting.** The coefficient is strongly significant at all conventional levels. Similarly, **TCA effects register a difference of 8.7 percentage points**. In both cases, this is evidence of significant negative effects of Brexit on innovation, in a sector that features a high patenting intensity. To build confidence on this result, the remainder of this section shows estimates equations (2),(3),(4), for the chemistry sector.

Table 15 reports the estimates from the Binarised DiD estimation, performed on the chemistry sector. The overall pattern is unchanged: the effects are negative and highly significant for both events; symmetrically to the TWFE results, effects are stronger for the referendum than for the TCA. **On average, chemical products that were exposed to Brexit experienced a 22 percentage point lower patent growth relative to their minimally exposed counterparts, in the immediate post-referendum period. The TCA counterpart is a 6 percentage point lower patent growth.** Note that both estimates refer to average level treatment effects under the assumption of parallel trends.

Table 15: **Binarised DiD — Chemicals sector effects**

	(1) Referendum	(2) TCA
Treated	-0.2259***	-0.0612***
	(0.0177)	(0.0217)
Years used	2016, 2018	2021, 2024
Observations / Clusters	126	126
R^2	0.0104	0.0005

Notes: Dependent variable is $\Delta Y = \log(1 + \text{patents})_t - \log(1 + \text{patents})_{\text{pre}}$. Column (1) compares 2018 to 2016 (referendum window). Column (2) compares 2024 to 2021 (TCA window). Treated equals one for chemicals SITC4 products with Brexit exposure greater than 0.05. Standard errors clustered at the SITC4 level in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Interesting insights can also come from the **non-parametric estimations**. Due to the lower sample size stemming from the isolation of the chemistry sector, the discretised estimates use quartiles, rather than deciles. **Table 16** reports the quartile-specific average treatment effects. The referendum has a clear trend of negative and statistically significant effects across all exposure quartiles. In particular, **after the referendum, chemical products at the highest quartile of Brexit exposure exhibit a 23 percentage point lower patent growth, relative to their minimally exposed counterparts**. The effects are statistically significant at all conventional levels. It is noteworthy that negative effects are also quite prominent at earlier quartiles. **Figure 20** displays the predicted margins plot of the discretised treatment regression [15], and the associated spline. Both figures confirm the trend described above. Moving to discretised estimates for the TCA, all coefficients are negative, yet smaller in magnitude and only significant for the third quartile of exposure. In that case, **after the TCA, chemical products at the third quartile of Brexit exposure exhibit a 10 percentage point lower patent growth, relatively to minimally exposed counterparts**. Importantly, the TCA does not feature a positive coefficient at the highest-exposure quartile in the chemical sector - which was the case for pooled regressions. Although statistically insignificant, the observed negative coefficient is more in line with theoretical predictions about the effects Brexit on innovation. However, recalling the co-occurrence of COVID and the fact that the magnitude of the third quartile effect is an outlier compared to the other quartiles, any extrapolation should be careful for what concerns TCA effects in the chemistry sector. **Figure 19** plots the predicted margins for the discretised treatment regression, and exhibits the spline counterpart.

[15] The regression table of predictive margins for the Referendum and the TCA can be found in the Appendix, Table 20.

Overall, **the effect of Brexit on innovation in the chemicals sector seems to be a negative one**. The empirical estimates for the chemical sector help building confidence towards the methodology adopted by this thesis. Remarkably, effects are in the direction predicted by theory, as well as by institutional knowledge on the NTBs applied to the chemistry sector. This can be understood in light of a combination of two channels discussed in the mechanism taxonomy, namely, uncertainty and regulatory alignment. The estimates are coherent with a story in which the outcome of the 2016 referendum triggered substantial uncertainty in the chemical sector, which until then operated under the EU REACH regulatory framework. REACH regulated the registration, evaluation, testing and certification required for chemical products to access the European market. In this context, **the referendum introduced uncertainty regarding the future regulatory environment and the degree of continued alignment with EU rules.** Since innovation represents a forward-looking and partially irreversible investment, heightened policy uncertainty may have led firms to postpone or scale down innovative activity already in the immediate post-referendum period. At the same time, Brexit **eventually materialised into regulatory divergence through the implementation of UK-REACH in 2021**. Firms wishing to sell chemicals simultaneously in the EU and the UK faced the need to duplicate compliance procedures, incurring additional fixed costs and regulatory complexity. Interestingly, while the post-TCA estimates remain negative, their magnitude appears smaller than the effects observed after the referendum. Taken at face value, this pattern suggests that the anticipation of regulatory divergence — and the uncertainty surrounding how it would be implemented — may have exerted a stronger short-run impact on innovation than the realised regulatory costs themselves. This section therefore provides tentative evidence that Brexit-induced uncertainty, in anticipation of regulatory divergence, slowed the innovative pace of the sector. More broadly, the results highlight the **importance of regulatory alignment not only in reducing compliance costs, but also in stabilising expectations in innovation-intensive industries**. Ultimately, while results for other sectors remain noisy and harder to interpret structurally, **the chemicals sector provides a coherent case where institutional knowledge and econometric evidence align closely.**

Table 16: Discretised DiD — Chemicals sector effects by exposure quartile

	(1) Referendum	(2) TCA
Quartile 1	-0.1895***	-0.0678
	(0.0494)	(0.0542)
Quartile 2	-0.2440***	-0.0490
	(0.0284)	(0.0355)
Quartile 3	-0.2409***	-0.1032**
	(0.0337)	(0.0450)
Quartile 4	-0.2303***	-0.0245
	(0.0252)	(0.0358)
Years used	2016, 2018	2021, 2024
Observations (SITC4)	126	126
R^2	0.0226	0.0147

Notes: Dependent variable is $\Delta Y = \log(1 + \text{patents})_{t_1} - \log(1 + \text{patents})_{t_0}$ computed over the indicated window. The control group (bin 0) includes chemicals SITC4 products with BrExposure ≤ 0.05. Quartiles are constructed using BrExposure among treated observations only (BrExposure > 0.05). Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

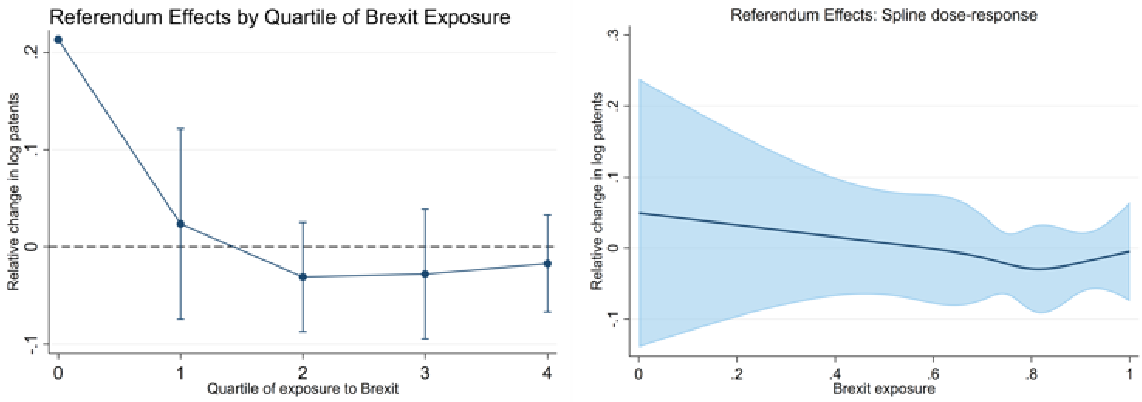

Figure 18: Referendum Effects in the Chemistry Sector

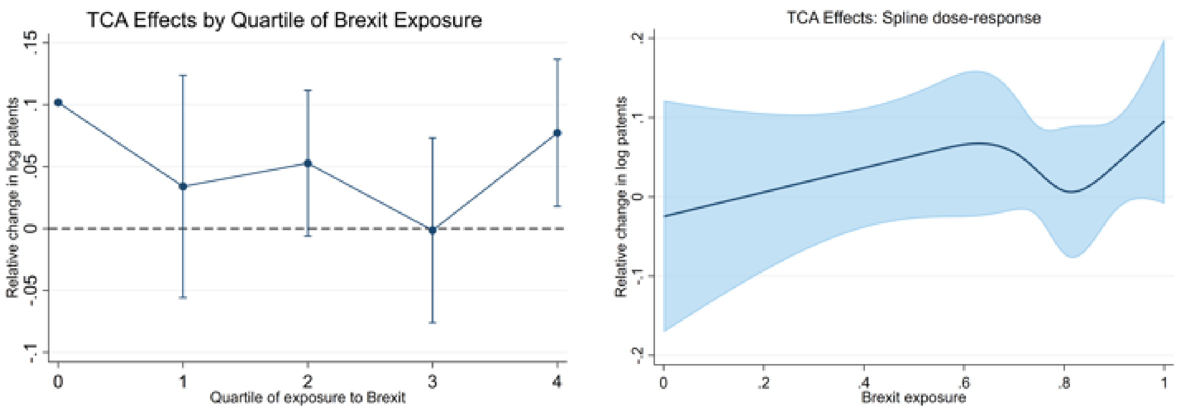

Figure 19: TCA Effects in the Chemistry Sector

As a last comment, the **chemical sector analysis can help putting into perspective and reconsidering the aggregate results discussed in the previous section**. In the pooled discretised specification, statistically significant effects tended to emerge primarily at the highest exposure decile, suggesting that only extremely exposed products experienced sizeable innovation slowdowns. By contrast, within the chemicals sector, the pattern is markedly different. Negative effects are observed across all four exposure quartiles, not only at the upper tail of the distribution. This may suggest that, **in highly innovation-intensive sectors such as chemicals, even moderate levels of exposure to EU markets may have been sufficient to generate a slowdown in innovative activity, with the magnitude of the effect intensifying and then stabilising as exposure increases**. Taken together, this contrast indicates that the **pooled regression may partly reflect an averaging over sectorally heterogeneous vulnerability thresholds**. Different sectors are likely to exhibit different exposure thresholds beyond which Brexit-related frictions become economically binding. For instance, sectors with weaker regulatory integration may only experience measurable effects at very high exposure levels, whereas in chemicals — a sector deeply embedded in EU regulatory structures — even relatively low exposure may trigger innovation responses. Aggregating across such heterogeneous sector-specific dose–response functions mechanically concentrates statistical significance at the extreme of the exposure distribution. The sectoral evidence therefore helps reconcile the pooled and chemistry-specific results, and highlights the importance of accounting for heterogeneity in exposure thresholds when interpreting aggregate effects.

5.5 Limitations and Future Research

The results of this thesis are not immune to a number of **structural limitations**, that this section seeks to highlight. When possible, **directions for future research** are also suggested.

One problem is the absence of **control variables** in the specifications. This thesis decided to operate at a very granular sector level (4-digit SITC). On one hand, this allowed to capture the inherently sectoral effects of Brexit: Brexit is expected to manifest through multiple overlapping sector-level effects rather than a uniform one. On the other hand, it allowed identification to cleanly rely on comparisons of patenting trends between products differently exposed to Brexit, yet belonging to the same broad sectors. The cost of the decision to operate at such granular level is that it is complex to find the relevant control variables that vary at the SITC 4-digit level. However, this choice was done in the awareness that the DiD design would control for time invariant sector characteristics (unit FEs), common time shocks (time FEs). What is left out is time-varying sector-specific shocks correlated with treatment intensity, namely exposure to Brexit. Omitting variables that control for these occurences may bias the estimates. Future research can compromise on this by choosing a coarser sectoral classification, but including some baseline controls.

The second problem that might be affecting this thesis' results is in the **aggregated nature of the data**. Clearly, the ideal scenario would be to estimate these effects with firm-level data. Crucially, this is because it is reasonable to expect that the firms most vulnerable to Brexit-induced uncertainty and NTBs are the smallest ones. Large firms are more resilient, navigate the complexities of diverging regulation and are able to seek alternative solutions with greater ease. This is not true for small firms. However, what drives the aggregate innovation trends of a country - captured by this thesis' data - is precisely the large firms, rather than the smaller ones. This means that micro-data is instrumental in capturing these effects, while aggregate data might risk obscuring the true underlying patterns. Future research may wish to repeat a similar exercise using firm-level data, to emphasise differential effects by firm size.

A third observation is that this thesis focused on **estimating net effects of Brexit**, without explicitly modeling the underlying mechanisms through the data. This implies that more insightful results could emerge from single-channel analysis, and that net effects may be the result of conflicting channels operating simultaneously. Future research could try to disentangle effects by underlying mechanism.

6 Conclusion

Debates over trade integration and the role of innovation have defined the economic landscape of the past decades. This thesis began by observing a puzzle associated with the interaction of the two phenomena. At a time where the role of innovation is discussed as a salient engine of long-run growth, the United Kingdom embarked a major protectionist episode through Brexit. **This thesis posed as an objective to assess the extent to which such a protectionist shock is compatible with innovation.**

By combining the intuition of four experts on Brexit and UK innovation, a taxonomy of mechanisms underlying this relationship was constructed. This exercise speaks directly to the question of what type of economic shock Brexit represents. The answer that emerges is that **Brexit is an inherently multidimensional shock, that passes through trade barriers, migration, knowledge spillovers, uncertainty, expectations, and policy responses, in potentially mutually offsetting ways**. In this conceptualisation, most theoretical predictions point to overall negative effects on innovation. However, this prediction leaves scope for significant sectoral heterogeneity, and for offsetting mechanisms through compensatory policies. The mechanism taxonomy also acts as a map of valuable direction for future research, namely in isolating the innovation effects of the specific mechanisms with clean identification. More broadly, it puts into perspective the widespread implications of protectionist shocks.

Beyond its conceptual contribution, this thesis provided an empirical assessment of how such a multidimensional protectionist shock translates into innovation outcomes. Exploiting pre-referendum variation in product-level exposure to EU trade, and leveraging recent advances in Difference-in-Differences with continuous treatments, the analysis uncovered a nuanced pattern. At the aggregate level, Brexit does not appear to have triggered a broad collapse in UK patenting activity. Instead, the results suggest that innovation displays a marked degree of persistence. **This persistence acts as a structural buffer, mitigating the immediate impact of protectionist shocks, particularly at low and moderate levels of exposure.** However, this buffer is not unlimited; **once exposure crosses sector-specific vulnerability thresholds, negative effects begin to materialise.** In the aggregate analysis, statistically significant slowdowns concentrate at the upper tail of the exposure distribution. When zooming into a highly innovation-intensive and regulation-dependent sector — that of chemicals —

the negative effects emerge already at moderate exposure levels and intensify before stabilising. This sectoral evidence confirms that Brexit's innovation effects are neither uniform nor linear, but instead reflect heterogeneous dose–response relationships shaped by the structure of each industry and the extent of their embeddedness in EU markets.

Taken together, these findings suggest that a protectionist shock does not mechanically halt a country's innovative pace. Rather, its effects depend on the depth of prior integration and the sectoral configuration of vulnerability. In this sense, **Brexit is compatible with continued aggregate innovation performance, yet simultaneously capable of generating meaningful slowdowns where exposure is highest**. The empirical analysis therefore reconciles the apparent contradiction between innovation resilience and protectionist disruption. Methodologically, the thesis also conveys a broader lesson. Analysing Brexit — and protectionist shocks more generally — through aggregate averages risks applying an incomplete analytical lens. **The innovation effects of economic disintegration become visible only when sectoral heterogeneity and variation along the exposure distribution are explicitly incorporated into the analysis.** Understanding the interaction between trade integration and innovation thus requires moving beyond summary coefficients and embracing the inherently uneven way in which such shocks propagate through an economy.

What tentatively emerges from this thesis is that innovation remains persistent in the aggregate despite Brexit, yet becomes vulnerable at higher levels of exposure, with sector-specific vulnerability thresholds. **More broadly, this underscores that the innovation effects of deep disintegration shocks do not surface as abrupt-breaks, but rather hide beneath aggregate stability, emerging only when examined with the appropriate lens.** The question that remains open is whether these underlying effects are temporary adjustments or markers of deeper structural change, and through which mechanisms they primarily operate.

References

Philippe Aghion and Peter Howitt. A model of growth through creative destruction. *Econometrica*, 60(2):323–351, 1992. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/2951599.

Philippe Aghion, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. Competition and innovation: an inverted-u relationship*. *The Quarterly Journal of Economics*, 120(2): 701–728, 05 2005. ISSN 0033-5533.

Philippe Aghion, Antonin Bergeaud, Matthieu Lequien, and Marc J. Melitz. The impact of exports on innovation: Theory and evidence. Working Paper 678, Banque de France, 2018.

Ufuk Akcigit and Marc J. Melitz. International trade and innovation. In Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, editors, *Handbook of International Economics*, volume 5, pages 377–404. Elsevier, 2022. ISBN 9780323988896. doi: 10.1016/bs.hesint.2022.02.006.

Ufuk Akcigit and John Van Reenen, editors. *The Economics of Creative Destruction: New Research on Themes from Aghion and Howitt*. Harvard University Press, 2023. doi: 10.2307/jj.4820341.

Danyal Arnold, Shania Bhalotia, and Swati Dhingra. Deglobalisation in disguise? brexit barriers and trade in services. CEP Discussion Papers dp2110, Centre for Economic Performance, LSE, Jun 2025. URL https://ideas.repec.org/p/cep/cepdps/dp2110.html.

Andrew Atkeson and Ariel Tomás Burstein. Innovation, firm dynamics, and international trade. *Journal of Political Economy*, 118(3):433–484, 2010. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/10.1086/653690.

Jan David Bakker, Nikhil Datta, Josh De Lyon, Luisa Opitz, and Dilan Yang. How Brexit has raised UK food prices. CentrePiece - The magazine for economic performance 628, Centre for Economic Performance, LSE, June 2022.

David Bell. Brexit at the bookies. Centre on Constitutional Change blog post, 2016. URL https://www.centreonconstitutionalchange.ac.uk/blog/brexit-at-the-bookies. Accessed: 2026-03-07.

Nicholas Bloom, Mirko Draca, and John Van Reenen. Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *Review of Economic Studies*, 83(1):87–117, 2016. doi: 10.1093/restud/rdv039.

Nicholas Bloom, Philip Bunn, Scarlet Chen, Paul Mizen, Pawel Smietanka, and Gregory Thwaites. The impact of brexit on uk firms. Working Paper 26218, National Bureau of Economic Research, September 2019. URL http://www.nber.org/papers/w26218.

Nicholas Bloom, Philip Bunn, Paul Mizen, Pawel Smietanka, and Gregory Thwaites. The economic impact of brexit. NBER Working Paper 34459, National Bureau of Economic Research, Cambridge, MA, November 2025. URL https://www.nber.org/papers/w34459.

Paula Bustos. Trade liberalization, exports, and technology upgrading: Evidence on the impact of mercosur on argentinian firms. *The American Economic Review*, 101(1):304–340, 2011. ISSN 00028282. URL http://www.jstor.org/stable/41038790.

Brantly Callaway, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. Difference-in-differences with a continuous treatment. NBER Working Paper 32117, National Bureau of Economic Research, Cambridge, MA, February 2024. URL https://www.nber.org/papers/w32117.

David T. Coe and Elhanan Helpman. International rd spillovers. *European Economic Review*, 39(5):859–887, 1995. ISSN 0014-2921.

Federica Coelli, Andreas Moxnes, and Karen Helene Ulltveit-Moe. Better, faster, stronger: Global innovation and trade liberalization. *The Review of Economics and Statistics*, 104(2): 205–216, 03 2022. ISSN 0034-6535.

Swati Dhingra and Thomas Sampson. Expecting brexit. CEPR Discussion Paper 16970, Centre for Economic Policy Research (CEPR), Paris and London, 2022. URL https://cepr.org/publications/dp16970.

Swati Dhingra, Gianmarco Ottaviano, Thomas Sampson, and John Van Reenen. The consequences of brexit for uk trade and living standards. CEP Brexit Analysis Paper 02, Centre for Economic Performance, London School of Economics, 2016. URL https://EconPapers.repec.org/RePEc:cep:cepbxt:02.

Swati Dhingra, Rebecca Freeman, and Hanwei Huang. The impact of non-tariff barriers on trade and welfare. Lse research online documents on economics, London School of Economics and Political Science, LSE Library, 2023. URL https://EconPapers.repec.org/RePEc:ehl:lserod:117225.

Jonathan Eaton and Samuel Kortum. Technology, geography, and trade. *Econometrica*, 70 (5):1741–1779, 2002. ISSN 00129682, 14680262.

Rebecca Freeman, Marco Garofalo, Enrico Longoni, Kalina Manova, Rebecca Mari, Thomas Prayer, and Thomas Sampson. Deep integration and trade: Uk firms in the wake of brexit. CEP Discussion Paper 2066, Centre for Economic Performance, London School of Economics and Political Science, December 2024. URL https://cep.lse.ac.uk/pubs/download/dp2066.pdf. Revised December 2025.

Michael Gasiorek, Ilona Serwicka, and Alasdair Smith. Which manufacturing industries and sectors are most vulnerable to brexit? *The World Economy*, 42:21–56, 01 2019. doi: 10.1111/twec.12757.

Gene M. Grossman and Elhanan Helpman. *Innovation and Growth in the Global Economy*. MIT Press, Cambridge, MA, 1991.

Giammario Impullitti and Omar Licandro. Trade, firm selection and innovation: The competition channel. *The Economic Journal*, 128(608):189–229, 2018. ISSN 00130133, 14680297. URL http://www.jstor.org/stable/45023220.

Daniel K. N. Johnson. The OECD Technology Concordance (OTC): Patents by Industry of Manufacture and Sector of Use. OECD Science, Technology and Industry Working Papers 2002/5, OECD Publishing, March 2002.

L. B. Jones and C. J. Burns. Reaching for divergence? uk chemical regulation post-brexit. *Integrated Environmental Assessment and Management*, 20(5):1529–1538, September 2024. doi: 10.1002/ieam.4941.

Agnes Norris Keiller. Brexit and investment. CEP Discussion Papers dp2025, Centre for Economic Performance, LSE, Aug 2024. URL https://ideas.repec.org/p/cep/cepdps/dp2025.html.

Wolfgang Keller. International technology diffusion. *Journal of Economic Literature*, 42(3): 752–782, September 2004.

Samuel Kortum and Jonathan Putnam. Assigning Patents to Industries: Tests of the Yale Technology Concordance. *Economic Systems Research*, 9(2):161–176, 1997.

Travis J. Lybbert and Nikolas J. Zolas. Getting patents and economic data to speak to each other: An 'algorithmic links with probabilities' approach for joint analyses of patenting and economic activity. *Research Policy*, 43(3):530–542, 2014. ISSN 0048-7333.

Marc J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1555536.

Luis A. Rivera-Batiz and Paul M. Romer. Economic integration and endogenous growth. *The Quarterly Journal of Economics*, 106(2):531–555, 1991. ISSN 00335533, 15314650. URL http://www.jstor.org/stable/2937946.

Paul M. Romer. Endogenous technological change. *Journal of Political Economy*, 98(5):S71–S102, 1990. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/2937632.

Pian Shu and Claudia Steinwender. The Impact of Trade Liberalization on Firm Productivity and Innovation. *Innovation Policy and the Economy*, 19(1):39–68, 2019.

United Nations Conference on Trade and Development. Non-Tariff Measures to Trade 2019, 2019. URL https://unctad.org/system/files/official-document/ditctab2019d5_en.pdf. UNCTAD DITC/TAB/2019/5.

United Nations Statistics Division. *Standard International Trade Classification, Revision 4*. Number 34, Rev.4 in Statistical Papers, Series M. United Nations, New York, 2006. ISBN 92-1-161493-7. URL https://unstats.un.org/unsd/publication/SeriesM/SeriesM_34rev4E.pdf. ST/ESA/STAT/SER.M/34/Rev.4.

World Intellectual Property Organization. *Guide to the International Patent Classification (IPC)*. WIPO, Geneva, ipc edition 2023.01 edition, 2023. URL https://www.wipo.int/classifications/ipc/en/.

World Intellectual Property Organization. *PCT Applicant's Guide: Introduction to the International Phase*. WIPO, Geneva, 2025. URL https://www.wipo.int/pct/en/guide/index.html. Updated continuously.

World Trade Organization, editor. *World Trade Report 2012: Trade and Public Policies — A Closer Look at Non-Tariff Measures in the 21st Century*. World Trade Organization, Geneva, Switzerland, 2012. ISBN 978-92-870-3815-9. URL https://www.wto.org/english/res_e/booksp_e/anrep_e/world_trade_report12_e.pdf. Annual publication on international trade issues and policy trends.

7 Appendix

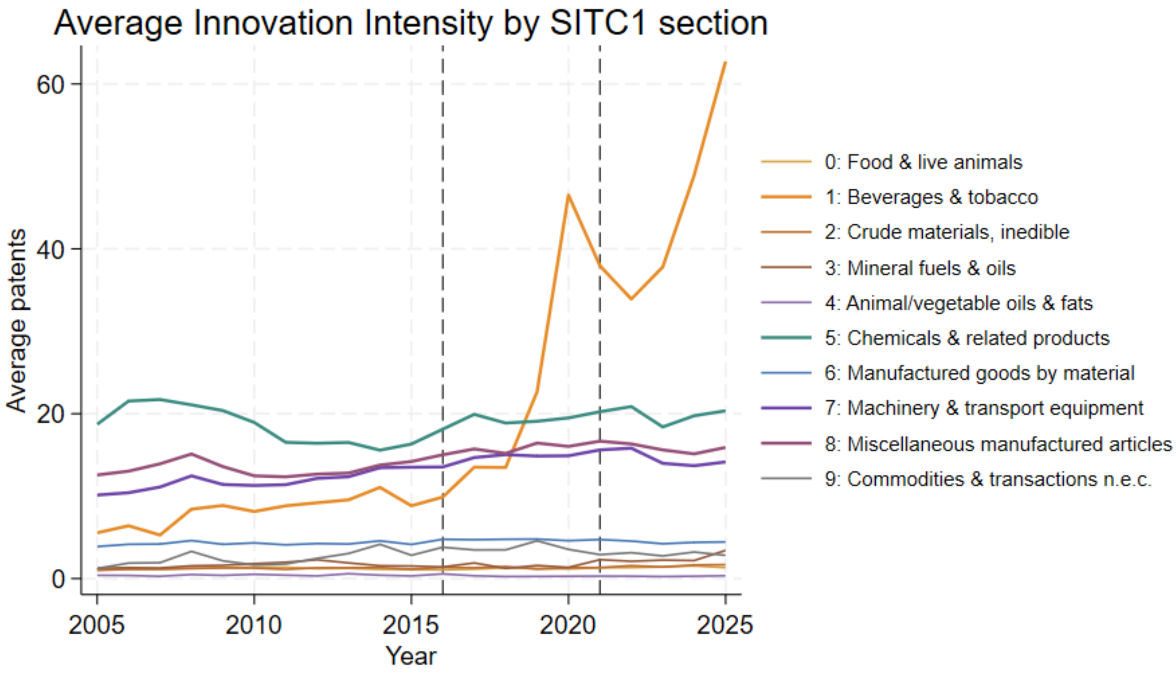

Figure 20: Evolution of Patent Intensity by SITC1 Sector

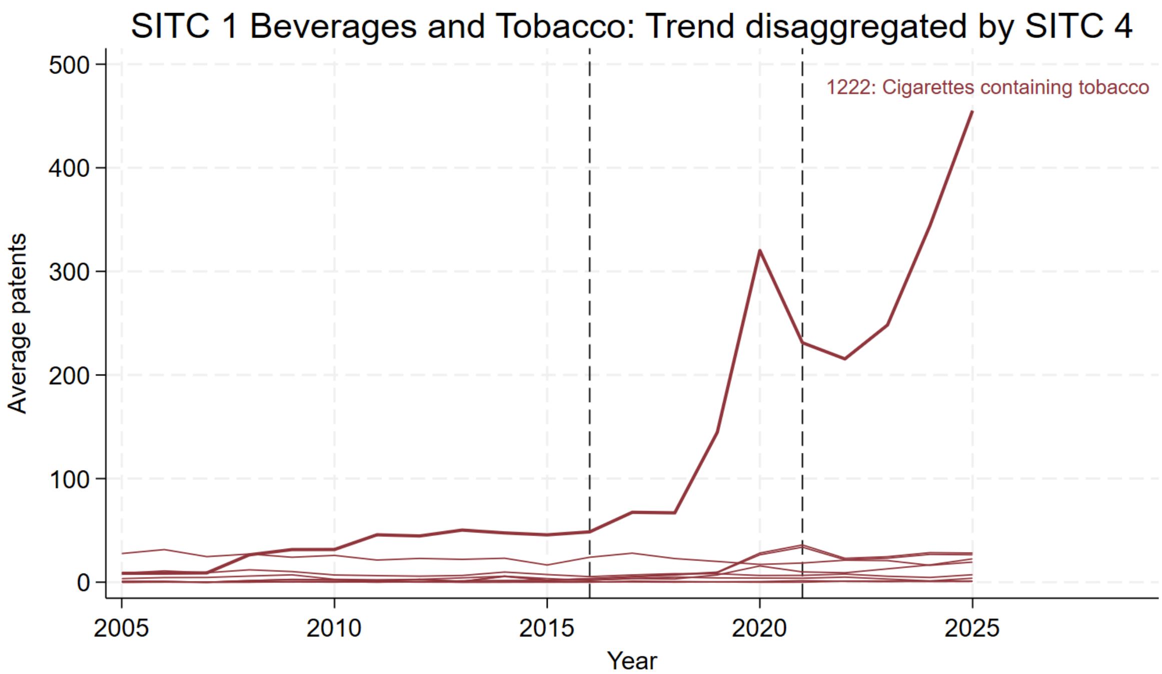

Figure 21: SITC 1 Trend is Driven by Cigarettes

Table 17: SITC4 sectors excluded from exposure matching

S4	Description	S4	Description	S4	Description
0015	Horses, asses, mules (live)	4312	Animal/veg. fats & oils (proc.)	8459	Other knitted garments
0173	Liver, prepared/preserved	5114	Sulphonated derivatives	8511	Footwear (metal toe cap)
0174	Poultry meat, prepared	5791	Scrap polymers (ethylene)	0372	Crustaceans, preserved
0175	Swine meat, prepared	5931	Prepared explosives	0542	Leguminous veg., dried
0176	Bovine meat, prepared	5932	Fuses, detonators	0547	Veg. provisionally preserved
0179	Other prepared meat	5984	Mixed alkylbenzenes	0589	Fruit/nuts preserved
0221	Milk, not concentrated	5991	Municipal waste	0722	Cocoa powder
1122	Fermented beverages n.e.s.	6291	Rubber pharma articles	1211	Tobacco, not stemmed
2123	Fur cuttings	6332	Agglomerated cork	2231	Copra
2474	Coniferous wood, rough	6352	Wooden barrels etc.	2475	Non-coniferous wood, rough
2479	Wood n.e.s., rough	6742	Iron/steel plated tin	2482	Coniferous wood, sawn
2483	Wood strips & friezes	6745	Iron/steel coated	2484	Non-coniferous wood, sawn
2515	Chemical wood pulp	6751	Silicon-electrical steel	2613	Raw silk
2614	Silk waste	6752	High-speed steel	2741	Sulphur
2771	Industrial diamonds	6754	Alloy steel (hot-rolled)	2772	Natural abrasives
2784	Asbestos	6756	Alloy steel (cold-rolled)	2814	Roasted pyrites
2861	Uranium ores	6757	Alloy steel n.e.s.	2862	Thorium ores
2876	Tin ores	6831	Nickel alloys, unwrought	8831	Photographic plates
9110	Postal packages n.e.s.	6851	Lead alloys, unwrought	6994	Steel springs
7465	Roller bearings	6871	Tin alloys, unwrought	7528	ADP units
7598	ADP parts	7616	TV reception apparatus	7922	Aircraft >15,000kg
7928	Aircraft parts	7933	Fishing vessels	7937	Floating structures
8431	Men's overcoats	8432	Men's suits	8437	Men's shirts
8438	Men's underwear	8441	Women's overcoats	8442	Women's suits
8447	Women's blouses	8448	Women's underwear	5114	Sulphonated derivatives
5791	Scrap polymers (ethylene)	7528	ADP units	7598	ADP parts

Notes: SITC4 products with no exposure match. Descriptions abbreviated.

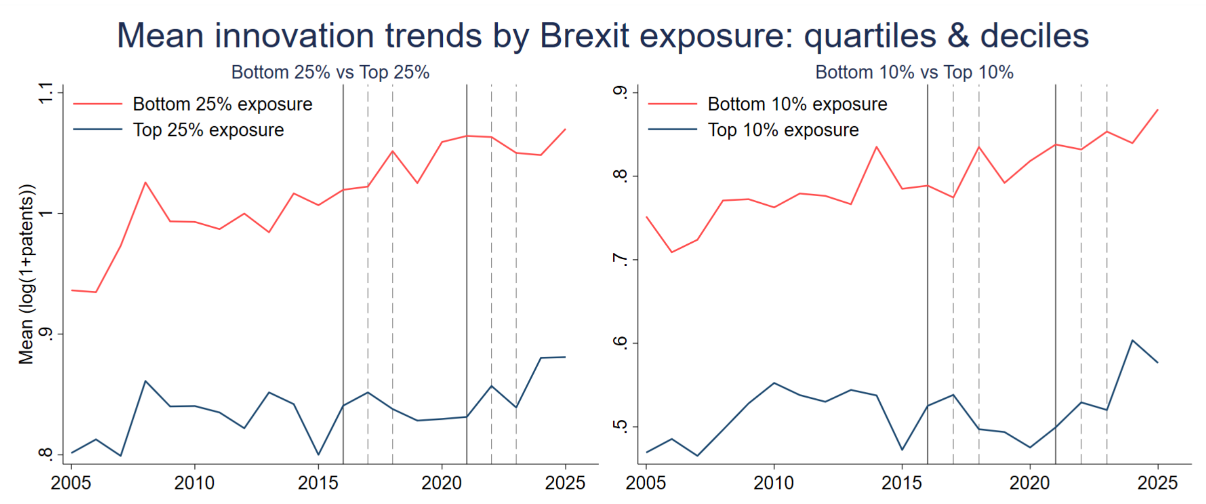

Figure 22: Innovation Trend by Brexit Exposure Quartiles and Deciles

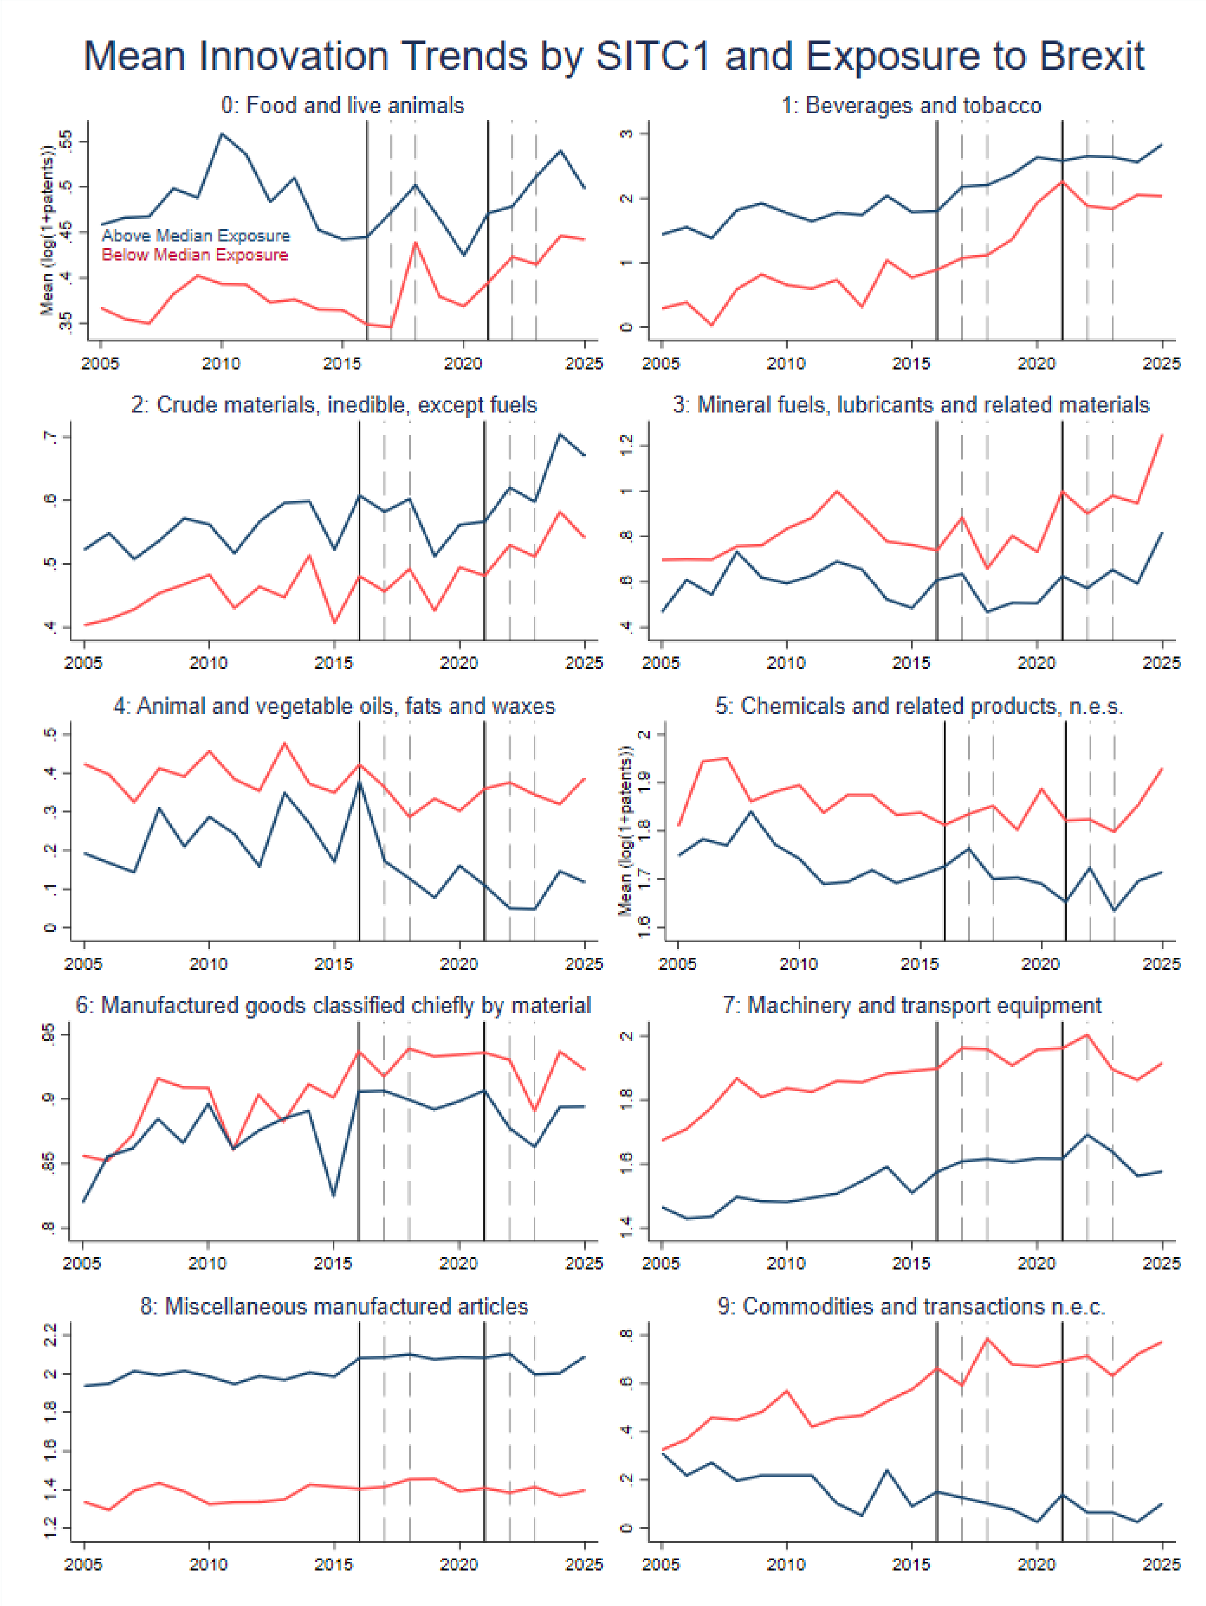

Figure 23: Innovation Trend by Above- and Below-Median Brexit Exposure, by SITC1 Sector

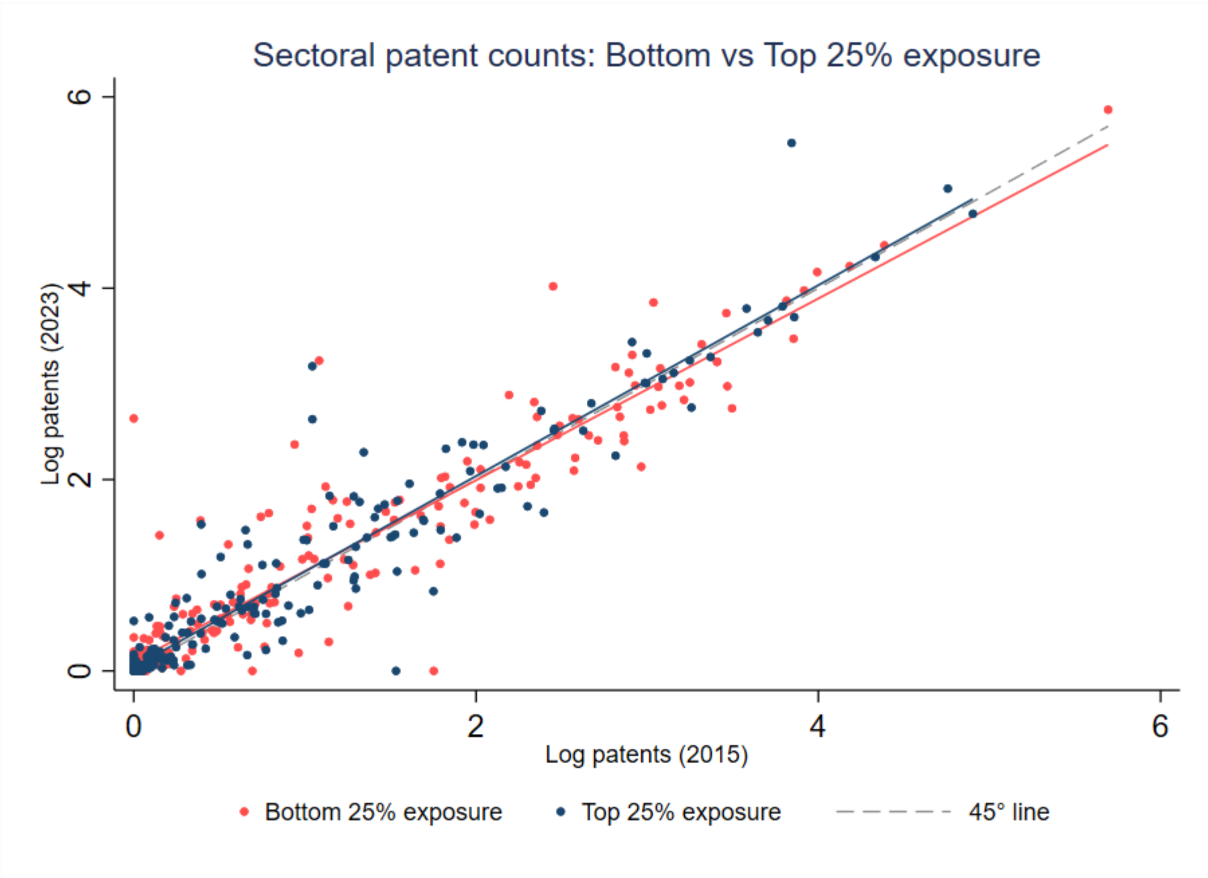

Figure 24: Scatterplot of SITC4 Sectors by Top and Bottom Quartile of Exposure to Brexit

Table 18: **Discretised Treatment by Decile — 2016 Referendum Effects**

	(1) 2018	(2) 2018	(3) 2019	(4) 2019
Bin 0 (untreated)	0.0766432	0.09185	0.00707	0.00612
	(0.05455)	(0.05641)	(0.02957)	(0.03814)
Bin 1	0.0366673	0.01900	-0.01709	0.00162
	(0.02956)	(0.03422)	(0.03181)	(0.03092)
Bin 2	0.0475839	0.02894	0.02581	0.01639
	(0.02885)	(0.02539)	(0.02880)	(0.02761)
Bin 3	-0.0251332	-0.00311	-0.02526	-0.00717
	(0.02112)	(0.02267)	(0.02773)	(0.02733)
Bin 4	0.0746823***	0.06579***	0.02091	0.01768
	(0.02554)	(0.02441)	(0.02610)	(0.02494)
Bin 5	0.0444956*	0.03926	0.02899	0.02865
	(0.02509)	(0.02466)	(0.02556)	(0.02552)
Bin 6	0.0186182	-0.00298	0.01037	-0.00345
	(0.02802)	(0.02715)	(0.02483)	(0.02511)
Bin 7	0.0428287*	0.03446	0.01167	-0.00454
	(0.02514)	(0.02692)	(0.02615)	(0.02662)
Bin 8	0.0320969	0.02121	-0.00692	-0.00798
	(0.02410)	(0.02360)	(0.02920)	(0.02828)
Bin 9	0.0108512	0.02193	-0.00259	0.00670
	(0.02464)	(0.02275)	(0.02828)	(0.02763)
Bin 10	-0.0399046	-0.03244	-0.03960*	-0.03435
	(0.02473)	(0.02524)	(0.02259)	(0.02671)
SITC1 FEs	Yes	No	Yes	No
Years used	2016,2018	2016,2018	2016,2019	2016,2019
Observations	894	948	894	948

Notes: reported values in column (1) are **predictive margins** from the regression. Robust standard errors (Delta-method) in parentheses. Untreated defined as $\mathrm{BrExposure} < 0.05$. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 19: **Discretised Treatment by Decile — 2021 TCA Effects**

	(1) 2023	(2) 2023	(3) 2024	(4) 2024
Bin 0 (untreated)	-0.03980	-0.00141	0.00810	0.05264
	(0.03412)	(0.03330)	(0.04889)	(0.05176)
Bin 1	0.02474	0.01370	-0.02975	-0.03375
	(0.03567)	(0.03355)	(0.03831)	(0.03681)
Bin 2	-0.01491	-0.02557	0.03111	0.00045
	(0.02910)	(0.02702)	(0.03174)	(0.03088)
Bin 3	-0.02887	-0.03931	-0.01344	-0.02123
	(0.02894)	(0.02840)	(0.03396)	(0.03376)
Bin 4	-0.04487	-0.03002	0.00204	-0.01025
	(0.03254)	(0.03164)	(0.03059)	(0.02967)
Bin 5	-0.03244	-0.04056	-0.02533	-0.04490
	(0.02773)	(0.02632)	(0.02955)	(0.02809)
Bin 6	-0.03945*	-0.04258**	-0.03602*	-0.03602
	(0.02172)	(0.02044)	(0.02162)	(0.02273)
Bin 7	0.01464	0.01204	0.00676	0.00979
	(0.02931)	(0.02910)	(0.02735)	(0.02715)
Bin 8	-0.03459	-0.03968	-0.02634	-0.02716
	(0.02633)	(0.02515)	(0.02730)	(0.02727)
Bin 9	0.00520	0.02618	0.00107	0.03421
	(0.02353)	(0.02149)	(0.02966)	(0.02759)
Bin 10	0.02510	0.02095	0.09302***	0.10405***
	(0.03105)	(0.02741)	(0.03614)	(0.03180)
SITC1 FEs	Yes	No	Yes	No
Years used	2021,2023	2021,2023	2021,2024	2021,2024
Observations	894	948	894	948

Notes: reported values are **predictive margins** for all exposure deciles; robust standard errors in parentheses; untreated observations defined as $\mathrm{BrExposure} < 0.05$; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 20: **Predictive Margins by Exposure Quartile — Chemicals Sector**

	(1) Referendum	(2) TCA
Control (Bin 0)	0.2131	0.1018
Quartile 1	0.0236	0.0340
	(0.0494)	(0.0542)
Quartile 2	-0.0309	0.0528
	(0.0284)	(0.0355)
Quartile 3	-0.0278	-0.0014
	(0.0337)	(0.0450)
Quartile 4	-0.0172	0.0773**
	(0.0252)	(0.0358)
Years used	2016, 2018	2021, 2024
Observations (SITC4)	126	126

Notes: Entries report **predictive margins** from quartile-based DiD regressions within the chemicals sector. Dependent variable is $\Delta Y = \log(1 + \text{patents})_{post} - \log(1 + \text{patents})_{pre}$. Control group (bin 0) corresponds to chemicals products with BrExposure ≤ 0.05. Quartiles are defined among treated observations only (BrExposure > 0.05). Robust standard errors in parentheses. ${}^{*}p < 0.10$, ${}^{**}p < 0.05$, ${}^{***}p < 0.01$.