Banc Ceannais na hÉireann
Central Bank of Ireland
Eurosystem

Research Technical Paper

Cross-Border Trade Fictions: The Effect of Negotiation Commitments on Ireland's Response to Brexit

Michael O'Grady

Vol. 2025, No. 2

Non-Technical Summary

This paper examines the effects of Brexit negotiations on trade between Ireland and the UK, focusing on the commitment to maintain a "no hard border" trade arrangement between Northern Ireland and the Republic of Ireland. We argue that this commitment significantly mitigated the negative trade effects typically associated with Brexit for other EU countries. Using an advanced statistical method known as the Augmented Synthetic Control Method (ASCM), we compare Ireland's actual trade patterns with the UK to a hypothetical "synthetic" Ireland that was not subject to the special border arrangement negotiations. This synthetic version is constructed using data from a basket of other EU countries.

Our main findings suggest that Irish exports to the UK were 16.3% higher than the synthetic counterfactual between 2016 and 2019, representing an additional €9.465 billion in export value over the period. Similarly, Irish imports from the UK were 15.5% higher than the counterfactual, amounting to an extra €8.459 billion of imports. These positive effects contrast with the negative trade impacts observed for other EU countries following the Brexit vote. We attribute these differences to the repeated assurances from both UK and EU negotiators about maintaining an open border on the island of Ireland. We argue that Irish firms likely believed they could reorganize supply chains via Northern Ireland, maintaining tariff-free trade with the UK even after Brexit.

We conduct several robustness checks to validate our findings, including changing the timing of the Brexit effect in their model and altering the group of countries used for comparison. Overall, our findings suggest that the unique arrangements for Ireland in the Brexit negotiations had substantial positive effects on Irish-UK trade, contrary to the negative impacts experienced by other EU member states in the aftermath of the Brexit vote. We believe that these results demonstrate the importance of negotiated policy commitments in shaping economic outcomes, even prior to their formal implementation.

Cross-Border Trade Fictions: The Effect of Negotiation Commitments on Ireland's Response to Brexit

Michael O'Grady*

February 2025

Abstract

Employing an augmented version of the synthetic control method, we estimate the effects of the Brexit vote shock and Irish border guarantees on trade patterns between Ireland and the UK, relative to other euro area economies. We let a matching algorithm determine a combination of comparison economies that best resembles the path of Irish bilateral trade with the UK before (i) the Brexit referendum and (ii) guarantees from both the UK and the EU to deliver a "no-border solution" to the island of Ireland. The differences between imports and exports for Ireland and its synthetic doppelganger represent this commitment to a "no-border solution" between Northern Ireland and the Irish Republic, mitigating the negative effects of the Brexit vote shock on bilateral Irish-UK trade. We show that, contrary to the prevailing narrative, Irish imports and exports did not respond to the uncertainty surrounding the Brexit vote to the same extent as other euro area member countries.

JEL Classification: [E65, F13, F42, F68]
Keywords: [Brexit, European Union, Synthetic Control Methods, Trade Policy Uncertainty, Globalization]

---

*Central Bank of Ireland, Irish Economic Analysis Division. michael.ogrady@centralbank.ie.
We thank Benjamin Born, Silvia Calò, Robert Kelly, Martina Lawless, Ciarán Mac Domhnail, Cian Ruane, members of the IRC EU-UK Network, and participants at the CBI-IEA divisional seminar, the IAAE Annual Conference 2023, the ICMAIF Annual Conference 2023, and the IEA Annual Conference 2023 for helpful discussions, comments and suggestions on previous drafts that contributed to the paper. The views in this paper are those of the authors only and not of the Central Bank of Ireland or the European System of Central Banks. All errors are our own.

1 Introduction

While the exposure of the Irish economy to the UK has decreased over the last half-century, it remains substantial across several metrics. Part of the decoupling of the Irish economy from the UK has been attributed to Ireland's decision to join the EEC first and the EC later (O'Rourke, 2019). These moves have brought Ireland closer to its other EU trading partners, and consequently decreased the importance of the UK and its own domestic market. This paper investigates the effect that the decision of the UK to leave the European Union, its Single Market and its customs union, had on the external sector of the Irish economy.

Notwithstanding the decrease in the role of the UK market for Ireland in relative terms, it has remained an important trading partner for Irish firms, particularly in certain economic sectors (e.g. the construction and agriculture industries). In 2021, the UK was Ireland's second-largest merchandise export market (accounting for 11% of total merchandise exports), and the largest merchandise import market (with 19% of all Irish merchandise imports coming from the UK). In comparison, the aggregate euro area accounted for 25% of merchandise exports and 13% of merchandise imports in 2021.

Table 1 presents import and export data for Ireland's main trading partners. In 2018, almost 80 per cent of Irish exporting firms exported goods to the UK.[1] Within this subset, almost half of firms had no other outward external trade linkages, only exporting goods to the UK. Similarly, 85 per cent of all importing enterprises purchased goods from the UK. Within this subset, 40 per cent of firms imported goods exclusively from the UK.

There are number of structural characteristics that made the financial and economic interconnections between the Irish and UK economies more substantial than the other member states of the European Union. Primarily, there is a shared land border

[1] Central Statistics Office, Profile of Trading Enterprises 2018.

Table 1: Main Irish Trading Partners, millions of euro, 2021

	Exports		Imports	
Country	Merchandise	Services	Merchandise	Services
USA	€52,567	€41,117	€18,125	€143,667
UK	€18,182	€40,241	€19,512	€23,405
Germany	€17,774	€21,174	€7,297	€7,075
Netherlands	€9,387	€12,972	€4,499	€12,497
China	€11,197	€9,417	€8,452	€5,645
France	€5,348	€11,397	€10,315	€4,808
Belgium	€13,530	€4,926	€2,223	€5,842
Switzerland	€2,813	€6,156	€5,140	€7,180
Italy	€4,366	€9,388	€2,193	€3,673
Japan	€2,715	€11,497	€1,430	€2,355

Source: Central Statistics Office, Value of Merchandise Trade Statistics, 2021

between the UK and Ireland, making trade linkages between Ireland and Northern Ireland costly to replicate. The common language reduces the need for translation services, allows for more direct communication and provides a sense of shared culture. Similarly, the common legal system reduces administrative and legal costs, preventing firms from having to draft alternative sets of contracts for their traded and non-traded merchandise.

In this paper, we examine the relative effects of the Brexit vote on the Irish economy, in light of the specific guarantees of "no hard border on the island of Ireland" and "no customs border in the Irish Sea", which potentially drove the expectations of Irish firms regarding Irish-UK trading arrangements that would ultimately arise following extensive withdrawal negotiations.[2] We focus, in particular, on the anticipation effects of Brexit and the Northern Ireland Protocol on the Irish external sector. To do so, we use an approach similar to that adopted by Born et al. (2019) in analyzing the effects of Brexit on the GDP of the UK. We apply an augmented form of the synthetic control methodology

[2] The various forms of the Northern Ireland Protocol suggested keeping Northern Ireland in some aspects of the Single Market, to prevent the development of trade barriers between the Republic of Ireland and Northern Ireland.

(SCM) to build a counterfactual Ireland, which was subject to the Brexit vote shock, but not to the special "no hard border" negotiation commitments that linked Ireland to the UK economy in a stronger manner than other EU member states.

In doing this, we deviate from one of the major assumptions in Born et al. (2019): the authors assume that the result of the Brexit referendum only affected the UK economy, and its effects did not reverberate beyond UK borders. In this work, we test the hypothesis that Ireland was instead affected by the decision of the UK to leave the EU, but to a differing extent than other EU countries. We justify our hypothesis with the repeated guarantees, given by both the European Commission and the UK government, that free movement of goods and labour between Northern Ireland and Ireland would be a central bloc of post-Brexit trade arrangements.[3]

Consequently, we propose the hypothesis that the effects of this commitment by the EU and UK, to avoid a "no-border solution", combined to prevent the Brexit vote shock from negatively affecting bilateral trade between Ireland and the UK to the same extent that trade with other EU member states was impacted. Effectively, we argue that if Ireland was a continental European economy, with a near-similar set of demographic, economic and financial conditions, but without expectations of a treaty that would reinforce post-Brexit trade linkages with the UK, the effects on bilateral UK imports and exports following the Brexit referendum shock would have been substantially more negative.

To estimate our counterfactual hypothesis, we use the augmented synthetic control method (ASCM) of Ben-Michael, Feller, and Rothstein (2021): a transparent, unbiased and fully data-driven approach that corrects for biases present in the standard synthetic control method. The ASCM algorithm determines which combination of other European

[3] Statements regarding the preference for a "no-border solution" were made as early as July 2016, less than one month after the Brexit referendum vote.

economies matches the time path of Irish import and export series, before the Brexit vote, with the highest degree of accuracy. The better the synthetic doppelganger constructed by the algorithm for the Irish economy (from the weighted combination of other EU economies) before the Brexit vote "treatment", the more precise our causal estimate of the divergence of the Irish trade response following the Brexit vote. We use as large as possible a dataset, subject to restrictions preventing heterogeneous economic and financial conditions from biasing our results, to obtain the best match possible from our algorithm.

Overall, our estimates suggest that Ireland did indeed respond to the Brexit vote in a manner that was substantially different to other EU economies. While there was limited difference in the estimated response of net exports, both gross imports and exports were significantly higher than the counterfactual doppelganger for Ireland proposed by the ASCM algorithm. These results suggest that, over the 2016q3 - 2019q4 period, imports would have been €8.459 billion (16.3%) lower, and exports would have been €9.465 (15.5%) billion lower, had Irish firms responded to Brexit in a similar way to their European counterparts.[4]

Given the uniqueness of the United Kingdom's EU membership referendum as a natural experiment in economic disintegration, considerable work has been devoted to estimating the macroeconomic and financial effects of Brexit, with a particular focus on UK economy. The research most comparable to our analysis is Born et al. (2019), who use a multi-estimation approach to identify the causal effect of the Brexit vote on the UK's macroeconomic performance and the underlying transmission channels through which these effects operate. Using the synthetic control method, they create a doppelganger country that evolves in the same manner as the UK economy in the absence of the Brexit

[4] Values denominated in real 2019 euro terms.

vote, estimating the difference in GDP between the real and synthetic UK to be between 1.7% and 2.5% by the end of 2018.

Graziano, Handley and Limão (2021) estimate how shocks to the probability of Brexit affected bilateral export investments and trade flows between the UK and the EU. Decomposing monthly variation in UK exports prior to the June 2016 referendum, they identify uncertainty, demand and supply shocks to both bilateral UK-EU exports and export entry investment series, allowing for the quantification of trade effects from large permanent changes in the probability of Brexit. They find that shocks to the probability of Brexit reduce trade flows and trade participation, with these effects larger in products where the potential for upward tariff revisions are highest. At the average WTO tariff rate of 4.5%, they find that the post-referendum increase in uncertainty implies a reduction of $11 - 20\%$ of UK-EU exports.

Siedschlag and Koecklin (2019), also using the synthetic control method, estimate the impact of Brexit uncertainty on FDI-related employment in Northern Ireland after the Brexit vote. Using regional data from the EU-27, they generate a synthetic Northern Ireland that was not subject to the events of Brexit. Results indicate that uncertainty from the outcome of the Brexit vote caused a decline in new FDI-related jobs in Northern Ireland, with a reduction of 31% estimated across the eight quarters following the Brexit vote in 2016q3.

Douch and Edwards (2021) analyse the effects of uncertainty and anticipation shocks from Brexit on trade between the UK and 14 EU and 14 non-EU trading partners. Controlling for exchange rate and GDP changes, UK exports to both groups of countries declined in the aftermath of the Brexit referendum, but also in the period between the Conservative party general election win (May 2015) and the referendum. Their estimates suggest that UK exports to the EU fell by $20 - 25\%$ over the 2015q2-2018q3 period, while UK exports towards non-EU countries declined by 15%. The authors attribute the decline in UK-EU exports to an inward supply shift, in anticipation of a demand decline in the event of Brexit. Similarly, the reduction in exports to non-EU countries is linked to concerns regarding UK competitiveness, likely from increased supply chain costs.

Crowley, Exton and Han (2018) estimate the effect of uncertainty due to trade agreement renegotiation on the export participation decision of UK firms. Using the Handley and Limão (2017) model of exporting under trade policy uncertainty, they develop measures of the trade policy uncertainty facing firms exporting from the UK to the EU after June 2016. Applying a dataset containing all UK export transactions at the firm-product level, and the range of WTO tariffs those transactions would face in the event of renegotiation terminating with no deal, the authors estimate the counterfactual actions of firms exporting from the UK to the EU in the absence of Brexit-related trade policy uncertainty. Their results suggest that reduced entry accounts for a decline of £$0.2 - 1.5$ billion of export value from the UK to the EU in 2016, while the reduction in value from the increased exiting was between £$0.19 - 1.4$ billion.

Our work also intersects with research into the effects of free trade agreements (FTAs) on trade dynamics. Baier and Bergstrand (2014) use a theoretically-motivated gravity equation using differenced panel data to estimate the dynamic effects of FTAs on bilateral trade flows. Using supporting arguments from theoretical models of trade liberalization, the authors account for the multilateral price variables in a panel gravity equation. Estimating the average treatment effect of FTAs, they find that a free trade agreement increases bilateral trade by about 58 per cent on average. First-differencing the data and controlling for concurrent, lagged and future changes in the FTA, the average treatment effect on trade flows between a pair of countries rises to 86 per cent.

Anderson and Yotov (2016) estimate the volume and terms of trade effects of free trade agreements implemented between 1990 and 2002. In contrast to much of the empirical gravity literature, the authors estimate trade gravity equations disaggregated at the 2 digit ISIC level across manufacturing sub-sectors. Using a Poisson pseudo-maximum likelihood (PPML) estimator, the approach controls for multilateral resistances, FTA endogeneity and the gradual phasing-in of the FTA effects. Across product categories, the average treatment effect of FTAs on bilateral trade flows is estimated to range from 33 per cent to 264 per cent, with almost all product categories showing positive and significant effects nine years after the introduction of the FTA.

The remainder of the paper is structured as follows. Section 2 presents a brief overview of the Irish-UK trade relationship, and the background to the UK's withdrawal from the EU. Sections 3 presents the methodology behind the augmented synthetic control method, and the improvements its use delivers above and beyond standard synthetic control approaches. Sections 4 provides a description of the data used in our empirical analysis. Section 5 presents the results of our ASCM estimation, and how these results can be viewed in relation to our hypothesis. Section 6 concludes the paper.

2 Ireland and the UK Exit from the EU

2.1 Brexit and the Complication of the Irish Border

From an EU-wide perspective, considerable economic uncertainty has derived from the result of the Brexit referendum in 2016. In Ireland, policymakers devoted substantial efforts to identifying the channels through which trade between the UK and Ireland could be disrupted, and the degree to which financial and economic interlinkages would be affected as new trade agreements were negotiated. A key element of this concern related to Northern Ireland: the border between both countries represents the only land-border between the UK and the EU post-Brexit, and the freedom of movement of goods and people across this border were topics of high-level political discussion in the wake of the Brexit vote.

While there are several economic, political and sociological dimensions to Brexit's ramifications on both the UK and Ireland, trade protocols are potentially one of the more impactful elements for both countries. The final form of Northern IrelandâĂŹs post-Brexit trading arrangements, designed to replace EU membership, was viewed by many analysts as being a barometer for the new relationship that would exit between the EU and the UK. Consequently, statements from politicians and officials regarding the possible nature or content of such protocols received intense media coverage prior to and throughout the negotiation process. A timeline of the key events leading up to the separation of the UK from the EU is provided in Figure 1.

Figure 1: Timeline of Key Brexit Event Dates

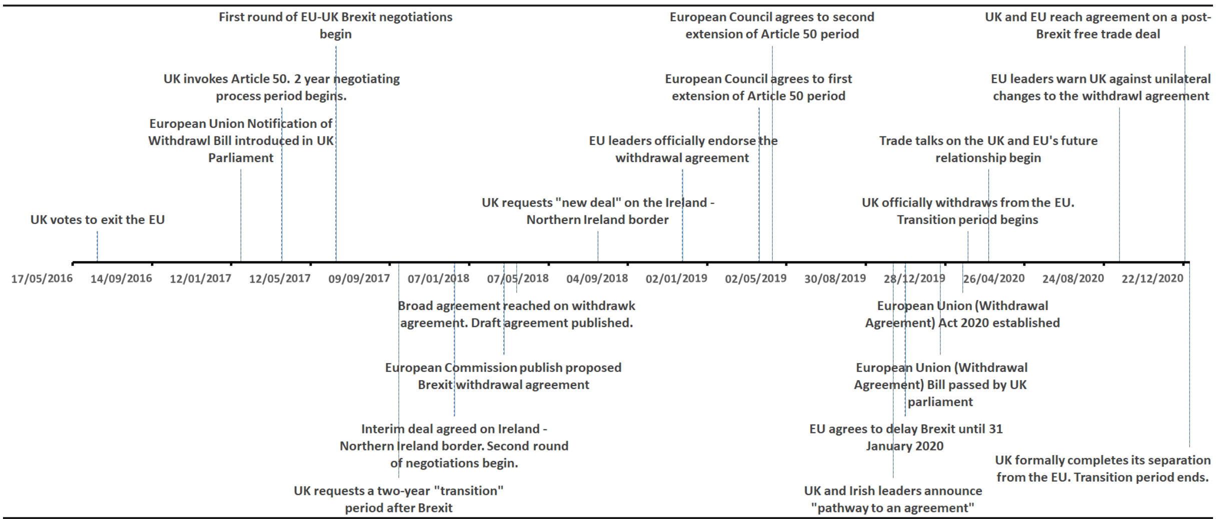

Following the serving of a withdrawal notice in March 2017, twenty-nine months of exit negotiations concluded with the UK formally relinquishing its EU member status in January 2020. Despite the 177 page Withdrawal Agreement that specified the framework for continued arrangements between the UK and Europe, this only represented a preliminary step in codifying the nature of long-term agreements on a range of matters, including: trade; financial services; security; legal jurisdiction; and immigration and mobility. The impact on Ireland and the issue of the Northern Ireland border (once it became an EU external border) were major complications in withdrawal process negotiations. The core issues related to the incompatibility of the three central objectives of the UK government:

I *Internal and external sovereignty with respect to control on internal regulation and developing an independent, more advantageous trade policy.*

II *No "hard border" between Northern Ireland and Ireland.*

III *No border solution that created separation between Northern Ireland and the rest of the UK.*

These objectives were referred to as the "Brexit Trilemma": delivering on any two of the objectives is entirely feasible, but only at the expense of the third objective.[5] These complications lead to several iterations of the Brexit withdrawal agreement that attempted to resolve the Brexit Trilemma, including the Irish Backstop, the Northern Ireland Protocol and the Windsor Framework.

For the majority of the negotiation period, discussions around the Irish border centered around keeping Northern Ireland in some aspects of the Single Market. However, in October 2019, the UK government renegotiated the terms of the Withdrawal Agreement, removing Northern Ireland from the EU customs union, but avoiding the introduction of tariffs or restrictions on goods crossing the Irish border in either direction. The UK formally left the EU in February 2020, but remained in the EU Single Market and Customs Union until December 2020. The final Withdrawal Agreement, implemented in October 2023, was a hybrid approach to

[5] As discussed by Springfield (2018), simultaneously delivering on the full set of objectives was fundamentally impossible. The UK could leave the single market and customs union, maintain a border-free Ireland, and pursue a whole-UK approach to Brexit, but not all three at the same time.

customs membership: goods could be imported tariff-free into Northern Ireland from either Ireland or Great Britain, subject to the condition that their final destination is not the counter-party customs union. By placing Northern Ireland in distinct customs and regulatory regimes versus other UK jurisdictions, the Protocol has introduced more barriers to trade in goods between Great Britain and Northern Ireland than trade in goods between Ireland and Northern Ireland.

We can view the events of the Brexit vote and withdrawal negotiations as a natural experiment in economic disintegration, as (i) the outcome of the vote and resulting negotiation requirements were an unanticipated shock, (ii) external trade performance was not a causal reason for the outcome of the Brexit vote, and (iii) Ireland was far more exposed to the results of the Brexit border negotiations than other EU member states. Negotiations to prevent a hard border between Northern Ireland and Ireland were designed to mitigate any risks that could destabilize the Good Friday Agreement, not to give Ireland a preferential trade relationship with the UK.

2.2 Qualifying the differing Effects of EU-UK Brexit Negotiations on Ireland

As an attempt to provide more certainty around Irish-UK trading arrangements, the "no border" commitment potentially mitigated the immediate negative trade effects that the Brexit vote could have had on the Irish economy. For the majority of the period between the Brexit vote (July 2016) and the exit of the United Kingdom from the European Union (January 2020), the negotiating positions of both the EU and the UK were strongly in favour of maintaining an open border between Ireland and Northern Ireland.

In July 2016, the head of the Irish government stated that "there will not be a hard border" on the island of Ireland. By August 2016, the secretary of state for exiting the

European Union had stated that the UK did not want "a hard border or unnecessary barriers to trade", while the Irish ambassador to the UK stated that the EU "would be very slow to do anything that would in any way cut across or create difficulties" for cross-border trade between Ireland and Northern Ireland. Even the lead EU negotiator, Michel Barnier, stated in December 2016 that he was "extremely aware" of border issues on the island of Ireland and said he was determined that the Brexit negotiations would "find a way" to preserve existing relationships. This was despite his stated reluctance to discuss negotiating positions prior to the formal commencement of negotiations.

With the appointment of negotiating teams and the formal notification of the UK's intention to withdraw from the EU (Article 50), these guarantees were reinforced by the prioritisation of the Irish border issue.[6] Repeated guarantees were given during the negotiating period regarding the intransigence of either party to let customs checks or controls inhibit cross-border trade on the island of Ireland.

The EU focused initial discussions on exit issues (including the border in Northern Ireland), and would only agree to start talks on the future relationship once "sufficient progress" had been achieved. As a direct consequence of this, UK and EU negotiating teams developed a supplemental treaty as part of the draft Brexit withdrawal agreement. Known as the Northern Ireland protocol, this treaty was designed to prevent both a hard border on the island of Ireland and a customs border in the Irish Sea. It gave quota and tariff-free status to a component of EU-UK trade flows: goods traveling from Northern Ireland to Ireland or other UK jurisdictions. This arrangement was unique across all possible bilateral trade pairings between the two unions, and showed the commitment by the UK government and European Commission to a "no-border solution".

---

[6] Resolving the "complex impacts of Brexit on the Northern Ireland - Ireland border" was one of only two priority items (the other being citizens' rights) listed by both sides upon invocation of Article 50.

Ultimately, this form of Northern Ireland Protocol did not become part of the final Brexit withdrawal agreement. While UK and EU negotiators supported the supplemental treaty, its ratification was rejected by UK parliament on three separate occasions in 2019.[7] In October 2019, a final set of border negotiations removed the Northern Ireland protocol from the Withdrawal Agreement, replacing it with an alternative version where Northern Ireland formally remains within the UK customs territory and internal market for goods, but is also required to comply with EU customs, internal market and valued added tax rules. By placing Northern Ireland in a separate customs and regulatory regime relative to other UK jurisdictions, there are greater barriers to trade in goods between Great Britain and Northern Ireland than between Ireland and Northern Ireland.

Similarly, the centrality of Irish border issues (as one of three areas of focused negotiation in the Withdrawal Agreement) received considerable news coverage in both Ireland and the UK; from initial 2016 proposals to harmonise external Irish border controls with the UK, via the 2017-2019 negotiations to develop a Northern-Ireland specific backstop, through to the official implementation of the Northern Ireland Protocol within the Withdrawal Agreement in December 2020, it was perceived that arrangements to maintain the *status quo* trade conditions between Ireland and Northern Ireland (and the wider UK) could be delivered by the EU and UK negotiating teams. A central facet of the EU negotiating strategy was to be transparent about its objectives and red lines. By keeping the public informed on negotiations, it increased the leverage of EU negotiators during discussions, while also making it difficult for the UK to hold bilateral talks with individual EU member states. Thus, the public was fully aware that

[7] Despite these rejections, Michel Barnier further committed to the concept of an open Irish border, stating in January 2019 that the backstop was "part and parcel" of the UK's Brexit withdrawal agreement and would not be renegotiated.

the EU sought a “no-border solution” on the island of Ireland for the majority of the negotiating period.

Given the continued compliance of the most geographically accessible part of the UK with EU customs, internal market and valued added tax rules, Irish firms trading with the UK since the 2016 vote likely did not view Brexit with the same finality as firms in other EU member states. If Irish firms could re-organise supply chains, to have goods sourced or routed via Northern Ireland without the imposition of tariffs, then existing trade patterns could be maintained until the UK formally exited the single market. Once Brexit occurred, and trade with Great Britain became subject to tariffs and non-tariff barriers, Irish firms could either (i) export (import) goods directly to (from) Northern Ireland, or (ii) export (import) goods that would undergo processing in Northern Ireland before being sent to Great Britain (Ireland). Goods traded through these channels would not be subject to additional EU or UK tariffs as a consequence of Brexit. Anticipation of these conditions, essentially maintaining the *status quo* for trade on the island of Ireland and reducing the uncertainty around trading conditions, potentially lead to Irish firms responding in a different manner to the Brexit vote than their EU counterparts, with respect to trade with the UK.

It should be noted that, for this hypothesis to hold, we would need to see a material change in the trade patterns of Irish imports and exports with both Northern Ireland and the UK, once the UK exited the EU Single Market in January 2021. If Irish firms could not reorganize supply chains through Northern Ireland, or if their reliance on firms within Great Britain for imports and exports was too substantial, then the composition of trade with various UK jurisdictions would not be affected by the formal exit of the UK from the single market. Using two distinct datasets on Irish-UK trade, we show supportive evidence for such changes in the next section.

2.3 Mechanism Support

As discussed in Section 2.2, the hypothesis we present in the paper for a differing effect on Irish trade with the UK (relative to other EU economies) stems from the commitment to a "no hard border" outcome to Brexit, and the focus on the land border issue between Ireland and Northern Ireland throughout the Brexit negotiations. Had Irish firms not believed that there would be an avenue through which tariff-free trade channels could be maintained with the UK in a post-Brexit world, we argue that their decoupling from the UK goods exports and imports markets would have been more consistent with the patterns observed in other EU economies in the advent of the Brexit vote shock in 2016.

For this argument to hold, we would expect to see several patterns in the trade data, following both the Brexit vote shock in June 2016 and the realization of Brexit (specifically, the UK leaving the EU customs union) in January 2021:

I No significant decline in the share of Irish intra-UK trade flows to Northern Ireland following the Brexit vote, with an increase in the share of UK trade through Northern Ireland once the UK exits the Single Market

II No significant change in maritime shipping volume between Irish and UK ports following the Brexit vote, but an aggregate decline in shipping volumes once tariff and non-tariff barriers are introduced in post-Brexit.

To support the mechanism behind our hypothesis, we provide evidence from two separate datasets on Irish import and export flows with the UK: bilateral intra-UK trade value data from the Irish Central Statistics Office (CSO) and Eurostat maritime shipping volumes between Irish and UK ports. Together, these data should provide an overview of the impact of both the Brexit vote and the UK exit from the EU Single Market, which should inform the main analytical results presented in this paper.

2.3.1 Irish Intra-UK Trade Post Brexit

Figure 2 presents monthly export and import flows between Ireland and Northern Ireland between 2015 and 2023, covering the period before the Brexit vote shock and after the exit of the UK from the EU single market. From both the exports and imports sub-figures, there is no evidence of a structural break in the data around the timing of the Brexit vote. However, both series show a material increase in values once the UK exited the EU Single Market in January 2021. Relative to 2015-2019 values, average export flows increased by 102% over the 2021-2023 period. Similarly, average monthly import flows increased by 154% over the same reference periods. At a quarterly level, the differential in (real 2019) trade flows between the post-referendum, pre-Brexit period and the post-Brexit period amounts to €588 milion in quarterly exports and €751 million in quarterly imports.

Figure 3 presents the share of Irish imports and exports to the UK accounted for by Northern Ireland, over the 2000-2023 period. Prior to the Brexit vote shock, the annual share of UK exports to Northern Ireland averaged 9.7%, with a peak value of 11.7% in 2014. Similarly, average shares of total UK imports were 6.5%, peaking at 7.2% in 2004. In 2016 (the first full year following the Brexit vote shock, these maximal values were exceeded for both exports (12.8%) and imports (9%), with near-monotonic increases in each subsequent year of the sample. Again, once the UK exits the EU Single Market in 2021, we see the structural breaks from Figure 2 increase Northern Ireland's share of Irish merchandise trade with the UK, with average exports (21.8%) and imports more than double their historical average share of UK trade flows over the 2021-2023 period.

Not only do these results support the mechanism we propose in our analysis, they also eliminate a competing explanation from being the source of potential differences in the Irish trade response to the Brexit vote shock. If Irish exporters and importers

Figure 2: Export and Import Flows between Ireland and Northern Ireland

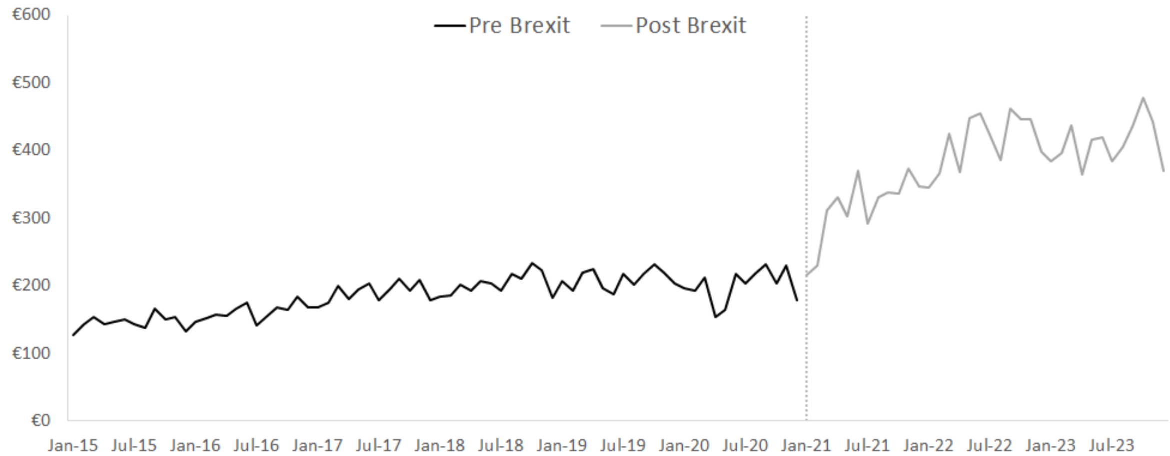

(a) Irish Exports to NI

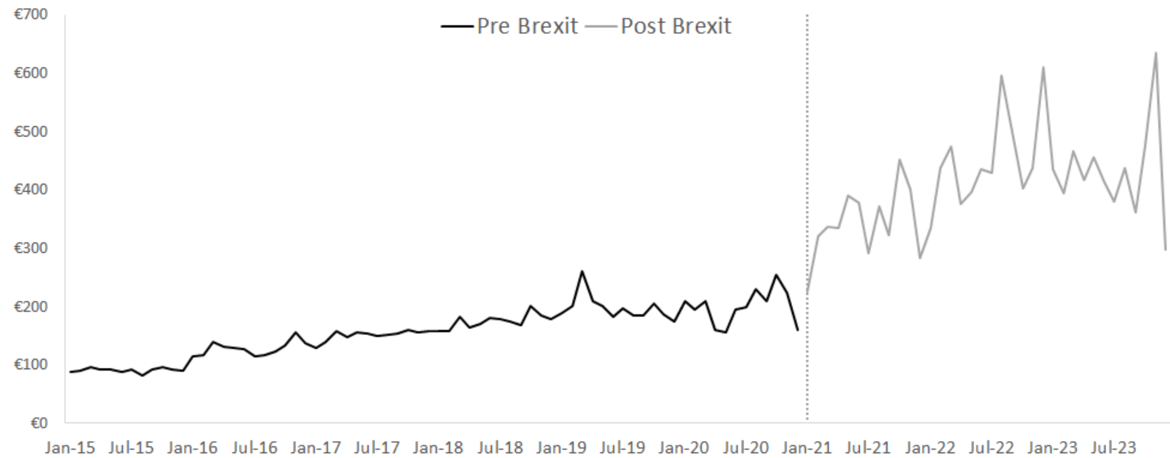

(b) Irish Imports from NI

Source: Central Statistics Office, External Trade database. Values in real 2019 € millions.

had a substantially greater reliance on (or inability to diversify away from) existing trade markets in Great Britain, relative to continental EU firms, then this could have accounted for the absence of a decline in trade flows between the UK and Ireland following the Brexit vote. However, such a reliance could not explain the compositional shift in trade flows from Great Britain to Northern Ireland, suggesting an alternative mechanism that presented Northern Ireland as a preferential trade destination during the period between the Brexit vote and the UK exit from the EU Single Market.

Figure 3: Share of Irish - UK Trade Flows accounted for by Northern Ireland

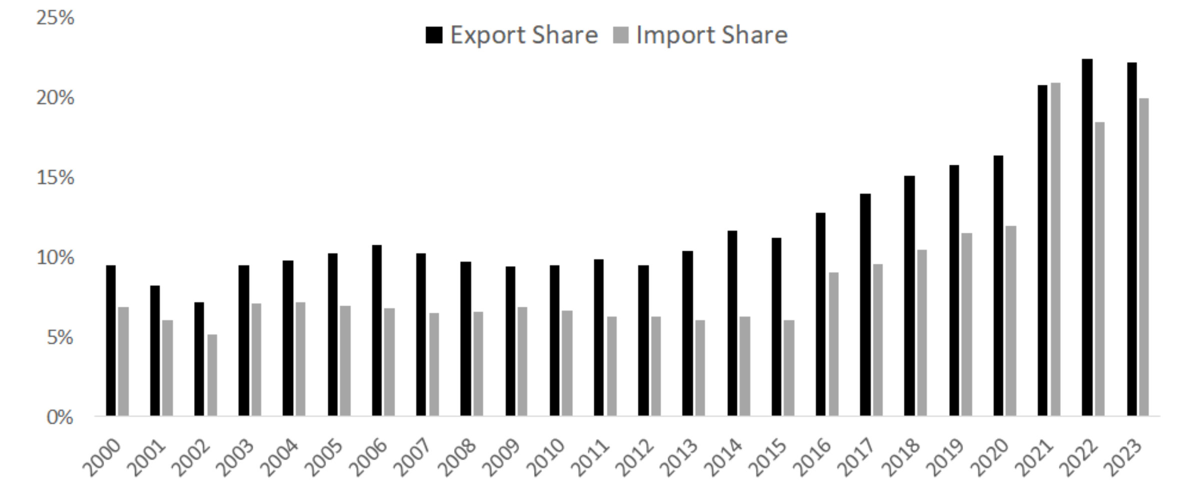

Source: Central Statistics Office, External Trade database.

2.3.2 Irish-UK Maritime Shipping Trade Post Brexit

To further reinforce the differing effects of supply chain reorganization that occurred following the Brexit vote shock, Figure 4 presents Eurostat maritime shipping data to / from UK ports. To compare the effects of both the Brexit vote shock and the UK exit from the single market, we present aggregate outwards (export) and inwards (import) flows with all UK ports, for both Ireland and the aggregate EU (excluding Ireland). For comparability, we exclude intra-country shipping volumes and index the average volume in the four quarters prior to the vote shock (2015q3-2016q2) to 1.

Figure 4: Outward and Inward Maritime Shipping Flows with UK Ports

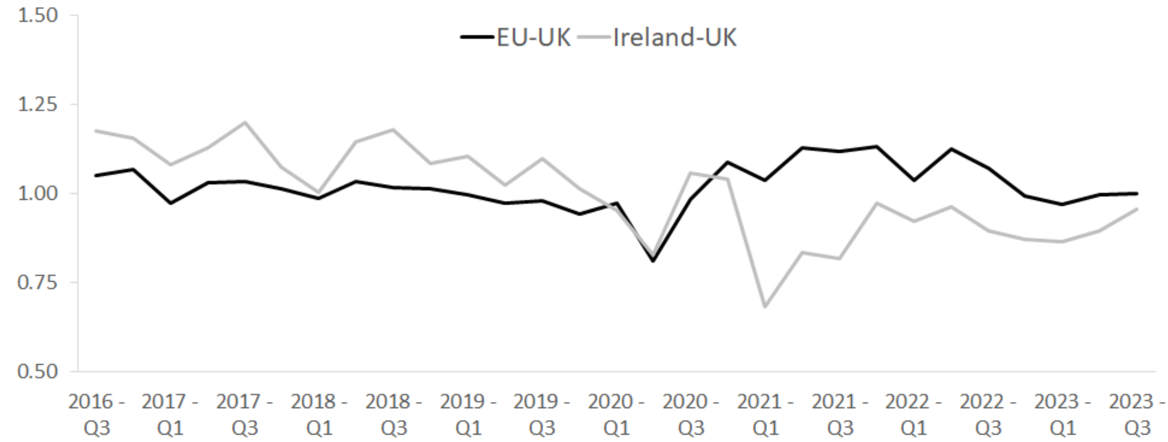

(a) Outward Flows to UK Ports

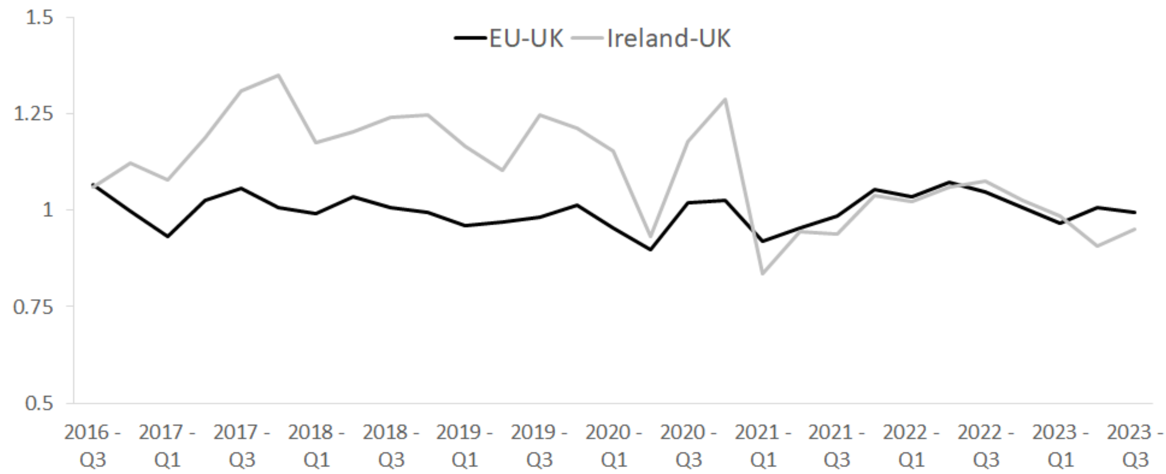

(b) Inward Flows from UK Ports

Source: Eurostat gross weight of goods transported to/from main ports. Values in thousand tonnes, indexed to 1 over 2015q3-2016q2 average.

As can be seen from both charts in Figure 4, Irish inward and outward flows remained substantially above their pre-vote volumes (and the comparable EU flow series) for the majority of the period between the Brexit vote and the introduction of UK trade tariffs in 2021q1. Neither Irish series show a material decline in the aftermath of the Brexit vote, in contrast to EU inflows and outflows, both of which show a marked drop in volumes by 2017q1. However, the effect of the UK's exit from the EU single market (2021q1) on Irish shipping flows is dramatic: quarter-on-quarter outwards flow volumes decline by 34%, while inward flow volumes contract by 35%. While inwards flows return to pre Brexit-vote levels by 2022, outwards flows do not return to either their pre Brexit-vote or pre-Brexit levels over the remaining sample period.

Again, these data reinforce the narrative driving our hypothesis: Irish firms retained their trading relationships with the UK, to a greater extent than other EU firms, due to the belief that alternative supply chains involving Northern Ireland could be developed. As EU firms sourced alternative markets, Irish-UK export and import flows increased, leading to record shipping volumes between Great Britain and Ireland. Once Brexit occurred, and trade tariffs with Great Britain were introduced, Irish firms adjusted supply chains, re-routing merchandise trade from Great Britain to Northern Ireland.[8]

3 Methodology

In keeping with the work of Born et al. (2019), Siedschlag and Koecklin (2019), and Breinlich et al. (2020) on the macroeconomic effects of Brexit, we treat the EU referendum vote of June 2015 as a unique natural experiment, with causality unrelated

[8] While some of this shipping decline is due to the reduction in use of the UK as a land-bridge route (i.e. transiting goods via the UK to alternate destinations), it cannot reflect the full effects observed in the data. Estimates by Lawless and Morgenroth (2017) put the maximum volume of Irish trade flows that use the UK as a land-bridge at 25% of total shipping volumes, well below the $\sim 35\%$ decline in shipping volumes post-Brexit.

to macroeconomic conditions, a precise date of the "treatment", and effects that are observable at the aggregate macroeconomic level. Additionally, we assume that the Irish border guarantees and negotiations were also part of this natural experiment, as (i) interest in the topic arose only after the referendum, (ii) the decision to include the Irish border as a central negotiation concern was not connected to economic performance, and (iii) the effects of the "no border" commitment in maintaining the open border caused changes at an aggregate macroeconomic level.

To quantify the effect of the "no border" commitment on Irish merchandise trade with the UK, we need to generate a counterfactual benchmark as our appropriate comparison economy. As our research focus is on the dynamic effects of this commitment on Irish-UK trade, we require our counterfactual economy to track the actual trade series as closely as possible prior to the result of the Brexit referendum. At the same time, it must be left unaffected by the Irish border negotiations, but impacted by the Brexit vote.

We follow Born et al. (2019) and use synthetic control methods to generate counterfactual import and export series between Ireland and the UK. The SCM is based on the idea that, when the units of observation are a small number of aggregate entities, a combination of unaffected units often provides a more appropriate comparison than any single unaffected unit alone. The synthetic control methodology formalizes the selection of the comparison units using a data driven procedure. See Appendix II for a more detailed discussion of synthetic control methods.

3.1 The Augmented Synthetic Control Method

As discussed by Abadie and Gardeazabal (2003), the Synthetic Control Method is recommended only in instances where the synthetic control's pre-treatment outcomes are near-exact matches with the pre-treatment outcomes for the treated unit. When it is not possible to construct a synthetic control that accurately fits pre-treatment outcomes, use of the SCM can lead to corner solutions with respect to the estimation of γ, and spurious outcomes for the post-treatment results.

The Augmented Synthetic Control Method (ASCM) of Ben-Michael, Feller, and Rothstein (2021) addresses, and corrects for, this specific issue. Analogous to bias correction for inexact matching, the approach takes the SCM values, estimates the bias due to imperfect pre-treatment fit using an outcome model, and then uses this estimate to de-bias the original SCM values. Under this approach, the ASCM estimator for $Y_{1T}(0)$ is:

$$\begin{aligned} \hat{Y}_{1T}^{aug}(0) &= \textstyle\sum_{W_i=0} \hat{\gamma}_i^{scm} Y_{iT} + \left(\hat{m}_{1T} - \textstyle\sum_{W_i=0} \hat{\gamma}_i^{scm} \hat{m}_{iT}\right) \\ &= \hat{m}_{1T} + \textstyle\sum_{W_i=0} \hat{\gamma}_i^{scm} (Y_{iT} - \hat{m}_{iT}) \end{aligned} \tag{1}$$

where $\hat{\gamma}_i^{scm}$ are the weights from the original SCM estimate and $\hat{m}_{iT}$ is an estimator for the post-treatment control potential outcomes $Y_{iT}(0)$.

There are two identified benefits to using the ASCM estimator over the standard SCM estimator. First, the conventional SCM estimate $\sum_{W_i=0} \hat{\gamma}_i^{scm} Y_{iT}$ is corrected by the imbalance in a function of the pretreatment outcomes $\hat{m}(\cdot)$. As $\hat{m}$ estimates the post-treatment outcome, it can be viewed as being an analogue of the bias correction method for inexact matching (Abadie and Imbens, 2011). As $\hat{m}(\cdot)$ represents the estimated bias component, the SCM and ASCM estimators will converge as $\hat{m}(\cdot) \rightarrow 0$. Second, the ASCM estimator is analogous to standard doubly robust (DR) estimation (Robins, Rotnitzky, and Zhao, 1994).

The improvement in performance over the standard SCM approach is dependent on the choice of estimator, and the nature of the relationship between pre-treatment outcomes and comparison units $\hat{m}$. If the estimator is a function of pre-treatment outcomes, e.g. $\hat{m}(\boldsymbol{X}) = \hat{\eta}_0 + \hat{\eta}_x \boldsymbol{X} + \hat{\eta_z \boldsymbol{Z}}$, the augmented estimator can be represented as

$$\hat{Y}_{iT}(0) = \sum_{i \neq 1} \hat{\gamma}_i^{scm} Y_{iT} + \sum_{t=1} \hat{\eta}_{xt} \left(X_{1t} - \sum_{i \neq 1} \hat{\gamma}_i^{scm} X_{it} \right) + \sum_{t=1} \hat{\eta}_{zt} \left(Z_{1t} - \sum_{i \neq 1} \hat{\gamma}_i^{scm} Z_{it} \right) \quad (2)$$

and pre-treatment periods will have larger (in absolute value) regression coefficients if they are more predictive of the post-treatment outcome; imbalances in these periods will cause a larger adjustment. Similarly, in an outcome model that is a combination of comparison units, e.g. $\hat{m}(\boldsymbol{X}) = \sum_{i \neq 1} \hat{\alpha}_i(\boldsymbol{X}) Y_{iT}$, the augmented estimator will be a weighting estimator that adjusts the weights to

$$\begin{aligned} \hat{Y}_{1t}(0) &= \sum_{i \neq 1} \left(\gamma_i^{scm} + \gamma_i^{adj} \right) Y_{it} \\ \gamma_i^{adj} &\equiv \hat{\alpha}(\boldsymbol{X}_1) - \sum_{j \neq 1} \gamma_j^{scm} \hat{\alpha}(\boldsymbol{X}_j) \end{aligned} \quad (3)$$

where γ_i^{adj} is the imbalance in a unit i-specific transformation of the lagged outcomes, which is itself dependent on the weighting function $\alpha(\cdot)$.

3.1.1 Choice of ASCM Estimator

Our choice of estimator for the ASCM is a diffusion-regression Bayesian state-space model. State-space models are systems, defined as a set of input, output and state variables, related by first-order difference equations. State variables are variables whose values evolve over time in a way that depends on the values they have at any given time, and on the externally imposed values of input variables. Output variables' values depend on the values of the state variables.

We can represent our state-space system as a pair of equations,

$$y_t = Z_t' \phi_t + \upsilon_t \tag{4a}$$

$$\phi_{t+1} = \Gamma_t \phi_t = R_t \nu_t \tag{4b}$$

where $\upsilon_t \sim \mathcal{N}(0, \sigma_{\upsilon t}^2)$ and $\upsilon_t \sim \mathcal{N}(0, Q_t)$ are independent of all other unknowns. Equation (4a) is the observation equation, linking the observed data, y_t, to a latent n-dimensional state vector, $\phi - t$. Equation (4b) is the state equation, governing the evolution of the state vector, ϕ_t, through time. In our model, y_t is a scalar observation, Z_t is a d-dimensional output vector, Γ_t is a $d \times d$ transition matrix, R_t is a $n \times m$ control matrix, υ_t is a scalar observation error with variance σ_t^2, and ν_t is a q-dimensional system error with a $q \times q$ state-diffusion matrix Q_t , where $q \geq d$.

There are two main components to the above system of equations: a local linear trend, and contemporaneous covariates with static coefficients. The local linear trend component of the model can also be represented as a set of equations

$$\begin{aligned} \mu_{t+1} &= \mu_t + \delta_t + \epsilon_{\mu t} \\ \delta_{t+1} &= \delta_t + \epsilon_{\delta t} \end{aligned} \tag{5}$$

where $\epsilon_{\mu t} \sim \mathcal{N}(0, \sigma_{\mu t}^2)$ and $\epsilon_{\delta t} \sim \mathcal{N}(0, \sigma_{\delta t}^2)$. The local linear trend model is a popular choice for modelling trends, as it quickly adapts to local variation, which is desirable when making short-term predictions.

The static coefficients covariates are critical to obtaining accurate counterfactual predictions, since they account for variance components that are shared by the series, including the effects of unobserved causes otherwise unaccounted for by the model.

We include these control series in the model through a static linear regression, where $Z_t = \beta' \boldsymbol{X}_t$ and $\alpha_t = 1$.

3.1.2 Prior Selection and Posterior Estimation

Let θ generically denote the set of all model parameters, and let $\phi = (\phi_1, ..., \phi_T)$ denote the full state sequence. We adopt a Bayesian approach to inference by specifying a prior distribution $p(\theta)$ on the model parameters, as well as a distribution $p(\phi_0|\theta)$ on the initial state values. We may then sample from $p(\phi_t, \theta|y)$ using MCMC methods.

Our model depends entirely on a set of variance parameters that govern the diffusion of the individual state components. We choose to model the variance terms as

$$\sigma_i^2 \sim \mathcal{G}(v, s) \quad \forall \ i \in \{\delta, \epsilon, \mu\} \tag{6}$$

where $v = 10^{-2}$, $s = 10^{-2}\rho_y$ and $\rho_y = \sum_t \frac{(y_t - \bar{y})^2}{(n-1)}$. Posterior inference in the model can be broken down into three components:

I Simulate draws of the model parameters, θ, and the state vector, α, given the observed data, $\{y_1, \ldots, y_{(T_0-1)}\}$.

II Set the posterior simulations to draw from the posterior predictive distribution of our outcome variable, $p(y_{T_0}, \ldots, y_T | y_1 \ldots y_{(T_0-1)})$, generating the counter-factual series $\{\hat{y}_{T_0}, \ldots, \hat{y}_T\}$ from the pre-intervention series $\{y_1, \ldots, y_{(T_0-1)}\}$.

III Use the posterior predictive samples to compute the posterior distribution of the pointwise impact $\{y_t - \hat{y}_t \ \ \forall t = 1, \ldots, T\}$.

4 Data

As we wish to examine whether there has been a Brexit-related structural change in any component of the Irish trade relationship with the UK (beyond the typical effect observed in other European economies) we make the decision to separately examine gross bilateral imports from the UK, and gross bilateral exports to the UK, in our analysis. Due to model complexity, and the increased accuracy of pre-treatment fit with sample-size, we require a substantial time dimension to our data. Fortunately, there are a number of international organizations that maintain trade data, back to the early 1990s, for a large panel of advanced and European economies.

Additionally, we need to incorporate a careful balance of auxiliary covariates to be used in parallel to the lagged outcome variables. Too few, and we risk a sub-optimal level of pre-treatment fit / covariance balance to our model. Too many, and we risk increased extrapolation bias for covariate combinations that lie outside the convex hull of the control units.

Furthermore, as we are attempting to identify the effects of a specific economic change (i.e. Brexit), there is a risk that extending the dataset to include other large structural economic developments could bias both the pre- and post-treatment outcomes. If there were asymmetric trade responses to such shocks, these effects could effectively pollute the ASCM estimates of the pre-treatment fit (if the event happens before the Brexit vote in June 2016), or attribute causal effects to the treatment when they derive from an alternative source (if the event happens after June 2016). To prevent structural biases prior to entry into force of the European Single Market, we start our sample in 1995. Similarly, to prevent heterogeneous responses to the Covid-19 crisis from being incorrectly captured as a Brexit effect, we end our sample prior to 2020.

We estimate our ASCM models on a panel of quarterly data over the 1995q1 - 2019q4 period, using a cross-sectional donor country sample limited to the other 18 member states of the euro area. We choose to focus on the euro area as our donor pool, given that they are part of the EU common market, to whom EU common trade policy applies, and the shared currency prevents bias in both pre-treatment fit and post-treatment outcome estimates due to heterogeneous monetary policy or exchange rate effects.

Our bilateral trade data comes from the IMF Direction of Trade Statistics (DOTS) series. We aggregate the data from monthly to quarterly frequency, to match with the lower frequency macroeconomic data that we use as auxiliary covariates. In keeping with other macroeconomic literature that implements the SCM approach, our set of covariates includes: (i) output, (ii) government spending, (iii) consumption, (iv) investment, (v) employment, (vi) and inflation. Data for output, government spending, consumption, and investment come from individual member states' quarterly national accounts (QNA) statistics series, while employment and inflation data come from Eurostat (to ensure a harmonized approach to measurement). All data are seasonally adjusted using the X-13ARIMA-SEATS program, implemented in the U.S. Census Bureau's software package.

5 Estimation Results

Our baseline specification of the model has all data entering the system in log-level form, with output, government spending, consumption and investment all in real 2019 values. However, we also conduct robustness checks with alternative forms of the data, including indices and (for euro-denominated values) as a percentage share of output. We also trial versions of the model with and without unit fixed effects, with and without pre-and post-treatment outcomes residualized against the auxiliary covariates, and with and without bootstrapped standard errors.

The general form of our augmented Synthetic Control Model has the representation:

$$\min_{\gamma \in \Delta^{N_0}} \quad \theta_x \parallel \boldsymbol{X}_1 - \boldsymbol{X}_0 \cdot \boldsymbol{\gamma} \parallel_2^2 + \theta_z \parallel \boldsymbol{Z}_1 - \boldsymbol{Z}_0 \cdot \boldsymbol{\gamma} \parallel_2^2 + \zeta \sum_{W_i=0} f(\gamma_i) \tag{7}$$

where $\boldsymbol{X}_i$ is the trade series under analysis, $\boldsymbol{Z}_i$ represent the potential set of time-invariant auxiliary covariates, and $f(\gamma_i)$ is a penalty function that penalizes deviations from the SCM weights.

As a first test to confirm that the use of the ASCM estimator is appropriate, we run a version of the standard Synthetic Control Method similar to the specification of Born, Müller, Schularick, and Sedláček (2019), with weights constrained to be on the simplex. As anticipated, there are problems with this approach, as there appears to be no convex combination of donor countries in our sample that closely approximates the treated units, leading to poor pre-treatment fit.[9]

5.1 Baseline Specifications

The optimal weighting matrices, $\{\boldsymbol{W}_x^*, \boldsymbol{W}_m^*\}$, for our bilateral export and import synthetic controls are presented in Table 2. As expected, the ASCM estimator gives a sparse weighting structure for each estimation, but sufficiently populated to mitigate potential concerns regarding the efficiency of the method to simulate an accurate doppelganger for Y_{1t}^N in both series. The synthetic control for exports (imports) is fit with only seven (six) countries from the donor pool, with approximately 85% (95%) of the fit coming from just four countries.

[9] This result is likely due to the UK's over-weighted importance in Ireland's trade data, as a share of imports and exports. In several cases, the results that were obtained constructed the doppelganger Ireland from a single member state in the donor pool.

Table 2: ASCM Weighting Structures, Baseline Specifications

	Exports	Imports
Finland	**32.4%**	**43.5%**
Belgium	**13.3%**	**27.6%**
France	**27.5%**	**11.3%**
Lithuania	**11.7%**	**11.6%**
Luxembourg	**6.6%**	**5.6%**
Portugal	**8.2%**	**0.4%**
Cyprus	**0.4%**	**0.0%**

Source: Authors' calculations using IMF DOTS, EU Quarterly National Accounts, and Eurostat data. Data series seasonally adjusted using X-13ARIMA-SEATS Seasonal Adjustment Program.

Figure 5 presents the results of the "gap plot" for gross bilateral exports: the difference between Ireland and its synthetic control doppelganger using a Bayesian causal impact estimator. Confidence intervals are calculated using a conformal inference procedure.[10] As can be seen from the plot, there is a considerable gap between the simulated export estimates and the realised data.

From the Brexit event until the end of the sample, our ASCM estimator suggests that the counterfactual Irish exports to the UK would have been lower than the realised value of exports in all 14 quarters of the validation period, with significant differences (at the 95% confidence level) in the point estimates for 10 of the 14 quarters. Furthermore, prior to the Brexit event, there were only 9 quarters (10.5% of the pre-treatment sample) in which the difference between the logged value of exports was above 0.1: in contrast, 11 of the 14 post-treatment quarters show the gap to be greater than 0.1.

[10] See Chernozhukov, Wüthrich and Zhu (2021) for a discussion of conformal confidence intervals.

Figure 5: ASCM Gap Plot for Counterfactual Bilateral Irish Exports to the UK

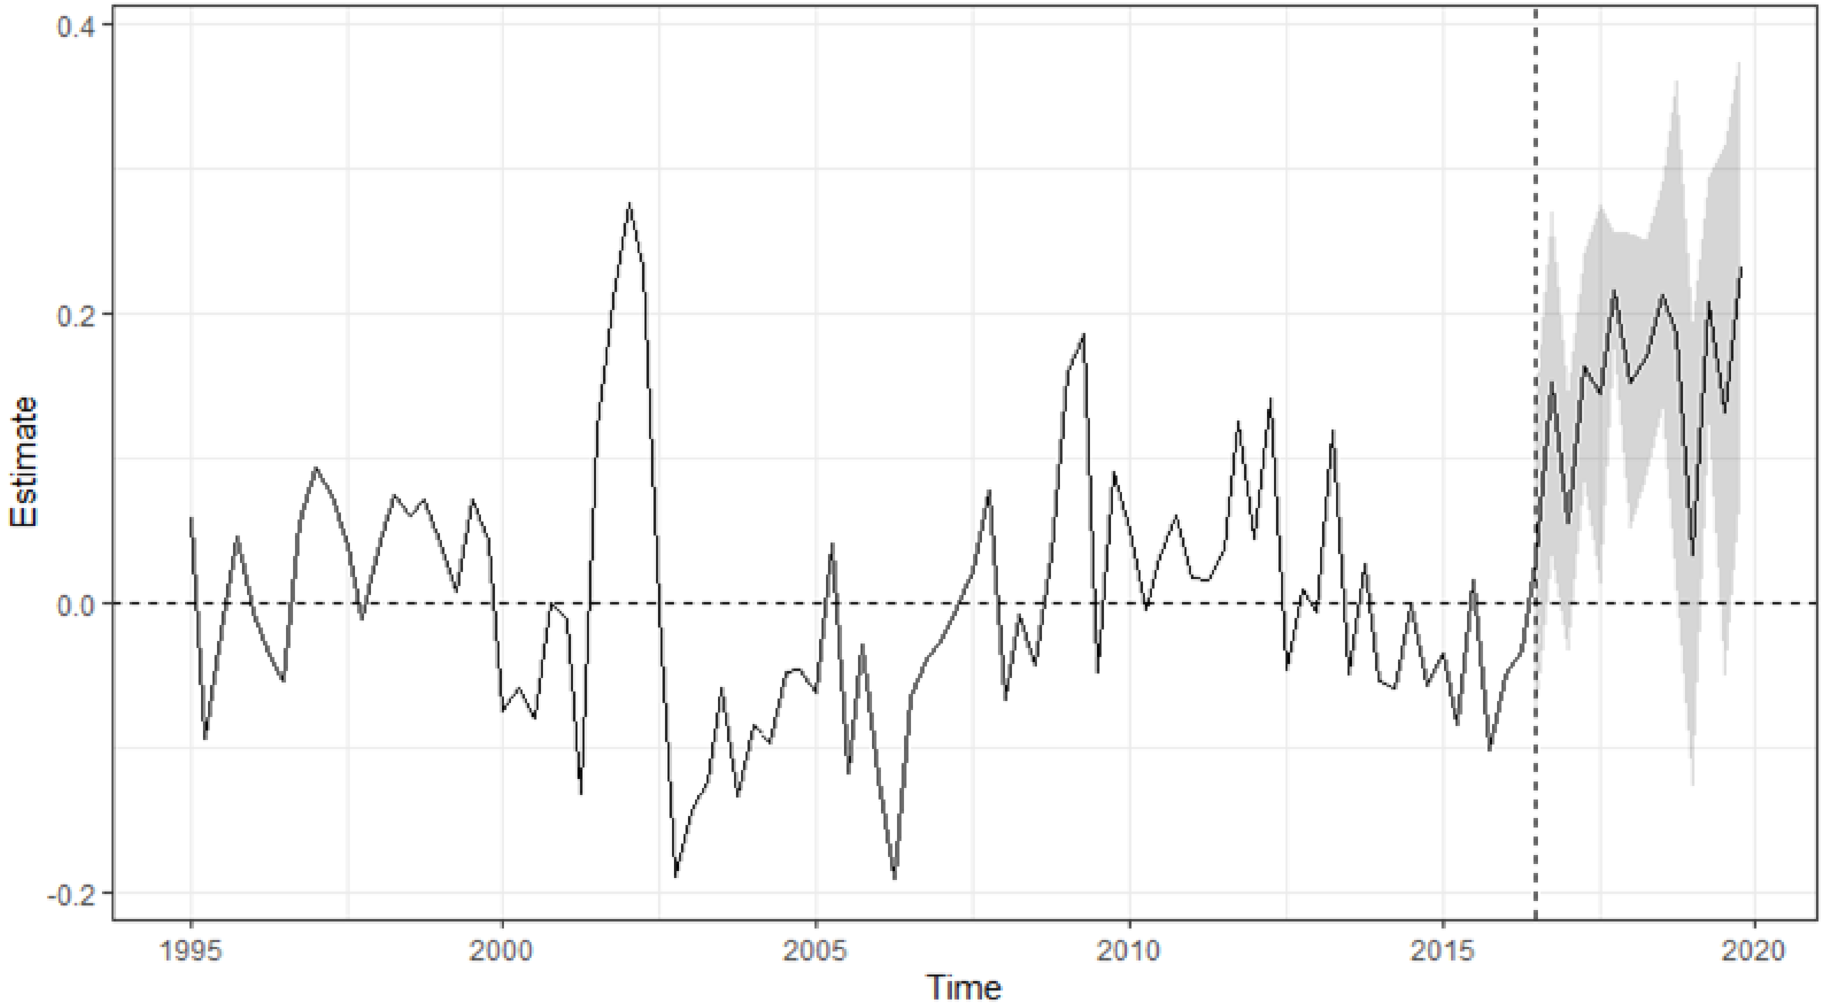

Note: Point estimates along with point-wise 95% conformal confidence intervals for the effect of Brexit on bilateral trade. Augmented SCM estimates derived using a Bayesian Causal Impact estimator and additional macroeconomic covariates.

Figure 6: ASCM Estimates for Counterfactual Bilateral Irish Exports to the UK

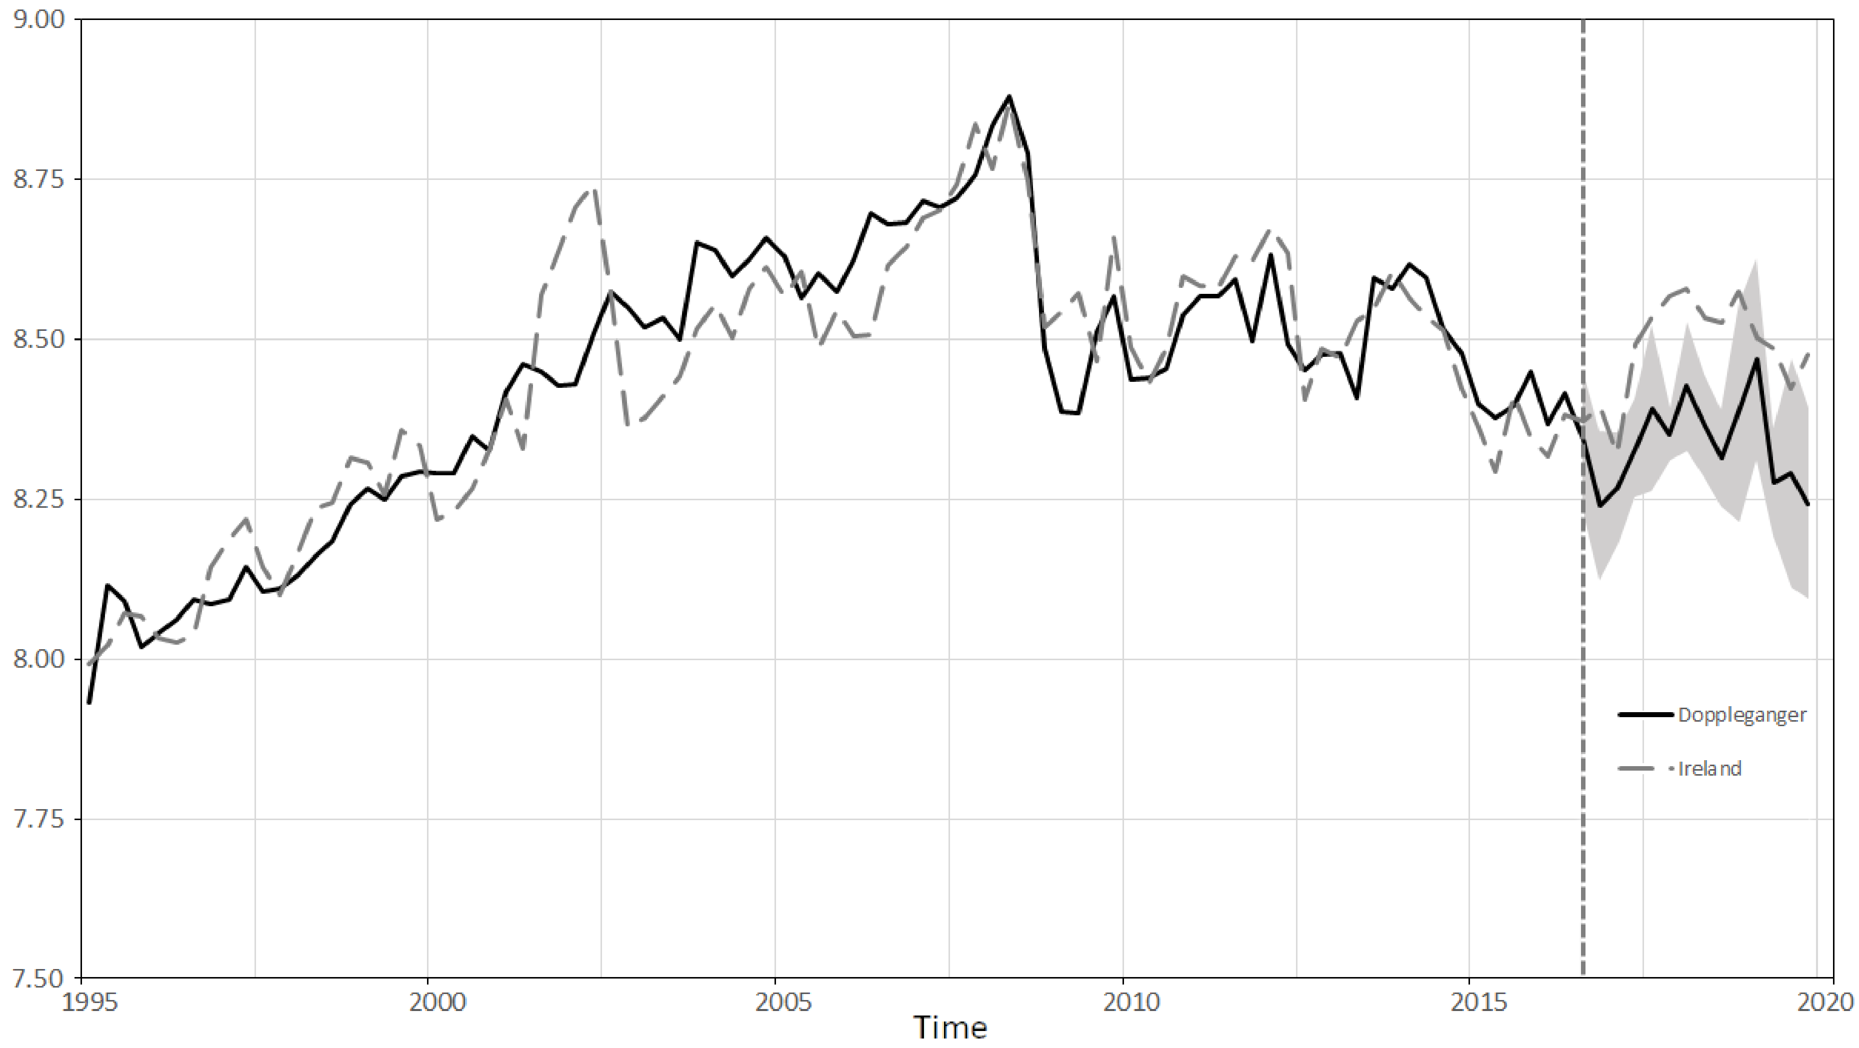

Note: Point estimates along with point-wise 95% conformal confidence intervals for the effect of Brexit on bilateral trade. Augmented SCM estimates derived using a Bayesian Causal Impact estimator and additional macroeconomic covariates.

Figure 6 shows log-level exports to the UK for both Ireland and its synthetic control using the ASCM approach. As seen from the gap plot in Figure 5, the estimator suggests that exports from the doppelganger should have been lower than the realised data for all points in the post-Brexit validation period. Additionally, the model also suggests that exports from the doppelganger should have been lower than the four-quarter average value of exports immediately prior to the Brexit vote, in 10 of the 14 periods, and by an average of 7.2% over the post-Brexit period. While the counterfactual results suggest that Brexit should have negatively affected Irish exports to the UK, the realised data shows that exports moved in the opposite direction; average export flows were more than 12% higher in the post-Brexit sample than the average of the four-quarter pre-vote period.

In (real) monetary terms, these results suggest that the announced "no hard border" negotiating stance adopted by the UK and EU translated into an average quarterly increase in Irish exports to the UK of €676 million. Irish exports during the post-Brexit vote period of our sample averaged €4.853 billion, while our counterfactual estimates suggest that average export flows of €4.177 billion would have been observed in our doppelganger Irish economy. Cumulatively, we estimate that the "no hard border" commitment supported an additional €9.465 billion of merchandise export flows over the 2016q3-2019q4 period. This represents an average quarterly differential of 16.3% between the true data and our synthetic counterfactual since the Brexit vote.

Gap plot results for gross bilateral imports are presented in Figure 7. Again, estimated levels of counterfactual imports are above realised values in all 14 quarters of the post-treatment sample. In the pre-treatment period there were never more than 8 continuous quarters in which the ASCM estimate was below the actual export data. Similarly, there are seven quarters in the post treatment sample (50%) where the estimated log difference between the realised and doppelganger values is above 0.15, and seven quarters in the pre-treatment period (8.1%) where the estimated difference is above 0.15.

Figure 8 shows the ASCM estimates and the realised values of log-level imports from the UK to Ireland. As with exports, the synthetic series present a markedly different directional trend to the realised data. Our model suggests that Irish imports should have declined in the aftermath of the Brexit vote, with average quarterly values 11.8% below the average value of imports in the four quarters prior to the Brexit vote. In contrast, the actual data series only recorded a transitory decline for the first two quarters of the post-Brexit period, before rebounding to levels above those observed in the previous four year. Over the full validation period of our sample, Irish imports from the UK were 4.7% above the four-quarter average just prior to the Brexit vote.

These effects of the "no hard border" commitments translate into an average quarterly increase in import flows of €604.2 million following the Brexit vote. Our synthetic results suggest that Irish quarterly imports from the UK would have been €3.98 billion in the absence of such commitments, while an average quarterly value of €4.585 billion was recorded in the actual import series for Ireland: an average quarterly differential of 15.5%. On a cumulative basis, this represents an incremental €8.459 billion of import flows from the UK to Ireland over the 2016q3-2019q4 period, which we believe would not have occurred had a hard Brexit been the objective of either party in the exit negotiations.

Overall, these results provide strong supporting evidence of our hypothesis, that the commitment to a "no hard border" solution underpinned bilateral trade between the UK and Ireland during the 2016q3-2019q4 period, significantly mitigating the trade effects of the Brexit vote relative to other EU countries. At the outset of the Brexit Withdrawal Agreement negotiations in 2017, both EU and UK negotiation teams made repeated references to the importance of maintaining an "open border" between Ireland and

Figure 7: ASCM Gap Plot for Counterfactual Bilateral Irish Imports from the UK

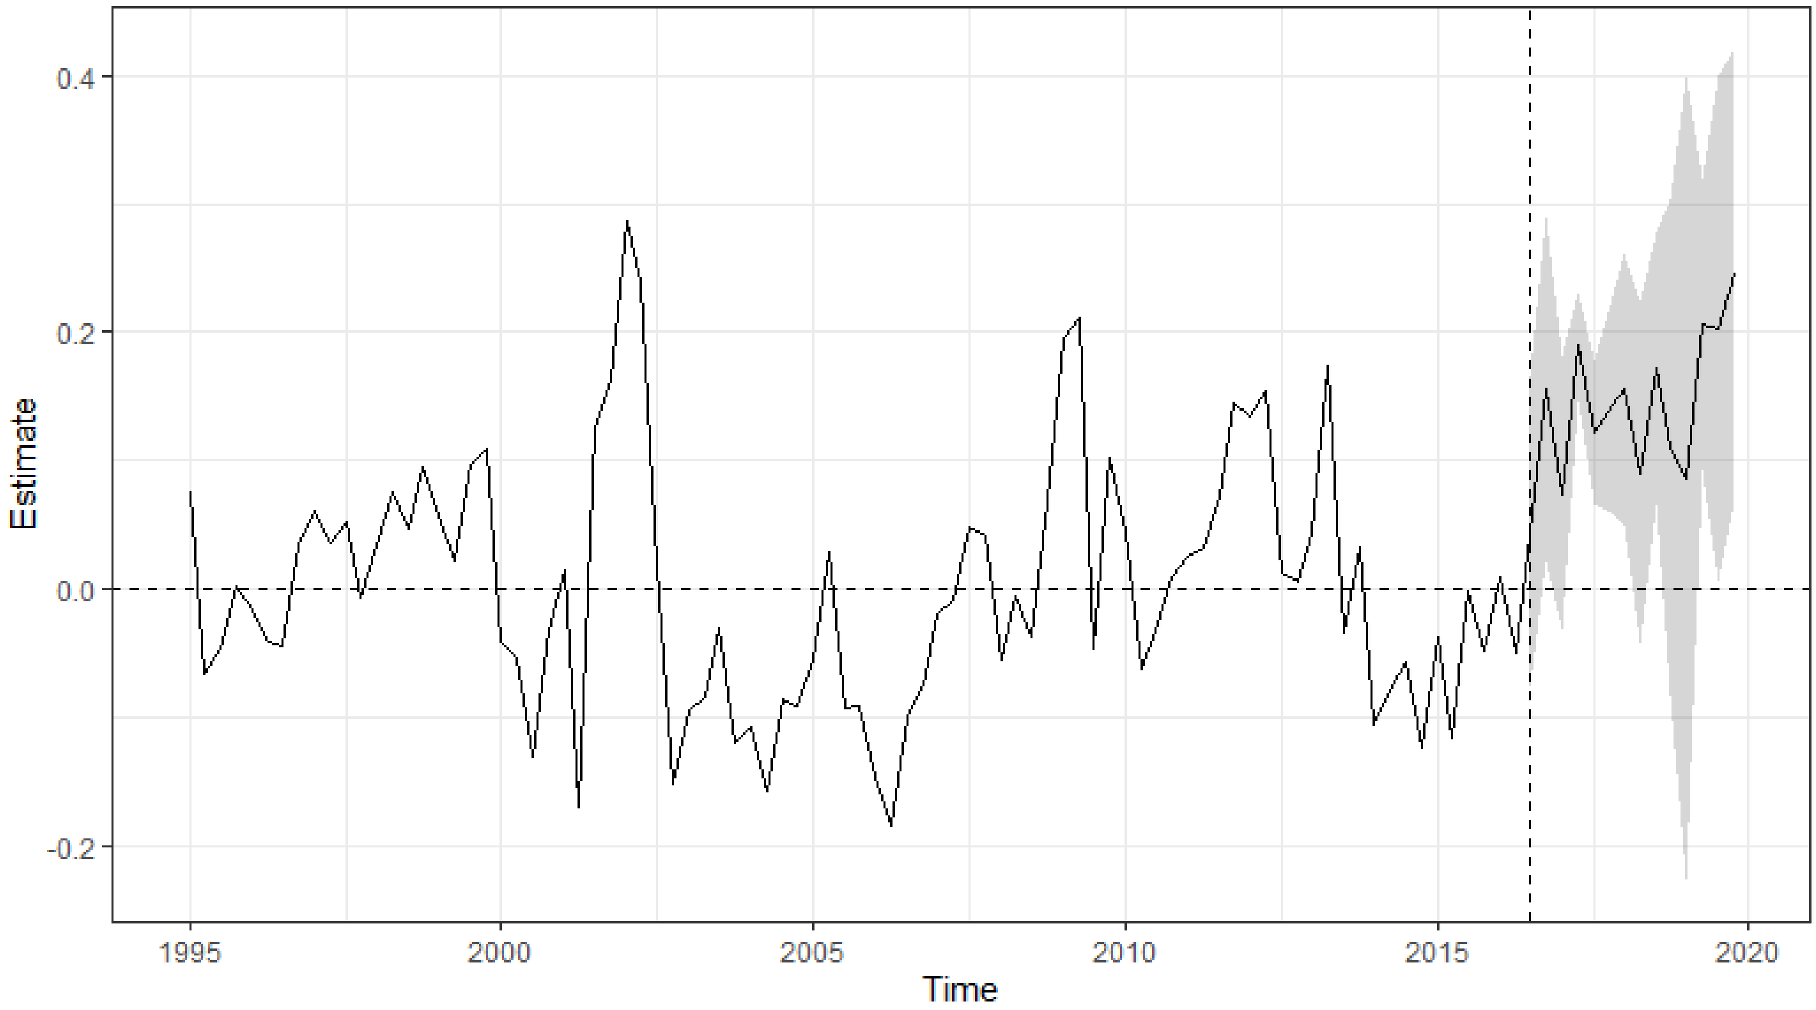

Note: Point estimates along with point-wise 95% conformal confidence intervals for the effect of Brexit on bilateral trade. Augmented SCM estimates derived using a Bayesian Causal Impact estimator and additional macroeconomic covariates.

Figure 8: ASCM Estimates for Counterfactual Bilateral Irish Imports from the UK

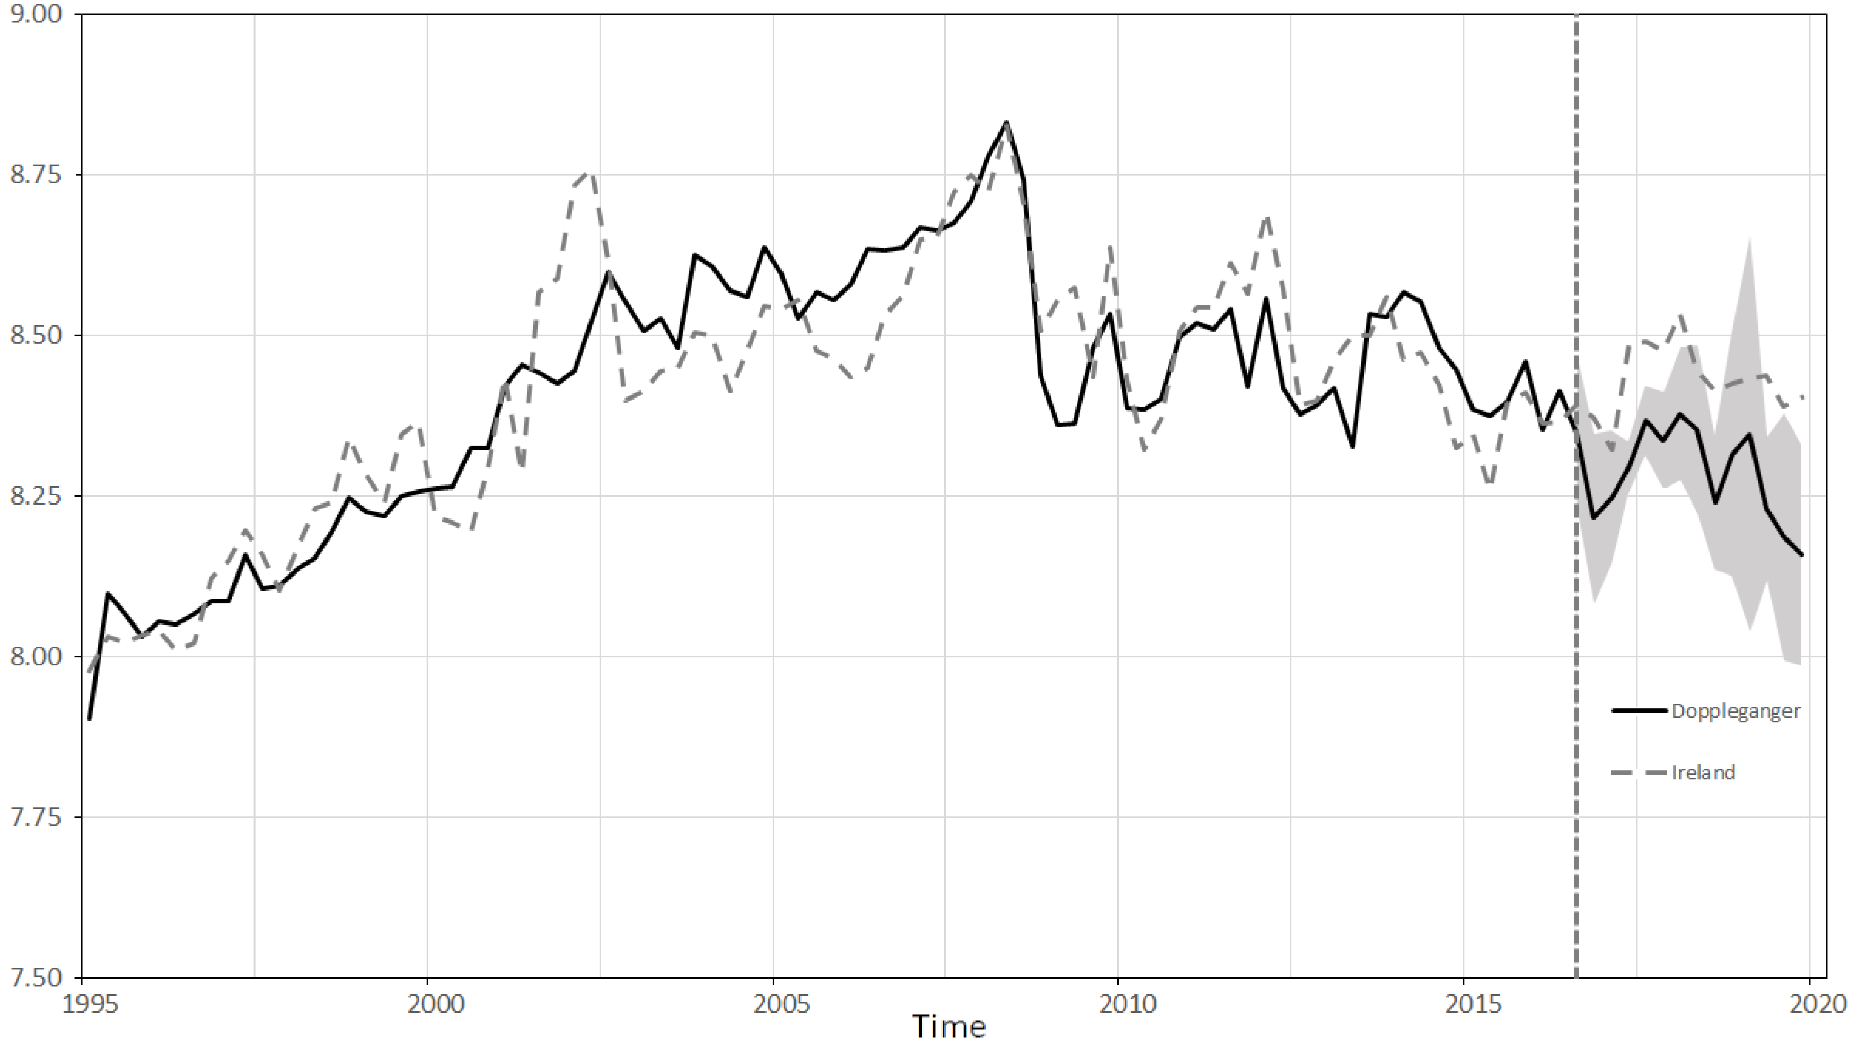

Note: Point estimates along with point-wise 95% conformal confidence intervals for the effect of Brexit on bilateral trade. Augmented SCM estimates derived using a Bayesian Causal Impact estimator and additional macroeconomic covariates.

Northern Ireland, while draft negotiating papers allowed for any permissible solution to be "unique" to the post-Brexit relationship between the UK and the EU.

Even prior to the formal Brexit negotiations, there were discussions between the British and Irish Governments in 2016 on proposals to avoid a hard border and potential technical solutions to trade issues. Given these efforts to maintain trade across the Irish border, it seems plausible that exporters and importers on both sides of the Irish border believed the short-run and long-term disruptions to trade linkages from the UK's withdrawal from the EU were less impactful than to the linkages with other EU countries, i.e. Brexit did not really mean Brexit.

Our analysis also shows the importance of separately examining the effects of Brexit on exports and imports, rather than just focusing on net export values. In aggregate terms, there is little evidence in the data to support the hypothesis that Ireland's trade response to Brexit differed from its European counterparts; there are very few periods of statistical difference between the realised Irish data and the synthetic doppelganger, when the bilateral net exports series is considered. However, gross positions tell a much different story, with the realised merchandise export and import data substantially above the corresponding counterfactual values estimated by our ASCM approach.

5.2 Robustness Checks

To test the credibility of the results from our baseline specification, we perform two robustness checks common to the synthetic control literature.

An important consideration in any synthetic control analysis is whether the gap in the outcome between the actual treated unit and the doppelganger is driven by pre-treatment effects. If pre-existing trends are present in the data before the timing of the treatment event, it may not be possible to infer causality from the event, and the SCM analysis will likely be spurious. Thus, we perform a placebo test on the date at which the Brexit event occurs in our sample.

Furthermore, it could be possible that our results are being driven by persistent policy differences in the treatment period, for some of the countries in our sample. If this is the case, and these countries have non-negligible weights in our $\boldsymbol{W}_x$ or $\boldsymbol{W}_m$ matrices, the synthetic control method may potentially be defining the absence of these policy responses in the validation period as Brexit outcomes. Thus, we re-run our earlier analysis for a constrained subset of our full donor pool.

5.2.1 Placebo Test for the Treatment date of the Border Commitment

As a first test of the validity of our results, we re-run our counterfactual analysis, changing the date at which the treatment effect of interest occurs in our sample. Analysis using an artificial treatment date should fail to reject the null hypothesis of in-time placebo effects prior to the event, but should be able to confirm the underlying post-treatment effect from the date of the intervention.

To conduct this placebo study, we re-estimate our doppelganger Ireland using the same ASCM approach outlined in the previous section, but with the date of the Brexit vote brought forward to 2015q1. This date represents a point after the September 2014 Conservative Party conference where the prospect of and EU membership referendum was first raised at a national level, but before the European Union Referendum Act was brought before parliament in May 2015. We use the same estimation approach to compute the augmented synthetic control, with the same set of variables across the new training and validation periods. The gap plot results for these placebo tests are presented in Figure 9 for Irish exports, and Figure 10 for Irish imports.

As shown in both charts, there is no significant divergence in the flow of imports or exports between our synthetic Ireland and actual Ireland, over the 2015q1 - 2016q2

Figure 9: Placebo Test Gap Plot for Counterfactual Bilateral Irish Exports to the UK

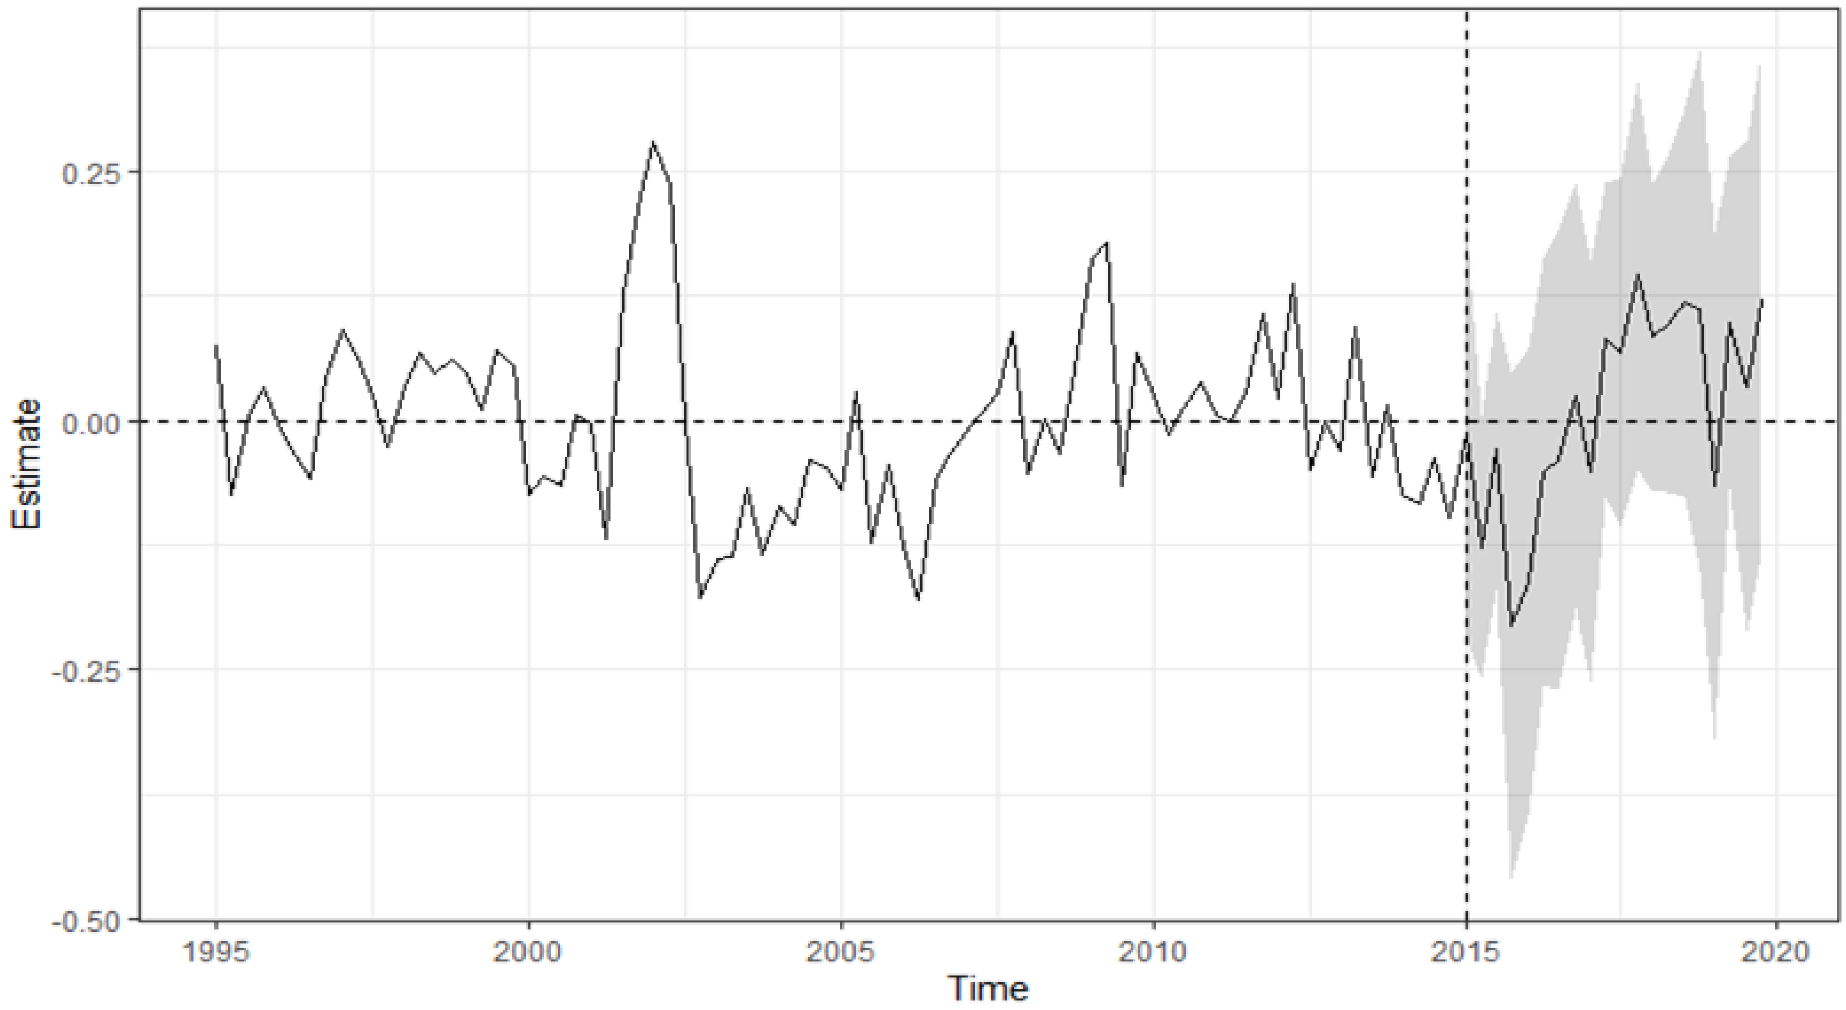

Note: Point estimates along with point-wise 95% conformal confidence intervals for a placebo Brexit effect on bilateral trade. Augmented SCM estimates derived using a Bayesian Causal Impact estimator and additional macroeconomic covariates.

Figure 10: Placebo Test Gap Plot for Counterfactual Bilateral Irish Imports from the UK

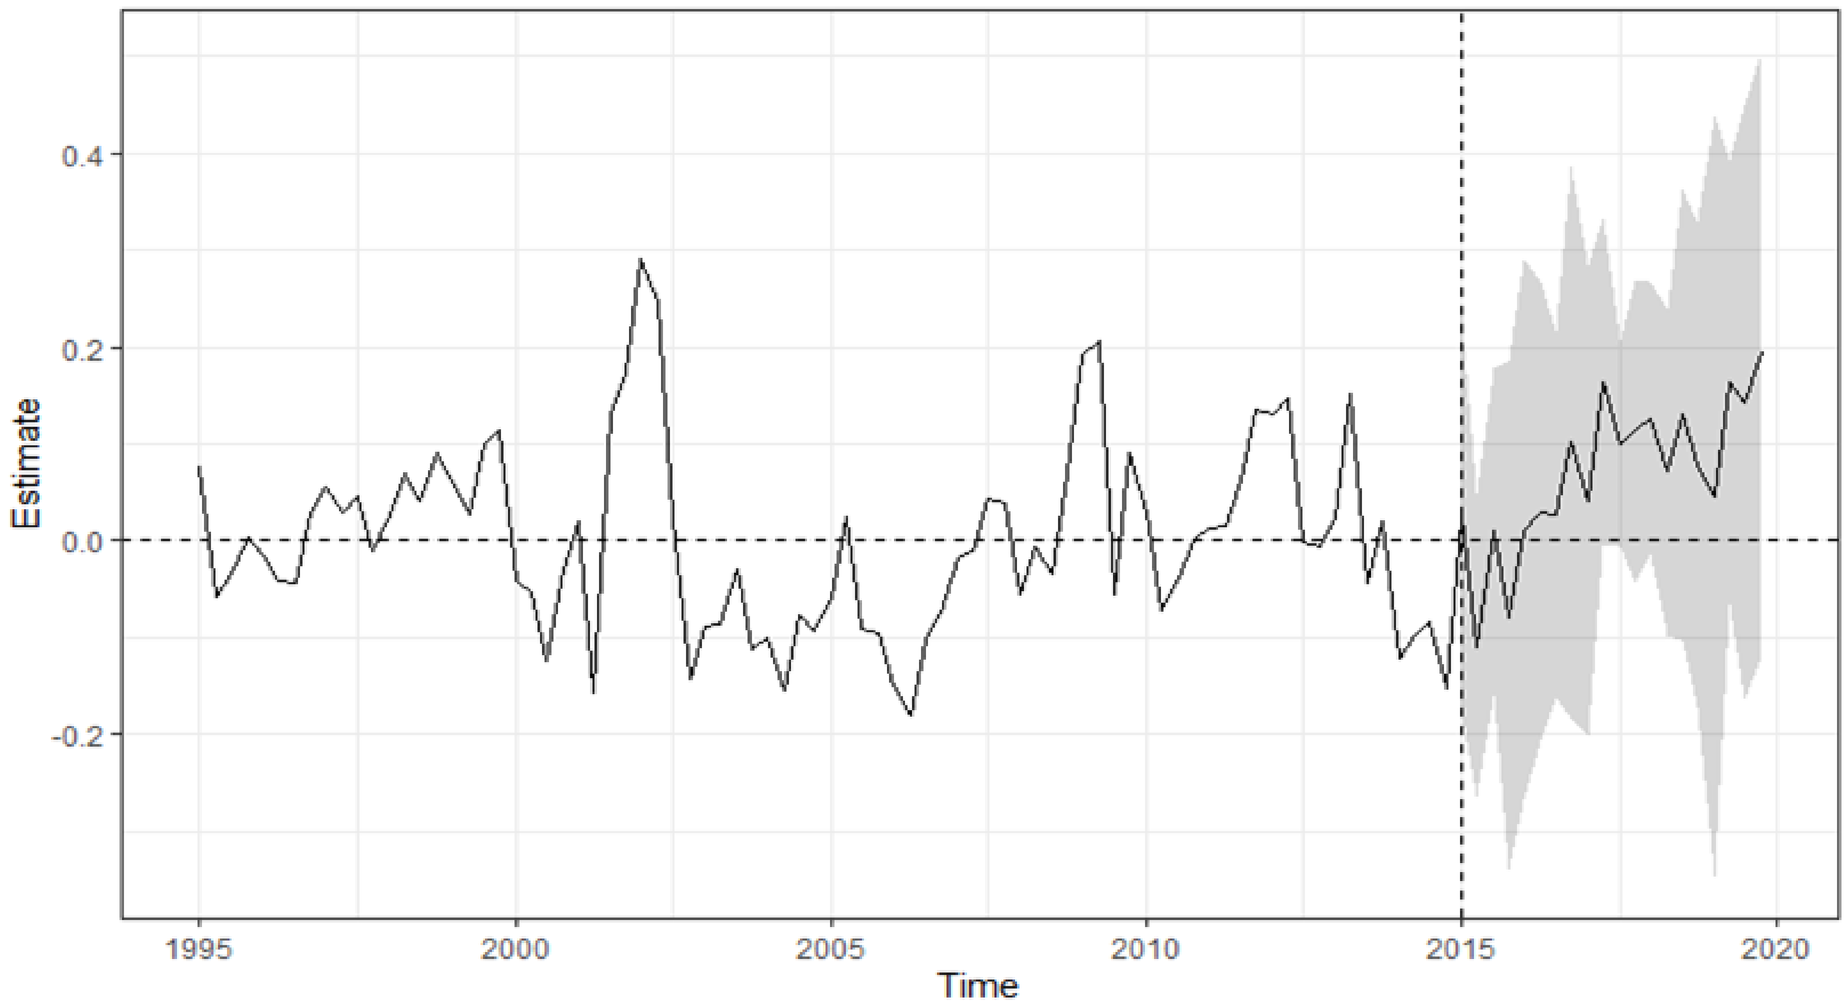

Note: Point estimates along with point-wise 95% conformal confidence intervals for a placebo Brexit effect on bilateral trade. Augmented SCM estimates derived using a Bayesian Causal Impact estimator and additional macroeconomic covariates.

period of our "in-time placebo" analysis. This is in contrast to the correctly dated 2016q3 timing of the "no hard border" commitments, which shows substantial effects on impact for both bilateral import and export flows. These results add supportive evidence to the argument that the gaps estimated in Figure 5 and Figure 7 reflect the effects of the "no hard border" commitment by both UK and EU officials (relative to other EU members states that were unaffected by the Irish border negotiations), and not a lack of predictive power from the augmented synthetic control method.

Additional placebo tests, extending the Brexit event back further in time and moving the Brexit event to the end of 2015, deliver similar gap plot results to those presented in Figure 9 and Figure 10.

5.2.2 Robustness Check on the Donor Pool Member States

As discussed above, an important condition of the synthetic control approach assumes that there are no heterogeneous policy shocks that may be driving differences in the synthetic counterfactual and the realised data. To minimize the risk of such shocks contaminating our estimates, we limited the donor pool to the sample of euro area member countries at the time of the Brexit vote event. However, it may be the case that for the set of countries in the sample that were not original euro area members, persistent differences in monetary or exchange rate policy in the pre-treatment period may be biasing estimates in the post-treatment period, where these differences no longer exist.

To confirm that our results are not driven by heterogeneous macroeconomic policy in the pre-intervention period, we restrict our donor pool to the sample of EA12 member states, dropping the seven countries that joined the euro area between 2007 and 2015. We re-estimate our ASCM approach, with alternative weighting matrices $\{\tilde{\boldsymbol{W}}_x^*, \tilde{\boldsymbol{W}}_m^*\}$, but keeping all other elements from our baseline specification fixed. Gap plot results for our counterfactual export and import series are presented in Figure 11 and Figure 12.

Overall, our results do not change materially from those obtained in the baseline specification. For the export series, the model now draws five countries from the reduced donor pool, with Finland now possessing the largest weighting share of 34%. While the absolute deviation between the realized data and the synthetic counterfactual series is (axiomatically) larger in the pre-treatment period, we still obtain nine post-treatment quarters where the difference is significant, and seven quarters where the deviations are larger than those observed in all but one quarter of the pre-treatment period. Over the full post-treatment period, these alternative estimates suggest that the cumulative difference between the actual data and the synthetic doppelganger translates to export values of €8.72 billion, or €622.9 million per quarter following the EU referendum vote.

For the import series estimated using the reduced donor pool, only four member countries are used to generate the doppelganger, with Finland retaining the largest weighting share (45.1%). As per the baseline model, the ASCM estimator indicates that imports would have been substantially lower under the synthetic counterfactual: positive, significant differences between both values are estimated in 9 of the 14 post-treatment periods. There are 5 quarters in the post treatment period where the estimated log-difference between the actual and counterfactual import series is above 0.15, but only 8 quarters where this is the case in the pre-treatment period. Converting the difference between both series into real 2019 euro values, the alternative estimates suggest that import values were €7.24 billion above the synthetic counterfactual over the full post-treatment period, equivalent to €517.1 million of additional imports each quarter following the Brexit vote.

Figure 11: Robustness Check Gap Plot for Counterfactual Bilateral Irish Exports to the UK

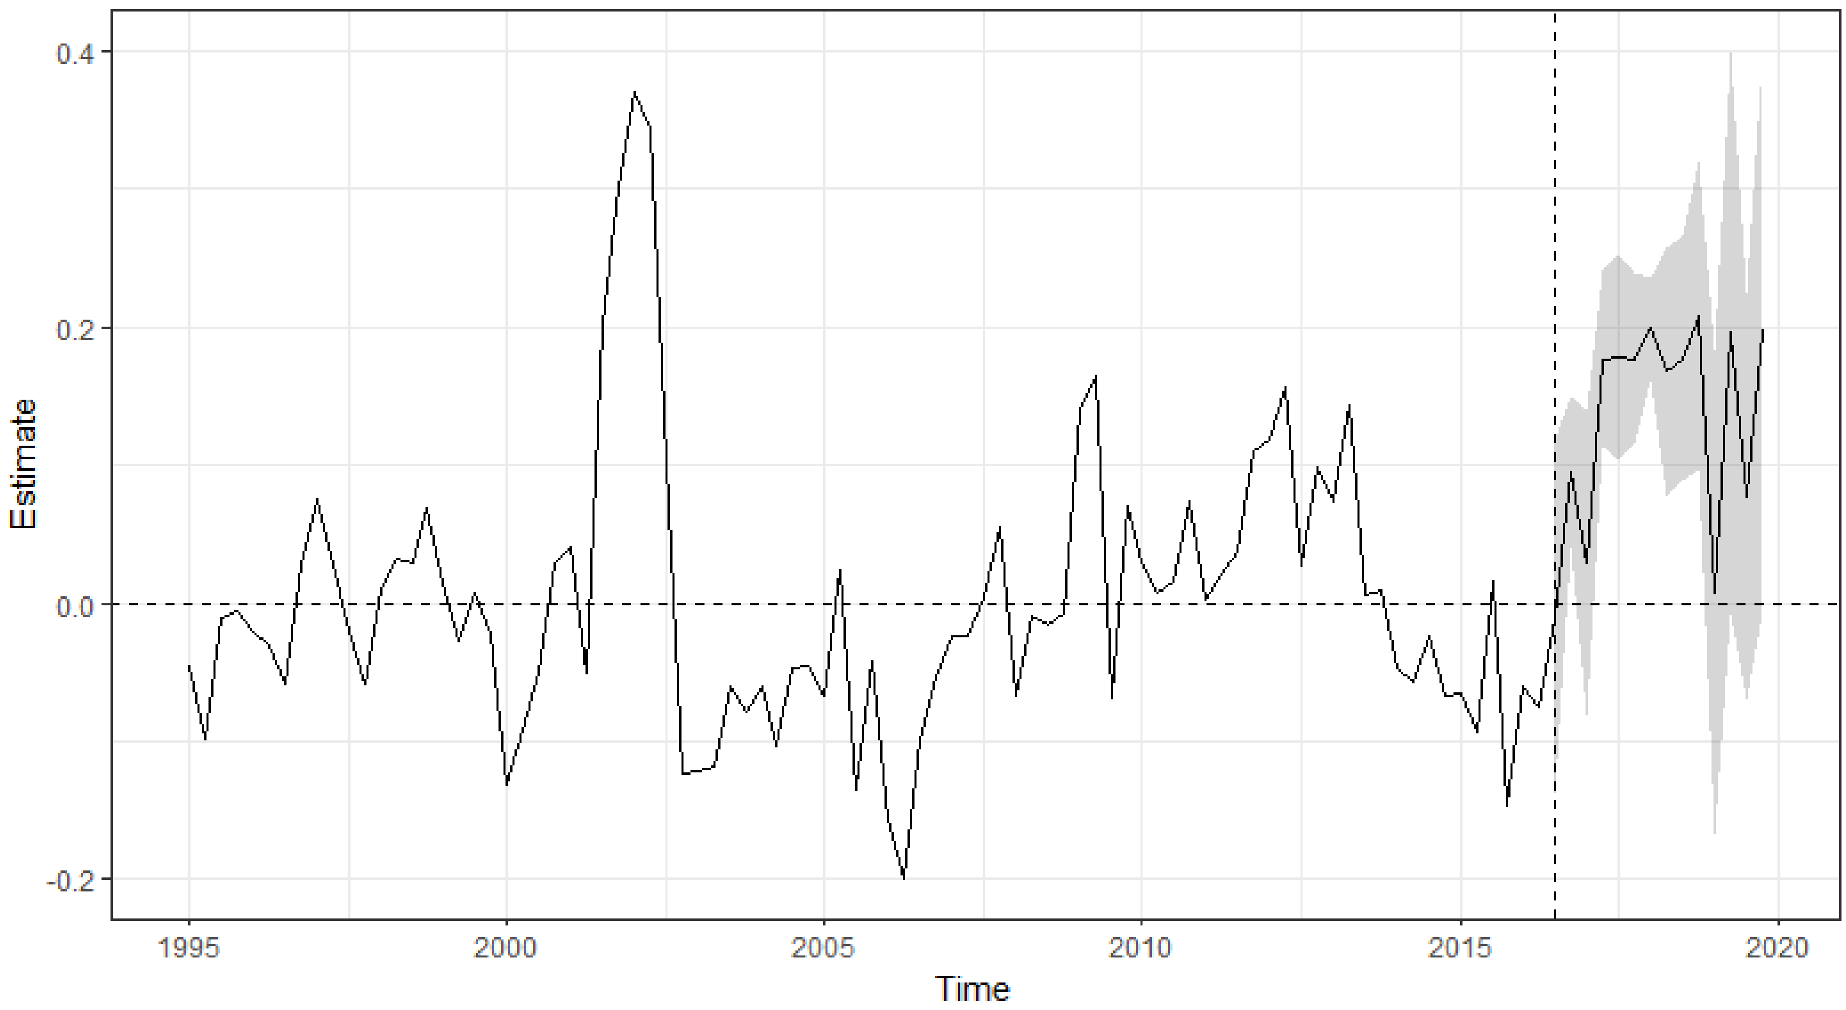

Note: Point estimates along with point-wise 95% conformal confidence intervals for a placebo Brexit effect on bilateral trade. Augmented SCM estimates derived using a Bayesian Causal Impact estimator and additional macroeconomic covariates.

Figure 12: Robustness Check Gap Plot for Counterfactual Bilateral Irish Imports from the UK

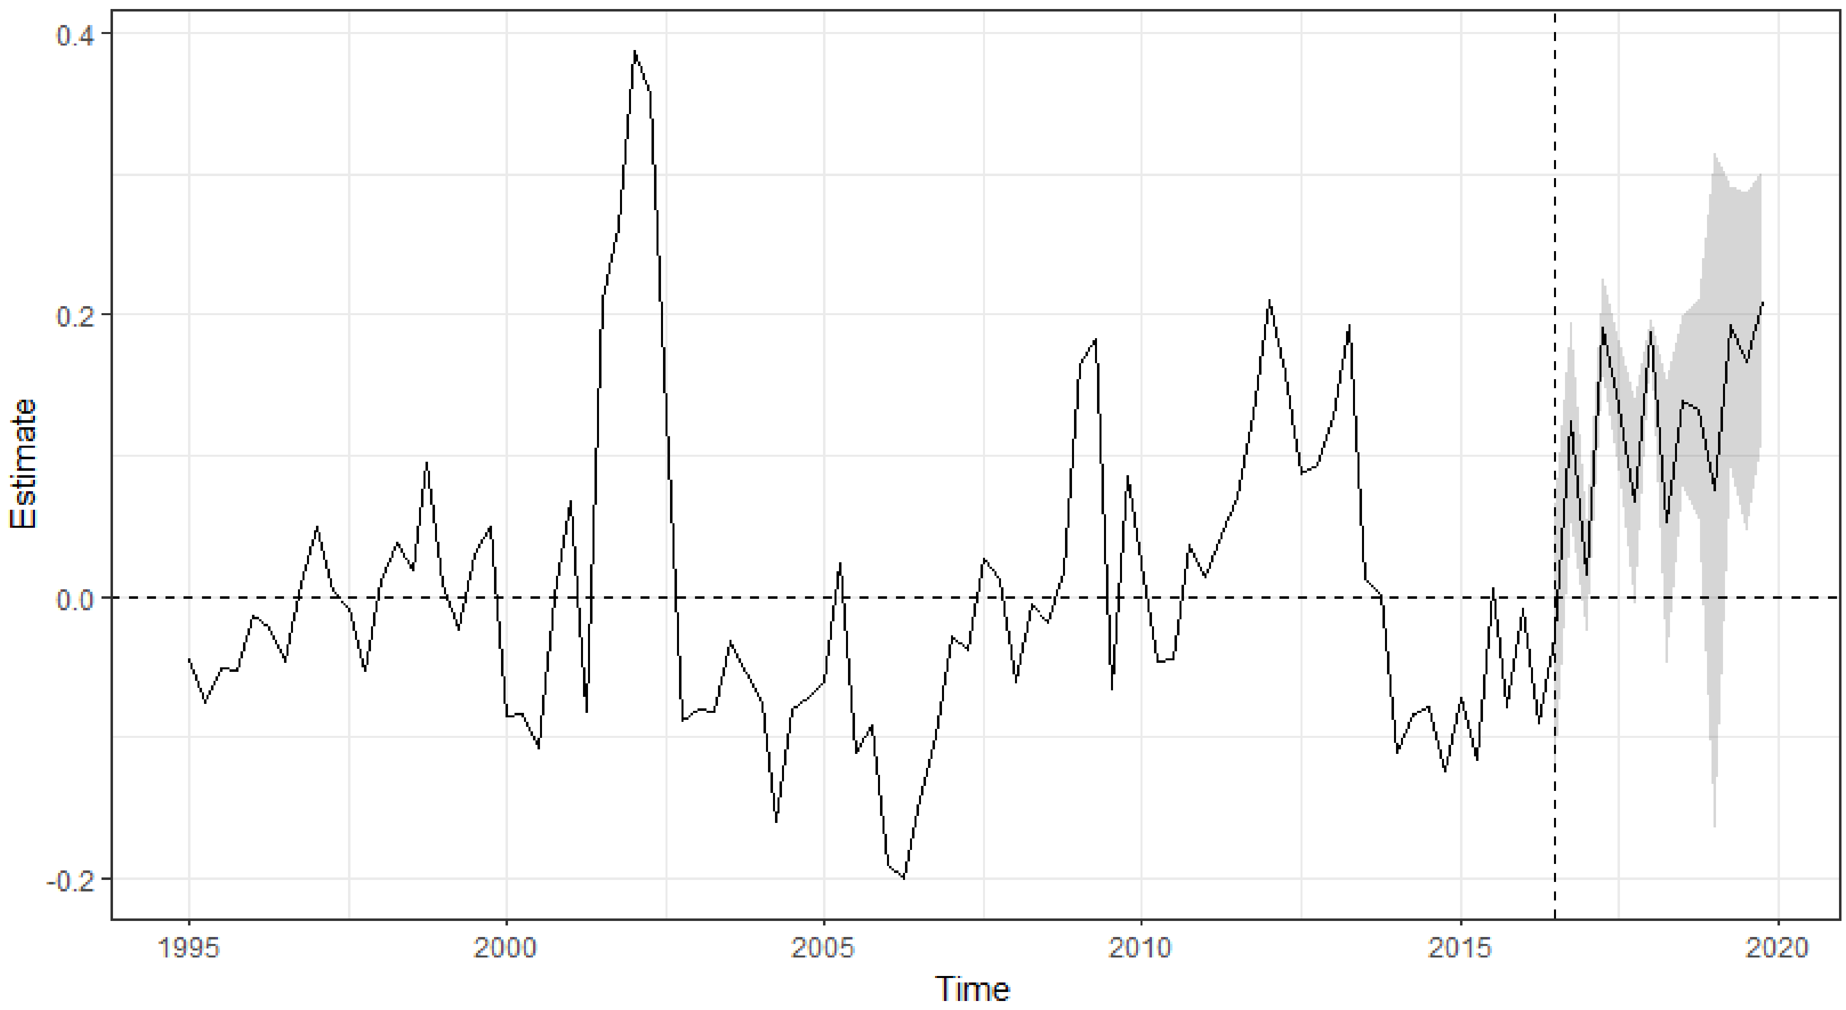

Note: Point estimates along with point-wise 95% conformal confidence intervals for a placebo Brexit effect on bilateral trade. Augmented SCM estimates derived using a Bayesian Causal Impact estimator and additional macroeconomic covariates.

6 Conclusions

Quantifying the aggregate effects of consequential policy interventions is a distinct challenge within the field of empirical macroeconomics. Following the Brexit vote, considerable uncertainty surrounded the UK's decision to leave the EU single market, primarily with respect to the impact on international trade flows. However, a unique intervention by the UK and EU supporting the retention of a "no-border solution" on the island of Ireland allowed for a mechanism through which Ireland could maintain tariff-free trade with the UK economy. This intervention reduced the uncertainty associated with Irish importers and exporters maintaining trade contacts with UK firms. Naturally, this outcome raises the question of the counterfactual: to what degree would Irish trade flows with the UK have developed had it not been for these guarantees of a "no-border solution"?

To answer this question, we capitalize on the natural experiment characteristics of the Brexit vote and Irish border guarantees, in order to assess the effects of policy interventions before the actual implementation of formal procedures and protocols. We employ a bias-corrected form of an approximate balancing weights estimator from the treatment evaluation literature: the augmented synthetic control method, which compares a treated country (Ireland) with an estimated counterfactual (other EU member states). Under this approach, the counterfactual is a linear combination of comparison units that are similar to the treated economy along covariates (economic, demographic and financial factors) and pretreatment realizations of the outcome variable, with an outcome model used to correct for bias due to imperfect pretreatment fit.

Using this ASCM, we estimate a counterfactual response of Irish trade flows with the UK, constructed from a donor pool of other EU member states. Given the stated importance of maintaining the "no border" status quo, we consider differences between realised and counterfactual import and export flows to be due to the causal effect of the "no hard border" commitment, which mitigated adverse trade effects between the Brexit vote in June 2016 and approval of the Brexit withdrawal agreement by UK an EU parliaments in January 2020. Results from our model suggest that, between 2015q3 and 2019q4, Irish exports to the UK were 16.3% above our synthetic counterfactual, while Irish imports from the UK were 15.6% above counterfactual estimates. Our results are robust to changing the timing of the effects of the Brexit vote and border commitments, and the EU member state sample used to construct our counterfactual doppelganger.

The analysis from our paper can be view in contrast to the majority of prior work on the effects of Brexit. Most natural experiment research on Brexit estimates the counterfactual without a policy shock that results in *negative* anticipation and uncertainty effects, to quantify the impact from lowered expectations of economic outcomes. Instead, we present our counterfactual without a policy shock that results in *positive* anticipation and uncertainty effects. Thus, in contrast to works like Börn et al. (2019) who estimate the losses from a change in expectations about the UK's economic future, we measure the economic gains resulting from the anticipated retention of tariff-free trade channels between Ireland and the UK in a post-Brexit world.

These gains are consistent with the identified effects of free-trade agreements identified in the international trade literature. Several studies, including Baier and Bergstrand (2014) and Anderson and Yotov (2016), find FTAs to have positive and significant effects on trade flows. While the short-run effects of FTAs are considered to be substantial, the long-run effects are typically much larger, suggesting that the effects of the negotiated Brexit Protocols between Ireland, Northern Ireland and Great Britain could extend well beyond the time period under analysis in this paper, should they be maintained.[11] Crucial to these effects is the reduction in policy uncertainty (Handley and Limão, 2007), which allows firms to make costly irreversible investment decisions without the concern that the negotiated trade agreements remain credible or face some probability of reversal.

References

ABADIE, A., A. DIAMOND, AND J. HAINMUELLER (2010): "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 105(490), 493–505.

——— (2015): "Comparative Politics and the Synthetic Control Method," *American Journal of Political Science*, 59(2), 495–510.

ABADIE, A., AND J. GARDEAZABAL (2003): "The Economic Costs of Conflict: A Case Study of the Basque Country," *American Economic Review*, 93(1), 113âĂŞ132.

AICHELE, R., AND G. FELBERMAYR (2015): "Costs and Benefits of a United Kingdom Exit from the European Union," Project report, Bertelsmann Stiftung.

ANDERSON, J. E., AND Y. V. YOTOV (2016): "Terms of trade and global efficiency effects of free trade agreements, 1990âĂŞ2002," *Journal of International Economics*, 99(C), 279–298.

BAGEHOT (2018): "Britain's Government Slides into Chaos," The Economist.

BAIER, S. L., AND J. H. BERGSTRAND (2007): "Do free trade agreements actually increase members' international trade?," *Journal of International Economics*, 71(1), 72–95.

[11] Baier and Bergstrand (2014) estimate that FTAs increase bilateral trade flows by approximately 35 per cent over a four year period, while their effects over a 15 year period increase trade flows by over 86 per cent.

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, 131(4), 1593–1636.

Ben-Michael, E., A. Feller, and J. Rothstein (2021): "The Augmented Synthetic Control Method," *Journal of the American Statistical Association*, 116(536), 1789–1803.

Berry, F. (2003): "Irish Economic Development over Three Decades of EU Membership," *Czech Journal of Economics and Finance (Finance a uver)*, 53(9-10), 394–412.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685.

Booth, S., C. Howarth, M. Persson, R. Ruparel, and P. Swidlicki (2015): "What if...? The Consequences, Challenges and Opportunities Facing Britain outside EU," Project report, Open Europe.

Border and Protocol Delivery Group (2016): "HM Treasury Analysis: The Long-Term Economic Impact of EU Membership and the Alternatives," Technical report, HM Government.

——— (2022): "The Border with the European Union: Importing and Exporting Goods," Technical report, HM Government.

Born, B., A. M. Dietrich, and G. J. Müller (2021): "The lockdown effect: A counterfactual for Sweden," *PLOS ONE*, 16(4), 1–13.

Born, B., G. J. Müller, M. Schularick, and P. Sedláček (2019): "The Costs of Economic Nationalism: Evidence from the Brexit Experiment," *The Economic Journal*, 129(623), 2722–2744.

Born, B., and J. Pfeifer (2014): "Policy risk and the business cycle," *Journal of Monetary Economics*, 68(C), 68–85.

CHERNOZHUKOV, V., K. WÜTHRICH, AND Y. ZHU (2021): "An Exact and Robust Conformal Inference Method for Counterfactual and Synthetic Controls," *Journal of the American Statistical Association*, 116(536), 1849–1864.

CORICELLI, F., N. CAMPOS, AND L. MORETTI (2016): "Economic Growth and Political Integration: Estimating the Benefits from Membership in the European Union Using the Synthetic Counterfactuals Method," Working Papers hal-01267266, HAL.

CROWLEY, M. A., O. EXTON, AND L. HAN (2019): "Renegotiation of Trade Agreements and Firm Exporting Decisions: Evidence from the Impact of Brexit on UK Exports," CEPR Discussion Papers 13446, C.E.P.R. Discussion Papers.

DHINGRA, S., H. HUANG, G. OTTAVIANO, J. P. PESSOA, T. SAMPSON, AND J. V. REENEN (2017): "The costs and benefits of leaving the EU: trade effects," *Economic Policy*, 32(92), 651–705.

DHINGRA, S., G. OTTAVIANO, AND T. SAMPSON (2017): "A hitch-hiker's guide to post-Brexit trade negotiations: Options and principles," *Oxford Review of Economic Policy*, 33(suppl1), 22–30.

DHINGRA, S., G. OTTAVIANO, T. SAMPSON, AND J. V. REENEN (2016): "The impact of Brexit on foreign investment in the UK," CEP Brexit Analysis Papers 03, Centre for Economic Performance, LSE.

DJANKOV, S. (2017): "The City of London after Brexit," Policy Briefs PB17-9, Peterson Institute for International Economics.

DOUCH, M., AND T. H. EDWARDS (2022): "The bilateral trade effects of announcement shocks: Brexit as a natural field experiment," *Journal of Applied Econometrics*, 37(2), 305–329.

DOUDCHENKO, N., AND G. W. IMBENS (2016): “Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis,” NBER Working Papers 22791, National Bureau of Economic Research, Inc.

EBELL, M., I. HURST, AND J. WARREN (2016): “Modelling the long-run economic impact of leaving the European Union,” *Economic Modelling*, 59(C), 196–209.

EBELL, M., AND J. WARREN (2016): “The Long-Term Economic Impact of Leaving the EU,” *National Institute Economic Review*, 236(1), 121–138.

FLYNN, E., J. KREN, AND M. LAWLESS (2021a): “Early reactions of EU-UK trade flows to Brexit,” Papers WP713, Economic and Social Research Institute (ESRI).

——— (2021b): “Initial impact of Brexit on Ireland-UK trade flows,” Papers WP714, Economic and Social Research Institute (ESRI).

GRAZIANO, A. G., K. HANDLEY, AND N. LIMÃO (2021): “Brexit Uncertainty and Trade Disintegration,” *The Economic Journal*, 131(635), 1150–1185.

HANDLEY, K. (2014): “Exporting under trade policy uncertainty: Theory and evidence,” *Journal of International Economics*, 94(1), 50–66.

HANDLEY, K., AND N. LIMÃO (2015): “Trade and Investment under Policy Uncertainty: Theory and Firm Evidence,” *American Economic Journal: Economic Policy*, 7(4), 189–222.

HAYWARD, K., AND M. KOMAROVA (2022): “The Protocol on Ireland/Northern Ireland: Past, Present, and Future Precariousness,” *Global Policy*, 13(S2), 128–137.

KIERZENKOWSKI, R., N. PAIN, E. RUSTICELLI, AND S. ZWART (2016): “The Economic Consequences of Brexit: A Taxing Decision,” Economic policy papers, OECD.

KREN, J., AND M. LAWLESS (2024): “How has Brexit changed EU-UK trade flows?,” *European Economic Review*, 161(C).

LAWLESS, M., AND E. L. W. MORGENROTH (2019): "The product and sector level impact of a hard Brexit across the EU," *Contemporary Social Science*, 14(2), 189–207.

MCGRATTAN, E., AND A. WADDLE (2020): "The Impact of Brexit on Foreign Investment and Production," *American Economic Journal: Macroeconomics*, 12(1), 76–103.

O'ROURKE, K. H. (2017): "Independent Ireland in Comparative Perspective," *Irish Economic and Social History*, 44(1), 19–45.

SAMPSON, T. (2017): "Brexit: The Economics of International Disintegration," *Journal of Economic Perspectives*, 31(4), 163–184.

WALKER, N. (2021): "Brexit Timeline: Events Leading to the UK's Exit from the European Union," Briefing Paper 7960, House of Commons Library.

Appendix: The Synthetic Control Method

I General Synthetic Control Methods

Consider a set of j units, where $j = 1, 2, \ldots, J$. Assume $j = 1$ is the treated unit, subject to a different state of the world than the other units in the sample. The remaining set of units, $j = 2, 3, \ldots, J$ is the "donor pool", who form the set of potential comparisons from which we can develop out "synthetic", or counterfactual, treated unit.

Let the data in the sample cover T time periods, with the first T_0 periods occurring before the intervention, and the remaining $T - T_0$ periods spanning the post-intervention period. For each unit j, in each time period t, we observe the outcome for our variable of interest, $Y_{j,t}$ and a set of exogenous predictor variables, $X_{1,j}...X_{k,j}$. The $k{\times}j$ vectors $X_1 \ldots X_J$ contain predictor values for each of the $j = 1, \ldots, J$ units. Similarly, the $k{\times}(J - 1)$ matrix, $X_0 = [X_2 \ldots X_J]$ contains values of the predictors for the $(J - 1)$ untreated units. For each unit, j, and time period, t, let Y_{jt}^N be the potential response without intervention. For the unit affected by the intervention $(j = 1)$, define Y_{1t}^I to be the potential response under the intervention, in the post intervention period, i.e. $\forall\ t > T_0$.

Under this structure, $\forall\ t > T_0$, the effect of the intervention of interest for the affected unit in period t is:

$$\tau_{1t} = Y_{1t}^I - Y_{1t}^N. \tag{8}$$

Because unit $j = 1$ is exposed to the intervention from period T_0 onwards, it is axiomatically the case that $Y_{1t} = Y_{1t}^I\ \forall\ t > T_0$. Thus, for the all time periods subsequent to the intervention, the only observable outcome for the treated unit(s) of interest is the potential outcome under the intervention. However, as is the case for most event study research, a key question of interest relates to the causal effect of the treatment. Counterfactuals are commonly used to asses these causal relationships through the modelling of alternate scenarios where the intervention did not affect the treated unit after $t = T_0$, i.e. estimating the unobserved value Y_{1t}^N. As permanent effects of the intervention may take time to be observed, it is important to allow τ_{1t} to vary across the time dimension.

Given the complexity of modern economic and financial systems, counterfactual analysis can be difficult to implement in macroeconomic research. It can be exigent to accurately match characteristics of the treated country or region to a single other untreated country, given that divergences in economic outcomes could be the result of pre-existing differences in monetary, fiscal, exchange rate or regulatory policies, or general developments in financial or labour markets. Consequently, it may be more appropriate to generate counterfactual scenarios using a combination of several countries from the donor pool, giving a better approximation to the values of Y_{1t}^N by more closely matching the treated country across multiple dimensions.

Definitionally, this is exactly what a synthetic control is: a weighted average of a combination of units from a donor pool. A synthetic control can be represented mathematically by a $(J-1) \times 1$ vector of weights, $\boldsymbol{W} = (\omega_2, \ldots, \omega_J)'$. Given a set of weights, $\boldsymbol{W}$, the synthetic control estimators of Y_{1t}^N and τ_{1t} have the form:

$$\hat{Y}_{1t}^N = \sum_{j=2}^{J} \omega_j Y_{jt} \qquad \text{and} \tag{9}$$

$$\hat{\tau}_{1t} = Y_{1t} - \hat{Y}_{1t}^N \tag{10}$$

Typically, weights are restricted to be non-negative and to sum to one. Restricting weights in this manner produces synthetic controls that are weighted averages of outcomes of the donor pool units, with weighting structures that are likely sparse, without the risk of extrapolation bias. As only a small number of units contribute to the estimate of the $\hat{Y}_{1t}^{N}$ counterfactual, and the contribution of each unit is uniquely represented by its weight, SCM counterfactual estimates are both intuitive and transparent.

Given the central role of the donor pool weights in determining the counterfactual estimate, construction of the weighting system is critical to the estimate of $\hat{\tau}_{1t}$. Several pioneering research papers in SCM estimation, including Abadie and Gardeazabal (2003) and Abadie, Diamond, and Hainmueller (2010) propose selecting $(\omega_2, \ldots, \omega_{J+1})$ so that the resulting synthetic control is an optimal fit for the outcome variable's treated unit of predictors in the pre-intervention period. With a set of non-negative constants, $(\upsilon_1, \ldots, \upsilon_k)$, the synthetic control $\boldsymbol{W}^* = (\omega_2^*, \ldots, \omega_J^*)'$ that provides this optimal fit will solve the function

$$\begin{gathered} min \parallel \boldsymbol{X}_1 - \boldsymbol{X}_0 \boldsymbol{W} \parallel = min \left(\sum_{h=1}^{k} \upsilon_h \left(X_{h1} - \omega_2 X_{h2} - \ldots - \omega_J X_{hJ} \right)^2 \right)^{\frac{1}{2}} \\ \text{s.t.} \quad \omega_j \geq 0 \quad \forall \; j = 2, \ldots, J \qquad \& \qquad \sum_{j=2}^{J} \omega_j = 1 \end{gathered} \tag{11}$$

Under this weighting structure, the estimated treatment effect for the treated unit at time $t = T_0 + 1, \ldots, T$ is

$$\hat{\tau}_{1t} = Y_{1t} - \sum_{j=2}^{J} \omega_j^* Y_{jt} \tag{12}$$

The vector of constants $(\upsilon_1, \ldots, \upsilon_k)$ from equation (11) can be considered a relative importance vector, defining the relevance that the synthetic control attaches to matching the realised values of each of the k predictors for the treated unit, $X_{11}, \ldots, X_{k1}$. For any choice of importance vector $\boldsymbol{V} = \upsilon_1, \ldots, \upsilon_k$, constrained quadratic optimization can be used to solve equation (11). Thus, each potential vector $\boldsymbol{V} = (\upsilon_1, \ldots, \upsilon_k)$

defines a synthetic control, $\boldsymbol{W}(\boldsymbol{V}) = (\omega_2(\boldsymbol{V}), \ldots, \omega_{J+1}(\boldsymbol{V}))'$, which can be identified by minimizing equation (11), subject to constraining the weights in $\boldsymbol{W}(\boldsymbol{V})$ to both be positive and sum to one. This approach shifts the problem from identifying the preferred weighting system $\boldsymbol{W}$ to identifying the optimal importance vector $\boldsymbol{V}$.

The purpose of the synthetic control is to approximate the true counterfactual dynamics in our variable of interest, which would have been observed in $t > T_0$ had the intervention not occurred, for Y_{1t}. Given these requirements, a weighting system $\boldsymbol{W}$ is chosen such that the resulting synthetic control optimally reflects the treated unit before the intervention, along the dimension of a set of characteristic variables $X_{11}, \ldots, X_{k1}$. The problem of choosing $\boldsymbol{V} = (\upsilon_1, \ldots, \upsilon_k)$ can then be interpreted as choosing values for each υ_{1h} that capture the relative importance of X_{h1} for predicting Y_{1t}^N in the post intervention period, $t = T_0 + 1, \ldots, T$.

As Y_{1t}^N is a latent variable for $t = T_0 + 1, \ldots, T$, it is not possible to directly estimate the relative importance of each predictor in approximating Y_{1t}^N once the intervention has occurred. However, given that Y_{1t}^N is observable in the pre-intervention period $t = 1, 2, \ldots, T_0$, pre-intervention data can be used to quantify the predictive power of the variables $X_{1j}, \ldots, X_{kj}$ in estimating Y_{1t}^N. As developed by Abadie, Diamond and Hainmueller (2015), a preferred approach to solving for $\boldsymbol{V}^*$ and $\boldsymbol{W}^*$ can take the following form:

I Partition the pre-intervention period into an initial training period $(t = 1, \ldots, t_0)$ and a subsequent validation period $(t = t_0 + 1, \ldots, T_0)$.

II For every candidate vector $\boldsymbol{V}$, let $\tilde{\boldsymbol{W}} = (\tilde{\omega}_2(\boldsymbol{V}), \ldots, \tilde{\omega}_{J+1}(\boldsymbol{V}))$ be the synthetic control weights computed with training period data on the predictors. The mean squared prediction error (MSPE) of this synthetic control with respect to Y_{1t}^N in the validation period is

$$\sum_{t=t_0+1}^{T_o} (Y_{1t} - \tilde{\omega}_2(\boldsymbol{V})Y_{2t} - \ldots - \tilde{\omega}_{J+1}(\boldsymbol{V})Y_{J+1t}) \tag{13}$$

III Choose $\boldsymbol{V}^*$ such that the MSPE in equation (13) is minimized.

IV Conditional on $\boldsymbol{V}^*$ and predictor data for the pre-intervention validation period, calculate $\boldsymbol{W}^* = \boldsymbol{W}(\boldsymbol{V}^*)$.

As the procedure is numerical in nature, its value is dependent on the degree to which the candidate $\boldsymbol{V}^*$ solves $Y_{1t} \approx \tilde{\omega}_2(\boldsymbol{V}^*)Y_{2t} + \cdots + \tilde{\omega}_J(\boldsymbol{V}^*)Y_{Jt}$ for the set of predictors used to calculate $\boldsymbol{W}^*$, and the sensitivity of results to different solution values of $\boldsymbol{V}$.

II SCM Representation

As popularised by Abadie and Gardeazabal (2003) and Abadie, Diamond, and Hainmueller (2010), the SCM problems imputes the missing potential outcome for the treated unit, $Y_{1T}(0)$, as a weighted average of the control outcomes, $Y_{0T}'\gamma$. Weights are chosen to balance a combination of pre-treatment outcomes and other auxiliary covariates. We can represent the SCM as a solution to the constrained optimization problem

$$\begin{aligned} \underset{\gamma \in \Delta^{N_0}}{min} \quad & \theta_x \parallel V_x^{\frac{1}{2}}(\boldsymbol{X}_{1\cdot} - \boldsymbol{X}_{0\cdot}'\boldsymbol{\gamma}) \parallel_2^2 + \theta_z \parallel V_z^{\frac{1}{2}}(\boldsymbol{Z}_{1\cdot} - \boldsymbol{Z}_{0\cdot}'\boldsymbol{\gamma}) \parallel_2^2 \\ \text{s.t.} \quad & \Delta^{N_0} = \{\gamma \in \mathbb{R}^{N_0} | \gamma_i \geq 0 \quad \forall\ i, \sum_i \gamma_i = 1\} \end{aligned} \tag{14}$$

where $\boldsymbol{V}_x$ and $\boldsymbol{V}_z \in \mathbb{R}^{T_0 \times T_0}$ are symmetric importance matrices, and $\parallel V_x^{\frac{1}{2}}(\boldsymbol{X}_{1\cdot} - \boldsymbol{X}_{0\cdot}'\boldsymbol{\gamma}) \parallel_2^2$ and $\parallel V_z^{\frac{1}{2}}(\boldsymbol{Z}_{1\cdot} - \boldsymbol{Z}_{0\cdot}'\boldsymbol{\gamma}) \parallel_2^2$ are the L^2 norms onto $\mathbb{R}^{T_0}$ after applying $V_x{}^{\frac{1}{2}}$ and $V_z{}^{\frac{1}{2}}$ as linear transformations.

We can view the SCM optimization problem in equation (14) as an approximate balancing weights estimator. As with all balancing estimators, a central question is what quantity to balance.

Following the recent methodological literature (see Doudchenko and Imbens, 2017; Ferman and Pinto, 2018), equation (11) directly optimizes for the pre-treatment period, minimizing the (possibly weighted) imbalance of pre-treatment outcomes between the treated unit and the weighted control mean. The SCM estimator for $t > T_0$ is given by

$$\text{SCM Estimator} = \hat{Y}_{1t} = \left(\sum_{i\neq 1} \gamma_i X_i\right)^{'} \hat{\theta}_x + \left(\sum_{i\neq 1} \gamma_i Z_i\right)^{'} \hat{\theta}_z + \left(\sum_{i\neq 1} \gamma_i \varepsilon_{it}\right) \tag{15}$$

We can decompose the bias of the SCM estimator into two elements

$$\begin{aligned} Y_{1t} - (\hat{\theta}_x \sum_{i\neq 1} \gamma_i X_i + \hat{\theta}_z \sum_{i\neq 1} \gamma_i Z_i) = \left[Y_{1t} - \hat{Y}_{1t}\right] + \\ \left[\hat{Y}_{1t} - (\hat{\theta}_x \sum_{i\neq 1} \gamma_i X_i + \hat{\theta}_z \sum_{i\neq 1} \gamma_i Z_i)\right] + \end{aligned} \tag{16}$$

where the $\left[Y_{1t} - \hat{Y}_{1t}\right]$ component of equation (16) is the extrapolation bias of the estimator, and $\left[\hat{Y}_{1t} - (\hat{\theta}_x \sum_{i\neq 1} \gamma_i X_{it} + \hat{\theta}_z \sum_{i\neq 1} \gamma_i Z_{it})\right]$ is the interpolation bias of the estimator.

Assuming the estimator to be a Lipschitz function, the magnitude of the extrapolation bias is bounded by

$$|\text{ExtBias}(\gamma)| \leq C \parallel \theta_x X_1 + \theta_z Z_1 - \sum_{i\neq 1} \gamma_i (\hat{\theta}_x X_i + \theta_z \hat{Z}_i) \parallel \; \equiv C_1 \times \text{Ext}(\gamma) \tag{17}$$

Additionally, assuming that $X_1 = \sum_{i \neq 1} \gamma_i X_i$ exists, the magnitude of interpolation bias will be bounded by

$$
\begin{aligned}
|\text{IntBias}(\gamma)| &= |\sum_{i \neq 1} \gamma_i (Y_{1t} - Y_{it})| \\
&\leq \sum_{i \neq 1} \gamma_i |Y_{1t} - Y_{it}| \\
&\leq C_2 \sum_{i \neq 1} \gamma_i |X_1 - X_i| \equiv C \times Int(\gamma)
\end{aligned}
\tag{18}
$$

By construction, the SCM estimator minimizes the quantity Ext(γ) that bounds extrapolation bias, but is susceptible to interpolation bias. Under the above properties, their exist two conditions under which the SCM estimator is unbiased,

$$
\begin{aligned}
\sum_{i \neq 1} \gamma^* Z_i = Z_1 \\
\sum_{i \neq 1} \gamma^* X_i = X_1
\end{aligned}
\tag{19}
$$

When the treated unit's vector of lagged outcomes, X_1, is inside the convex hull of the control units' lagged outcomes, X_0, the SCM weights in equation (11) achieve perfect pre-treatment fit, and the resulting estimator has many attractive properties, including a bias bound established by Abadie et al. (2010). Due to the curse of dimensionality, however, achieving perfect (or nearly perfect) pre-treatment fit is not always feasible with weights constrained to be on the simplex (see Ferman and Pinto, 2018). When the pre-treatment fit is poor or the number of pre-treatment periods is small, Abadie et al. (2015) recommend against using SCM. Even if the pre-treatment fit is excellent, Abadie et al. (2010, 2015) propose extensive placebo checks to ensure that SCM weights do not overfit to noise. Thus, conditionality of the analysis is essential to the appropriate application of SCM.