UNIVERSITY OF WARWICK

Economics

Measuring the Regional Economic Cost of Brexit: Evidence as of 2026

Eleonora Alabrese, Jacob Edenhofer, Thiemo Fetzer and Shizhuo Wang

June 2026
No: 1617

Warwick Economics Research Papers

ISSN 2059-4283 (online)
ISSN 0083-7350 (print)

Measuring the Regional Economic Cost of Brexit: Evidence as of 2026

Eleonora Alabrese Jacob Edenhofer Thiemo Fetzer
Shizhuo Wang *

June 16, 2026

Abstract

This paper examines the regional distribution of Brexit's economic costs across the United Kingdom. We apply a synthetic control approach to two complementary outcomes – real gross value added (GVA) and nominal gross disposable household income (GDHI) – covering multiple levels of UK spatial aggregation. We construct placebo-weighted ensemble counterfactuals using both post-2016 and post-2020 treatment windows. We document three main findings. Economic losses are large and geographically widespread: around 70% of local authority districts record output or income below their synthetic counterfactual. Losses are unevenly distributed. They are concentrated in initially prosperous and trade-integrated regions, particularly London, the South East, and Scotland, while less affluent areas experience comparatively smaller declines. This pattern is consistent with a process of "levelling down": Brexit has reduced regional inequality not by improving economic performance in lagging areas, but by disproportionately weakening leading ones. Northern Ireland is a clear exception, having been partly insulated from these losses by its continued access to key elements of the EU single market.

Keywords: Brexit, Regional Inequalities, Synthetic Control, Economic Integration, Subnational GDP
JEL Classification: F15, H72, R11, C21

*Alabrese is affiliated with the University of Bath, SAFE and QAPEC. Edenhofer is associated with the University of Oxford. Fetzer is affiliated with the University of Warwick, University of Bonn, CEPR, ECONtribute, STICERD, Grantham, CAGE, NIESR, and CESifo. Wang is associated with New York University. Junxi Liu and Ivan Yotzov provided excellent research assistance. We appreciate feedback from seminar audiences at the NIESR event on the economic consequences of Brexit. *Email addresses:* ea2143@bath.ac.uk, jacob.edenhofer@some.ox.ac.uk, t.fetzer@warwick.ac.uk, shizhuo.wang@nyu.edu. An interactive explorer of the estimates, their sensitivity along with broader narration that is specific to different areas can be found on https://www.brexitcost.org.

"That's Your Bloody GDP, Not Ours."

1 Introduction

The United Kingdom (UK)'s vote to leave the European Union (EU) in 2016 marked a watershed when it became clear that populists could succeed, even in one of the world's oldest democracies. The Leave campaign heavily relied on messaging that, among others, linked immigration and the UK's EU budget contributions to the ailing state of local public services amidst growing spatial inequalities. In doing so, the Leave campaign effectively mobilised the UK's large (Wiedemann, 2024) and growing regional or spatial inequalities.[1] By contrast, the Remain campaign's core message – that the *aggregate* economic consequences of leaving the EU would be dire – failed to resonate in many places, despite the fact that this warning proved to be largely correct (Born et al., 2019; Hassan et al., 2024; Steinberg, 2019; Sampson, 2017; Dhingra and Sampson, 2022; Breinlich et al., 2022; Bakker et al., 2023; Grassi, 2024; Novy, 2024). The *aggregate* economic effects were shrugged off by many voters, especially those in chronically deprived and long-declining regions in the Midlands and the North of England (Carreras et al., 2019). In fact, even in March 2021, more than 40% of respondents believed that Brexit was the right decision (YouGov, 2023). Since then, however, support for the referendum outcome has continued to erode: by June 2026, only 30% of respondents thought Britain was right to vote to leave the EU, compared with 57% who thought it was wrong (YouGov, 2026).

These twists and turns of the Brexit saga raise an as of yet unanswered question: what is the regional incidence of Brexit's economic costs within the UK? We

[1] Growing spatial inequalities may have been the result of policymakers' desire to realise agglomerations effects and thus achieve an economically more productive spatial allocation of people. The rather narrow focus on spatial economic efficiency, however, might have led academics and policymakers alike to ignore the distributional and political ramifications of these strong agglomeration effects. The UK's spatial divide was already unusually large before the referendum. Indeed, Stansbury et al. (2023) argue that this divide largely reflects a productivity gap between London and the Greater South East and the rest of the country, driven mainly by the weak productivity performance of major non-London cities.

study this (descriptive) question using a synthetic control approach (Abadie and Gardeazabal, 2003; Abadie et al., 2010; Pickett et al., 2026). For each spatial unit, we construct a placebo-weighted ensemble counterfactual across a broad family of donor pool configurations. In essence, we estimate many potential synthetic controls — each built from a different set of comparator economies — and weight them by how well each predicts actual UK regional trajectories in a pre-Brexit validation period (see Section 2.4). We exploit two distinct outcome measures — real gross value added (GVA) and GDHI (formally the B.6g balance of secondary distribution of income, which measures what households can spend or save after taxes and transfers)[2] — under both a post-2016 and a post-2020 treatment window. The Brexit-induced output or income gap is the difference between the actual series and its placebo-weighted synthetic control. Comparing the two series allows us to examine whether the production cost of Brexit and the household income cost align, and where the fiscal redistribution system — through which output losses in one place may be partly offset by transfers to households in another — has cushioned or amplified the shock.[3]

The joint production–income design also allows us to use Brexit as a shock to study how UK regions are economically connected. GVA records where production takes place; GDHI records where households receive their disposable income. These two measures are linked through commuting, wage flows, ownership claims, taxes, transfers, and residence choices — the connective tissue of a regional economy. A useful benchmark is provided by recent work on Disaggregated Economic Accounts (DEA), which reconstructs who receives income from production in each place by tracing the bilateral flows among firms, households, government, and the rest of the world (Andersen et al., 2026). We do not observe that full bilateral

[2]We use GDHI throughout as shorthand for this specific transaction. The distinction from gross primary income matters because the B.6g concept includes social benefits and transfers received and subtracts taxes paid, making it a closer analogue to household purchasing power.

[3]The estimates, interactive visualizations and collections of local impact stories can be found on `https://www.brexitcost.org`.

flow matrix for UK local authorities. Instead, where GVA and GDHI diverge after Brexit, the divergence provides indirect evidence about the channels through which a change in the UK's economic links with the EU was transmitted across regions. The sign of that divergence is itself informative. Where a region's production loss exceeds its income loss, part of the local output shortfall has been absorbed by households elsewhere through taxes, transfers, or claims on production beyond the region; where the income loss exceeds the production loss, residents bear losses that originate in production declines outside their own boundaries.

Using this approach, we document four main findings. First, the cost of Brexit — in the form of lower trend growth — is large and near-universal across all regions and constituent countries of the United Kingdom. Output losses range from 5 to 10 percentage points of GDP, relative to a large set of synthetic control estimates. Second, there is notable cross-country and regional heterogeneity in these costs. Our estimates show that Northern Ireland has to date not been adversely affected by Brexit. This is not surprising, given that Northern Ireland — unlike the rest of the UK — effectively continues to be part of the EU customs union, as stipulated in the EU-UK Trade and Cooperation Agreement. Yet, even across English regions, the costs are heterogeneous. While the vast majority of local authorities — around 70% — have experienced some cost of Brexit, only about 30% of districts appear to have outperformed their respective synthetic control. These figures refer to GVA gaps under the post-2016 window at the local authority district (LAD) level; shares for the GDHI measure and the post-2020 window follow a similar pattern.[4] The share of areas with output losses is notably higher in Scotland (93%), London (85%), and the South West (83%) — while the region with the highest share of areas with output gains is Northern Ireland (93%).[5] These findings point to a striking

[4] *Author verification needed:* confirm these percentages against Table 1 and the LAD-level distribution in Figure 3.

[5] The reconfiguration of trade routes towards the Island of Ireland from mainland Europe has likely resulted in geographic reallocation of economic activity within Northern Ireland, benefiting local economies in the hinterland, near the Irish border. By contrast, the Northern Irish port through which trade from mainland Europe was shipped via the England "land bridge" sees significant

conclusion. Brexit has contributed to levelling up the United Kingdom — but only by levelling down its economically more successful regions. Put differently, losses are concentrated in initially richer and more trade-integrated areas, narrowing the gap in relative terms while making the UK poorer in aggregate. Whether this constitutes a formal reduction in measured regional inequality would require a Gini or similar decomposition, which we discuss in the context of the cross-sectional results.

Third, we show that none of the covariates that — taken together or individually — were strong correlates of support for Leave in 2016 (Becker et al., 2016) explain a significant part of the cross-sectional variation in the cost of Brexit to date. The same is true for an area's exposure to the Covid-19 pandemic, as measured by mortality, or increased receipts of central government transfers under various "levelling-up" funds aimed at reducing spatial inequality.[6] That the gap estimates are orthogonal to these potential confounders supports the identifying assumptions of the synthetic control design.

The estimates are robust to an extensive battery of robustness checks. First, to allay concerns that the synthetic control estimates are biased by idiosyncrasies in the donor pool, we construct a placebo-weighted ensemble across the full set of $2^5 - 1 = 31$ donor pool combinations drawn from European NUTS2 regions, US states, G20 countries, OECD countries, and EU member states. Second, we score each potential synthetic control — that is, each pairing of a UK region with a specific donor pool — on out-of-sample predictive performance using pseudo-treatment windows that pre-date Brexit (pretending treatment happened in 2010–2013 and testing how well the counterfactual predicts the following years), so the primary

economic damage.

[6]There is, however, one notable exception, i.e. areas with higher levels of support for Leave in 2016 appear to experience — relatively speaking — lower economic cost of Brexit to date. This correlation is weak, but has gained in strength since 2017. In an earlier version of this paper, that used regional gross value-added estimates up to 2018, the conditional correlation pointed in the other direction. We believe this change in the correlation structure is attributable to methodological changes in the apportionment of regional GDP (Fetzer and Wang, 2020). The sign flip is recovered when focusing on within-region evidence of the cost of Brexit as evidenced in Appendix Table A1.

counterfactual is not chosen to minimise in-sample fit. Third, we validate the gap estimates by applying the same synthetic-control procedure to each untreated donor economy as if it had been treated, and we compare the UK gap to that distribution of placebo gaps. Together, these checks confirm that our estimates of the economic costs of Brexit are not confounded by donor-pool choice or overfitting.

The paper relates to a growing body of work on the aggregate, firm-level, and subnational economic consequences of Brexit. Born et al. (2019) provide the canonical estimate of the aggregate output cost using OECD donor economies. Hassan et al. (2024), Steinberg (2019), Dhingra and Sampson (2022), Breinlich et al. (2022), and Grassi (2024) document losses through trade, investment, uncertainty, and price channels at the national level. Closely related recent work by Bloom et al. (2025) combines cross-country macro counterfactuals with firm-level evidence from the Decision Maker Panel and estimates that Brexit reduced UK GDP by 6–8% by 2025, investment by 12–18%, and employment and productivity by 3–4%. At the subnational level, Javorcik et al. (2026) use high-frequency online job-advert data to estimate local labour-market effects of the Brexit vote, finding that areas more exposed to prospective barriers on professional-services exports experienced relative declines in job adverts, especially for high-skilled and professional/managerial occupations. To our knowledge, this paper is the first to estimate Brexit's economic costs across the full distribution of UK local authority districts while jointly comparing a production measure (GVA) and a disposable-income measure (GDHI) at this spatial resolution.[7] This allows us to speak to the regional incidence of Brexit's economic costs at a scale not available in prior work, and to examine how the fiscal redistribution system — through taxes, transfers, and commuting flows — shapes whether output losses in one place translate into household-income losses for the people who live there.

The paper also connects to emerging work that recognises regions as econom-

[7] Closest prior subnational work studies predicted local exposure or local labour-market outcomes rather than estimating realised GVA and GDHI synthetic-control gaps across the full distribution of local authority districts.

ically interdependent rather than self-contained: income generated in one place flows to residents of another through jobs, commuting, ownership, taxes, and spending. Andersen et al. (2026) develop Disaggregated Economic Accounts that directly reconstruct these bilateral flows among firms, households, government, and the rest of the world, making visible the income-distribution channels that standard regional statistics leave unobserved. Our setting is complementary: Brexit changed the UK's economic links with the EU, while the GVA–GDHI comparison reveals how that change was transmitted across the UK's production and residential geographies.

More broadly, the paper speaks to the literature on economic integration and spatial inequality. Becker et al. (2016) and Rodríguez-Pose (2018) document how spatial inequalities shaped the Brexit vote. Our findings speak to the reverse question: how have Brexit's costs been distributed across places, and what does that distribution imply for spatial inequality? The evidence that losses are disproportionately concentrated in initially prosperous regions — London, the South East, and Scotland — is consistent with a pattern of "levelling up by levelling down", with regional disparities narrowing, not because lagging areas improve, but because leading areas are pulled down.[8]

The remainder of this paper is organised as follows. Section 2 presents the data, the two outcome series (GVA and GDHI), the divergence between production and household income as a diagnostic for how shocks travel across regional economies, the synthetic control estimator, the donor pool design, the placebo-weighted ensemble, placebo-based inference, and the two treatment windows. Section 3 presents the spatial distribution of the estimated costs at the country, ITL1, ITL2, ITL3, and LAD levels, including aggregation consistency checks, a comparison with country-level estimates from the literature, an interpretation of what the GVA–GDHI divergence reveals about production and income incidence, and a systematic exploration of the

[8]This pattern has parallels in the US literature on the distributional effects of tariff policy (Fajgelbaum et al., 2020; Autor, 2024), where protection tends to benefit exposed sectors while imposing costs on economically more dynamic regions.

cross-sectional correlates of the cost distribution. Section 4 summarises the findings and their implications.

2 Data, Measurement, and Estimation Strategy

This section describes the data and estimation strategy used to quantify the regional economic cost of Brexit.

We begin with the two outcome series — GVA and GDHI — and the conceptual tension between them. We then discuss what the divergence between the two measures reveals about how shocks travel across regional economies, before specifying the synthetic control estimator, the donor pool design, the placebo-weighted ensemble, placebo-based inference, and the two treatment windows.

2.1 Outcome measures: GVA and GDHI

Assessing the regional economic cost of Brexit requires at minimum two distinct measures: one recording where output is produced and one recording where households receive their disposable income. GVA captures productive activity at the workplace; GDHI captures what residents can spend or save at home. The two need not move together when trade barriers change. A factory town hit by reduced export demand may see sharp GVA losses, while nearby residential areas lose income through commuting and wage channels. Conversely, a residential suburb whose residents commute to an affected city may see GDHI fall even if local GVA appears stable because the production loss is recorded at the workplace, not where the workers live. This divergence is most clear-cut where the shock — Brexit's change to the UK's economic links with the EU — affects traded sectors and workers who cross local boundaries every day.

Gross Value Added. Our first outcome is real *gross value added* (GVA), available at annual frequency from the Office for National Statistics (ONS) for local authorities, ITL3, ITL2, and ITL1 geographies, as well as the four constituent countries.[9] GVA is a *production* concept. It measures the value of goods and services produced within a geography, net of intermediate consumption. Because it is recorded on a *workplace* basis it captures the economic activity occurring in a region regardless of where the producing workers live. The series is expressed in real chained-volume terms (that is, adjusted for inflation), removing the confound of general price level movements. GVA is therefore the closest subnational analogue of aggregate output and is the measure used by most prior country-level studies of Brexit's economic impact (Born et al., 2019; Steinberg, 2019; Hassan et al., 2024).

Gross Disposable Household Income. Our second outcome is *gross disposable household income* (GDHI), specifically the B.6g transaction in the household sector accounts, available from the ONS at the same geographic levels.[10] B.6g is the *balance of secondary distribution of income*: it starts from primary incomes (operating surplus, compensation of employees, and property income) and adds social benefits and other current transfers received by households, then subtracts current taxes on income and wealth and social contributions paid. GDHI is therefore a *residence-based income* concept that reflects what households can actually spend or save, after the redistribution operated by the public sector. It is reported in *nominal* terms.

[9] ONS, *Regional gross value added (balanced) by industry*, available at https://www.ons.gov.uk/economy/grossvalueaddedgva/datasets/nominalandrealregionalgrossvalueaddedbalancedbyindustry. The series runs from 2000 through 2023 and is expressed in real chained-volume terms at 2019 prices. The 2024 release covers 382 local authority districts. A notable feature of subnational economic data is that it is published with considerably longer lags than the corresponding national aggregates. The ONS has faced repeated delays in releasing sub-regional GVA and GDHI revisions, and the data vintage available to researchers lags the national accounts by one to two years. This paper uses data through 2023, the latest available at the time of analysis. In addition, the ONS discontinued its quarterly *Regional Economic Activity by Gross Value Added* publication in 2020, removing the only source of quarterly sub-regional GDP estimates that had previously been used in earlier versions of this work. All results are therefore based on annual data.

[10] ONS, *Regional gross disposable household income by local authority: 1997 to 2023*, available at https://www.ons.gov.uk/economy/regionalaccounts/grossdisposablehouseholdincome. The series is nominal.

The tension between the two measures. GVA and GDHI-B6 answer different welfare questions, and they need not move together in response to a trade shock. Several forces can drive a wedge between them.

First, *geography of residence versus workplace*. GVA registers where output is produced while GDHI registers where earners live. In regions with large net commuting inflows — London being the prime example — a fall in productive activity reduces GVA at the location of production, but the household-income effect is spread across the wider commuting zone where workers actually live. London residents may also be exposed to income losses through finance, property income, and self-employment channels that are not fully captured in workplace-based GVA. Conversely, a residential suburb or commuter region may see GDHI fall even when local GVA is stable, because its residents earn most of their income in an affected workplace region; the production loss registers in GVA elsewhere but the income loss is recorded where those workers live.

Second, *fiscal redistribution*. GDHI-B6 includes the full flow of taxes and transfers. An aggregate supply shock that reduces incomes may partly be cushioned by an automatic expansion of in-work benefits, tax credits, or unemployment insurance. Regions in which more households are near the margin of eligibility for means-tested transfers will, all else equal, see a smaller decline in GDHI relative to GVA. At the same time, if Brexit-related uncertainty reduces property income or self-employment income—both sizeable in London and the South East—GDHI could fall further than GVA.

Third, *nominal versus real terms*. GVA is deflated; GDHI-B6 is nominal. If Brexit has had asymmetric price effects across regions—for example through the sterling depreciation that followed the referendum, which raised import prices disproportionately for households rather than for business-to-business production—then the gap between real GVA losses and nominal GDHI losses will differ across geographies in ways that are not purely measurement artefact.

Overall, therefore, the joint distribution of GVA and GDHI-B6 gaps provides more information about the economic incidence of Brexit than either series alone. GVA answers the question of where output has been lost; GDHI-B6 answers the question of where household purchasing power has been affected. We present both throughout, and we discuss cases where the two narratives diverge.

The GVA–GDHI-B6 wedge as flow incidence. The divergence between GVA and GDHI-B6 should therefore not be treated as a nuisance induced by measurement. It is a theoretically meaningful quantity in its own right. Let q_{it} denote production in region i at time t, measured by GVA, and let h_{it} denote residence-based disposable household income, measured by GDHI-B6. In a fully observed disaggregated accounting system, household income in region i would be linked to production in all regions j through a matrix of labour-income flows, commuting links, ownership claims, local consumption spillovers, fiscal transfers, and other interregional income channels. Schematically,

$$h_{it} = \sum_j B_{ijt} q_{jt} + r_{it}, \tag{2.1}$$

where B_{ijt} is a latent incidence matrix mapping production in region j into household income in region i, and r_{it} captures redistribution, property-income flows, and other income components not mechanically tied to contemporaneous local production. Standard regional accounts observe q_{it} and h_{it}, but not the full matrix B_{ijt}.

This interpretation connects our measurement strategy to the Disaggregated Economic Accounts framework of Andersen et al. (2026), which directly reconstructs bilateral flows among consumers, producers, government, and the rest of the world while satisfying national accounting identities. Our setting is more reduced form. We use Brexit as a disconnection shock that perturbs external economic integration and then observe how production and residence-based household income respond across space. Where GVA and GDHI-B6 gaps diverge, we interpret the divergence as the shadow of the latent flow network in (2.1): commuting zones, wage flows, ownership claims, transfers, and residence choices mediate the pass-through from production losses to household income losses.

GVA donor sources. Each GVA synthetic control is estimated against a dedicated donor panel built from the closest available international counterpart to the UK GVA series. The five base sources are: (i) real GVA by NUTS2 region from Eurostat (`nama_10r_3gva`), covering up to 223 European sub-national regions; (ii) real GDP by state from the U.S. Bureau of Economic Analysis (BEA SAGDP9N), covering all 50 states and the District of Columbia; (iii) G20 country-level real GDP from the Economist Intelligence Unit (EIU), 18 economies; (iv) OECD country-level real GDP from EIU, 33 economies; and (v) EU member state real GDP from EIU, 27 countries. At maximum, the pooled GVA donor universe contains 253 units.

GDHI donor sources. Each GDHI synthetic control is estimated against a corresponding income donor panel that uses the closest available international counterparts to the UK household-income series: (i) NUTS2-level household income in the secondary distribution account from Eurostat (`nama_10r_2hhinc`, indicator B6N, in EUR millions), covering up to 223 European sub-national regions; (ii) state-level disposable personal income from BEA SAINC51 (millions of dollars), covering all 50 states and DC; and (iii) at the country level, private household consumption expenditure at current prices (EIU series `CPRL`) for G20, OECD, and EU countries, used as the closest cross-country proxy where a national B.6g series is not available. At maximum, the pooled GDHI donor universe contains 322 units, reflecting the broader coverage of the Eurostat household income data relative to the NUTS2 GVA series.

Harmonisation to prevent double-counting. In both panels, where subnational data from the same country are included — for example, Germany represented through its NUTS2 regions — the national aggregate is excluded from the same donor-pool combination. This prevents the same economic activity from entering the synthetic control twice.

2.2 Synthetic Control Estimation

Let i index a treatment unit (a UK local authority, ITL region, or constituent country), $k \in \{\texttt{gva}, \texttt{gdhi\_b6}\}$ index the outcome series, and s index a donor-pool combination. All series are normalised to an index of 100 in 2000 before estimation. Let y_{it}^{k} denote the normalised outcome for the treated unit, let X_{st}^{k} denote the corresponding row vector of donor observations in pool s at time t, and let $\mathcal{T}_0$ denote the pre-treatment period.

For each triple (i, k, s), the isolated synthetic control estimator solves the simplex-constrained quadratic programme:

$$\widehat{\boldsymbol{w}}_{iks} = \arg \min_{\boldsymbol{w} \in \Delta^{J_s - 1}} \left(\boldsymbol{y}_i^{k,0} - X_s^{k,0} \boldsymbol{w} \right)' V_{iks} \left(\boldsymbol{y}_i^{k,0} - X_s^{k,0} \boldsymbol{w} \right), \tag{2.2}$$

where $\boldsymbol{y}_i^{k,0}$ and $X_s^{k,0}$ collect the pre-treatment observations, $\Delta^{J_s-1} = \{\boldsymbol{w} \geq \mathbf{0} : \mathbf{1}'\boldsymbol{w} = 1\}$ is the unit simplex over the J_s donors in pool s, and V_{iks} is a diagonal predictor-weight matrix chosen by a nested search that minimises the pre-treatment mean-squared projection error (MSPE). Note that $\boldsymbol{y}_i^{k,0}$ and $X_s^{k,0}$ contain the outcome series only; no additional pre-treatment covariates are included, because comparable covariate data at consistent quality are not available across the full range of donor economies and geographic tiers. The simplex constraint ensures the synthetic control is a convex combination of the donor units, which limits extrapolation and keeps the counterfactual interpretable as a weighted average of observed economies. Each series is estimated independently: the estimator is *isolated* in the sense that GVA and GDHI are fit separately for each treatment unit and donor pool, each with its own weight vector $\widehat{\boldsymbol{w}}_{iks}$.

The isolated construction is deliberate. A common-weight estimator would impose that the same donor combination reproduces both the production trajectory and the residence-based income trajectory of a UK region. That restriction is strong because GVA and GDHI-B6 capture different accounting concepts and may be governed by different spatial transmission mechanisms. Estimating the two outcomes separately preserves the GVA–GDHI-B6 wedge as a post-estimation object rather than mechanically suppressing it through the construction of the counterfactual.

The synthetic counterfactual for unit i under outcome k, donor pool s, and post-treatment period $t > T_0$ is

$$\hat{y}^k_{ist} = X^k_{st} \widehat{\boldsymbol{w}}_{iks}. \tag{2.3}$$

The corresponding *output gap* is

$$\hat{\delta}^k_{ist} = y^k_{it} - \hat{y}^k_{ist}. \tag{2.4}$$

A negative gap—actual below synthetic—corresponds to Brexit-related output or income loss. The average post-treatment gap, $\bar{\delta}^k_{is} = \frac{1}{|\mathcal{T}_1|} \sum_{t \in \mathcal{T}_1} \hat{\delta}^k_{ist}$, is our primary summary measure. As a descriptive comparison of production and income incidence, we also use the two estimated gaps jointly. Writing q for GVA and h for GDHI, define the post-estimation decoupling statistic

$$D_{it} = \hat{\delta}^{h,\text{ens}}_{it} - \hat{\delta}^{q,\text{ens}}_{it}, \qquad \bar{D}_i = \frac{1}{|\mathcal{T}_1|} \sum_{t \in \mathcal{T}_1} D_{it}. \tag{2.5}$$

Positive values of $\bar{D}_i$ indicate that household income is more resilient than local production: output falls relative to its counterfactual, but resident incomes are partly insulated by commuting, redistribution, or claims on production elsewhere. Negative values indicate that residence-based income falls more than local production, suggesting exposure through income channels — property income, self-employment earnings, commuting links to affected cities — not captured by the local production account. This is an interpretive comparison: $\bar{D}_i$ does not identify bilateral income flows but provides a descriptive diagnostic of how production and income losses are spatially aligned.

A natural alternative would be a joint synthetic control that matches both pre-treatment series using a common donor-weight vector. Where a common donor set is available, such an estimator solves

$$\widehat{\boldsymbol{w}}_{is}^{\text{joint}} = \arg\min_{\boldsymbol{w}\in\Delta^{J_s-1}} \left[\lambda_q \left(\boldsymbol{q}_i^0 - Q_s^0\boldsymbol{w}\right)' V_q \left(\boldsymbol{q}_i^0 - Q_s^0\boldsymbol{w}\right) + \lambda_h \left(\boldsymbol{h}_i^0 - H_s^0\boldsymbol{w}\right)' V_h \left(\boldsymbol{h}_i^0 - H_s^0\boldsymbol{w}\right)\right], \tag{2.6}$$

where Q_s^0 and H_s^0 are the pre-treatment donor matrices for GVA and GDHI-B6, and λ_q and λ_h scale the two outcomes to comparable importance. We do not use (2.6) as the primary estimator because it would make the GVA–GDHI-B6 wedge partly a maintained restriction rather than an object of study. The joint estimator is useful as a robustness benchmark, but the isolated estimator in (2.2) is the preferred baseline for studying whether production and residence-based income respond differently to Brexit.

Rationale and potential biases. The synthetic control estimator is well-suited to this setting: there is a single treated unit (the United Kingdom or a UK region) affected by a discrete policy shock, with no obvious control group from randomised assignment and a relatively long pre-treatment panel. The identifying assumption is that the donor economies used to construct the counterfactual share a common factor structure with the treated unit in the pre-treatment period, so that the projection of the treated trajectory onto the donor convex hull remains valid in the post-treatment period absent the shock.

This assumption is put under pressure by the specific economic narrative surrounding Brexit. The Leave campaign was partly a bet that the UK's future growth would come from deeper integration with faster-growing economies outside the EU—the United States, the Commonwealth, South and East Asia—rather than from a mature, relatively slow-growing European bloc. If that narrative were correct, a counterfactual constructed from EU peer economies would understate the trajectory the UK would have enjoyed, making the estimated Brexit cost appear too large. There is also a reverse concern: to the extent that the world has moved toward deglobalisation, trade fragmentation, and heightened geopolitical risk since 2016, some of the faster-growing G20 donors in the pool—China, India, Brazil—have themselves been subject to idiosyncratic structural forces that may not be good proxies for what the UK's trajectory would have been under continued EU membership. In both cases the composition of the donor pool is not a neutral choice, and the point estimate of the Brexit cost will depend on which economies are admitted into the comparison set.

We address this concern in two ways. First, rather than imposing a single donor pool, we enumerate the full set of $2^5 - 1 = 31$ combinations of the five base donor sources and estimate a separate synthetic control for each. We call each such estimate a *candidate*: one synthetic control for a given treatment unit, outcome series, and donor-pool specification. All candidates are then combined via the placebo-weighted ensemble (Section 2.4), which assigns higher weight to candidates whose donor pool predicted UK trajectories more accurately in a held-out pre-Brexit window. The spread of candidates across donor-pool specifications provides a direct sensitivity exercise: it shows how the headline gap estimate changes as the composition of the comparator group varies. Second, the ensemble weighting implicitly penalises donor pools that perform poorly before Brexit — for instance, pools dominated by rapidly industrialising economies whose own structural dynamics are a poor match for the UK — so those specifications receive low weight in the final estimate.

The GVA and GDHI-B6 series also carry different risks in this respect. A Brexit output cost that shows up strongly in GVA but modestly in GDHI-B6 may reflect a genuine distinction between where production has been lost (tradeable-sector hubs and ports) and where household welfare has fallen (which is damped by fiscal redistribution and commuting catchments). Conversely, if GDHI-B6 gaps are larger than GVA gaps in some areas, that may point to adverse income effects beyond what the local production account captures—for example through property and financial income channels that affect households directly. Comparing the two spatial signatures therefore provides diagnostic information about the economic mechanism linking Brexit to welfare outcomes.

2.3 Donor Pool Design

A synthetic control constructed from a single donor pool may be sensitive to idiosyncratic features of that pool: if one donor country experiences an episode specific to its own political economy, it will contaminate the counterfactual. To guard against this, we do not fix a single donor pool. Instead, we enumerate the full set of non-empty subsets of the five base datasets, yielding $2^5 - 1 = 31$ deterministic donor pool combinations for the annual analysis. These range from single-source pools (e.g. NUTS2 only, with 175 donors) to the comprehensive union of all five sources (253 donors). Appendix Table A2 lists every combination together with the corresponding donor pool size.

Within each combination, country-level and sub-national data are harmonised to prevent double-counting: if a country contributes regions at the NUTS2 level, its national aggregate is excluded from the same combination. The GVA series in the resulting donor matrices are normalised identically to the treated series.

For each triple (i, k, s) the estimator in (2.2) is solved, producing one synthetic counterfactual and one output gap per unit, per series, and per donor pool. Across the full geography ($\approx$ 360 LADs, 12 ITL1 regions, 4 constituent countries) and the two treatment windows, the production pipeline therefore estimates several thousand candidate synthetic controls for each outcome series.

The spread of gap estimates across the 31 donor-pool specifications constitutes a natural sensitivity exercise. Specifications that include only EU economies as donors may produce more negative gaps if the EU's trajectory exceeded what the UK would have achieved outside—but they may also produce less negative gaps if EU economies fared poorly in this period for their own reasons. Specifications that include the US states, OECD, or G20 introduce faster-growing economies as comparators and thus tend to make the counterfactual trajectory more demanding. The fact that we report the full ensemble across these alternatives, rather than a single selected specification, means the headline result does not depend on any particular view about which economies are the right comparators for the UK.

2.4 Placebo-Weighted Ensemble

Selecting the *best* synthetic control for each unit by minimising the pre-treatment fit metric—RMSPE or AAPE—carries a well-known risk: the selection criterion rewards in-sample fit and may therefore systematically favour specifications that overfit the pre-treatment period, leading to an upward bias in post-treatment prediction quality (Athey et al., 2019). We address this by constructing a *placebo-weighted ensemble* that scores each candidate on out-of-sample predictive performance using pseudo-treatment placebo windows rather than the actual treatment window.

Placebo windows. We exploit the fact that we can pose the same estimation problem for a set of false treatment dates, $\tau \in \{2010, 2011, 2012, 2013\}$, using data that pre-dates the actual Brexit vote. For each placebo year τ, we re-estimate (2.2) on the pre-placebo period $\mathcal{T}_0^\tau = \{2000, \ldots, \tau - 1\}$, treating year τ as though it were the treatment date. Since the true Brexit vote occurred in June 2016 and no comparable structural break is expected in 2010–2013, a good synthetic control should predict actual outcomes accurately in the *validation window* $\mathcal{V} = \{2012, 2013, 2014, 2015\}$ that follows the pseudo-treatment date. We use the post-2012 placebo window as our primary scoring window, which provides four validation years for every candidate.

Placebo scoring. For each candidate (i, k, s) — that is, each combination of treatment unit i, outcome series k, and donor-pool specification s — estimated in the placebo run, we compute the root mean square prediction error over the validation window:

$$\mathrm{RMSE}_{iks}^{\mathrm{placebo}} = \sqrt{\frac{1}{|\mathcal{V}|} \sum_{t \in \mathcal{V}} \left(y_{it}^{k} - \hat{y}_{ist}^{k,\tau} \right)^2}, \tag{2.7}$$

where y_{it}^{k} is the actual outcome for unit i and outcome k at time t, and $\hat{y}_{ist}^{k,\tau}$ is the counterfactual from the placebo estimation using pseudo-treatment year τ. Candidates with a lower $\mathrm{RMSE}_{iks}^{\mathrm{placebo}}$ predicted outcomes more accurately in the held-out window; these candidates receive higher weight in the ensemble.

Ensemble weights. Within each ensemble group $g = (i, k)$—all candidates for a given treatment unit and outcome series—we rank candidates by $\mathrm{RMSE}_{iks}^{\mathrm{placebo}}$ in ascending order (rank 1 = best predictive performance) and assign weights using a softmax transformation of the rank:

$$\tilde{w}_{iks} = \frac{\exp(-\kappa \cdot r_{iks})}{\sum_{s' \in \mathcal{S}_{ik}} \exp(-\kappa \cdot r_{iks'})}, \tag{2.8}$$

where r_{iks} is the placebo-RMSE rank of candidate (i, k, s) within group g, and $\kappa = 0.25$ controls how steeply weight falls with rank. By construction the weights sum to one and are non-negative. The softmax converts a ranking into a smooth weight schedule: the best-predicting donor pool (rank 1) receives the highest weight; the worst-predicting pool (rank 31) receives near-zero weight. A small κ produces approximately equal weights (pure averaging); a large κ concentrates almost all mass on the top-ranked candidate. The choice $\kappa = 0.25$ reflects a moderate degree of concentration, giving meaningful weight to the best several candidates while discounting poorly fitting specifications.

Placebo-weighted counterfactual. The primary counterfactual used in this paper is the placebo-weighted ensemble:

$$\hat{y}_{it}^{k,\mathrm{ens}} = \sum_{s \in \mathcal{S}_{ik}} \tilde{w}_{iks} \cdot \hat{y}_{ist}^{k}, \tag{2.9}$$

where $\hat{y}^k_{ist}$ is now the real-window counterfactual from (2.3), evaluated for the actual post-2016 or post-2020 period. Crucially, the weights $\tilde{w}_{iks}$ are fixed from the placebo scoring step using only pre-Brexit data, and they do not depend on the real post-treatment outcomes. This means the weights cannot be chosen to generate a larger or smaller Brexit gap: since they were assigned before looking at the post-2016 or post-2020 period, no data snooping is possible in the construction of the counterfactual. The corresponding ensemble gap is

$$\hat{\delta}^{k,\text{ens}}_{it} = y^k_{it} - \hat{y}^{k,\text{ens}}_{it}. \tag{2.10}$$

2.5 Inference via Restricted Placebo Units

Standard significance testing for the synthetic control is complicated at the aggregate UK or ITL1 level because there is only one treated unit and the distribution of the gap statistic under the null is unavailable in closed form. The intuition behind the inference approach is simple: if Brexit had no effect, the UK region's estimated gap should look no different from the gaps obtained by pretending untreated donor economies were themselves treated. A large actual gap — relative to the distribution of such placebo gaps — is evidence of a genuine Brexit effect.

We implement this via *restricted placebo unit* runs, following Abadie et al. (2010). The procedure proceeds in four steps. First, for each treated unit (i, k), we select the best-performing donor pool configuration, denoted $\hat{s}_{ik}$ and defined as the donor-pool specification that received the highest ensemble weight in Section 2.4. Second, for each donor unit j in that pool, we re-estimate the synthetic control as though j were the treated unit while keeping the remaining donors (excluding j) as the comparison group. Third, we compute the average post-treatment gap for each such placebo unit j, producing a distribution $\{\bar{\delta}^k_j\}_{j \in \mathcal{D}}$ of placebo gaps under the null. Fourth, we compare the actual ensemble gap $\hat{\delta}^{k,\text{ens}}_{it}$ to this distribution: the one-sided p-value is the share of donor units whose average placebo gap is as extreme or more extreme than the actual gap. This provides a non-parametric test of the hypothesis that Brexit had no effect on the region's economy.

2.6 Two Treatment Windows

We estimate the Brexit output gap under two distinct treatment codings, which capture different economic channels and are not interchangeable.

Post-2016 window. The first window codes the treatment as beginning in 2016, the year of the referendum. The pre-treatment period is $\mathcal{T}_0 = \{2000, \ldots, 2015\}$ (16 years), and the post-treatment period runs from 2017 through 2023. The motivation for this coding is that the economic effects of Brexit began well before the legal exit. The referendum result triggered an immediate depreciation of sterling of approximately 10–15%, a persistent rise in trade policy uncertainty (Bloom et al., 2019; Baker et al., 2020), and a reorganisation of supply chains and investment plans in anticipation of changed trading arrangements. Several studies have documented that a substantial fraction of the aggregate cost of Brexit had already materialised by the time the UK formally left the EU in January 2020 (Born et al., 2019; Hassan et al., 2024; Dhingra and Sampson, 2022). The post-2016 window is therefore designed to capture the full economic cost of Brexit, including its uncertainty and anticipatory channels.

Post-2020 window. The second window codes the treatment as beginning in 2020, the year of legal exit from the EU and the year of the COVID-19 pandemic. The pre-treatment period is $\mathcal{T}_0 = \{2000, \ldots, 2019\}$ (20 years), providing a longer and therefore typically more accurate calibration period for the synthetic control. The post-treatment period runs from 2020 through 2023. The motivation for studying this window is to isolate the effects of the actual institutional changes—the end of free movement of goods, services, capital, and people—from the earlier uncertainty channel. Because both the legal exit and the pandemic onset coincide in 2020, it is important to note that synthetic controls constructed for this window will conflate the two. However, since both shocks hit the UK and most donor economies in 2020, the synthetic control that best tracks the UK's pre-2020 trajectory should, under the identifying assumption, provide a valid counterfactual for the pandemic-only scenario that prevails in donor economies outside the UK. The excess gap over and above what donor-country trajectories imply is therefore attributable to the combination of Brexit and any idiosyncratic UK pandemic response, rather than to the pandemic per se.

Comparing results across the two windows also provides a disciplined diagnostic. If the post-2016 gap is large while the post-2020 gap is comparatively smaller, the evidence supports an important role for the uncertainty channel in driving aggregate costs. If the gaps are similar in magnitude, the dominant mechanism is likely the direct trade and regulatory disruption that took effect after January 2020. In either case, the two estimates together bracket the plausible range of the Brexit output cost.

3 The Spatial Distribution of Brexit's Economic Cost

3.1 Country- and Region-Level Estimates

We begin with the aggregate UK-level evidence from the quarterly synthetic control, which provides the highest-frequency picture of when Brexit effects began to materialise. Figure 1 displays the actual UK real GDP series against the full set of seven donor-pool synthetic counterfactuals, with the placebo-weighted ensemble highlighted as the primary estimate and the in-sample best-fitting specification (G20 + OECD donor pool) shown for comparison. All series are expressed as an index normalised to 100 at 2015 Q4, the final pre-treatment quarter; the log gap multiplied by 100 approximates the percentage cost in percentage points.

Two features of Figure 1 stand out. First, the synthetic series track actual UK

GDP closely in the pre-treatment period: across all seven donor pools, pre-period deviations are small and show no systematic trend, supporting the parallel-trends assumption on which counterfactual validity rests. Second, the divergence between actual and synthetic GDP opens immediately after the 2016 referendum and widens persistently through the sample. The loss area—shaded in pink—is substantial and grows over time, with the placebo-weighted ensemble suggesting a shortfall of approximately 3–5 percentage points by the end of the sample. The spread of the faint background lines, each representing a distinct donor-pool specification, constitutes a natural sensitivity exercise: the qualitative direction of the loss is consistent across all specifications, while the magnitude varies with the composition of the comparator group.

Notably, the best in-sample fit (G20 + OECD pool, lowest pre-period RMSPE) and the best out-of-sample fit (EU + G20 pool, top placebo rank) point to similar post-2016 trajectories, but the placebo-weighted ensemble lies closer to the out-of-sample best, reflecting that specifications favoured by pre-period fit do not necessarily generalise best beyond the fitting window.

We next turn to the annual constituent-country estimates. The Online Appendix presents the full annual placebo-weighted ensemble counterfactual paths for England, Scotland, Wales, and Northern Ireland under both the post-2016 and post-2020 codings, separately for GVA and GDHI-B6. The synthetic series track the actual series closely in the pre-treatment period, while notable post-treatment gaps emerge in England, Scotland, and Wales across both windows. Northern Ireland stands apart: its gap is negligible or positive, reflecting its effective continuation within the EU customs union under the EU-UK Trade and Cooperation Agreement. Scotland's large and robust losses merit particular note: they are partly driven by the distinct economic geography of Aberdeen and Falkirk, which host a disproportionate share of the UK's oil and gas industry.[11]

[11] Aberdeen is the centre of North Sea oil and gas extraction and its associated professional and technical services sector; Falkirk hosts major downstream oil refining and petrochemicals infrastruc-

Table 1 summarises the annual synthetic-control evidence across the two outcome concepts for the UK, its constituent countries, and the ITL1 regions. GVA is a real chained-volume production measure and GDHI is the nominal income overall (balance of secondary distribution) that reflects disposable household income after taxes and transfers. As discussed in Section 2.1, the two measures need not move in lock-step: GVA reflects where production occurs while GDHI reflects where incomes are received after redistribution. For each concept, the table reports the 2023 loss in money-metric terms, the corresponding 2023 percentage gap, a 2023 per-capita loss, the binary loser/winner coding, and a robustness measure, separately for the post-2016 and post-2020 windows. All monetary and percentage figures are point estimates for 2023; Appendix Table 3 reports the same table with all metrics expressed as averages over the full post-treatment window.

The Online Appendix presents the annual actual and synthetic-control paths for all twelve ITL1 regional GVA series under both treatment windows. Across regions, the synthetic series track the actual series closely in the pre-treatment period, while large post-treatment gaps emerge particularly in London and the Midlands. The spread of the candidate counterfactuals across the 31 donor-pool specifications is notably wider for London than for Northern regions, reflecting the greater sensitivity of London's counterfactual to whether fast-growing international economies are included in the donor pool.

Aggregation Consistency

As an internal validation, Table 2 aggregates the independently estimated lower-level gaps back to the country level and compares them against the directly estimated country-level synthetic controls. This exercise is more than a robustness

ture. Brexit created significant friction for these industries through regulatory divergence with the EU, increased barriers to the cross-border deployment of specialised labour and equipment, and tightened access to EU energy and trading arrangements under the TCA. The concentration of these sectors in Scottish sub-regions amplifies Scotland's aggregate GVA gap relative to its population-weighted share of UK output.

check: its interpretation depends critically on understanding what the synthetic control estimator can and cannot guarantee.

Why aggregation consistency matters. Each synthetic control in this paper is estimated independently for every spatial unit—LADs, ITL3 sub-regions, ITL2 regions, ITL1 regions, and constituent countries each receive their own counterfactual, calibrated to that unit's pre-treatment trajectory, without any cross-scale aggregation constraint imposed during estimation. This independence opens the door to several potential sources of inconsistency. First, *SUTVA violations*: if Brexit-induced losses in one region are offset by gains in adjacent regions through supply-chain reallocation, labour mobility, or trade diversion, the sum of sub-national gaps need not equal the aggregate gap even if all individual estimates are correct. Second, *measurement noise*: local authority GVA is estimated with greater uncertainty and subject to larger revisions than national aggregates; aggregating noisy lower-level estimates can amplify or attenuate discrepancies relative to the directly estimated aggregate. Third, *cross-value-chain spillovers*: production losses in one district propagate through input-output linkages to functionally connected districts in ways not captured when each unit is estimated in isolation. Fourth, *residence-workplace mismatches*: GDHI is recorded where workers live, GVA where they work, and these geographies diverge most sharply in large commuting zones such as London; this wedge can create mechanical differences between the two measures when aggregated at different spatial tiers.

Despite these potential sources of divergence, aggregation consistency is highly informative *precisely because* no aggregation constraint was imposed. Finding that independently estimated sub-national gaps, weighted by population and summed, closely match the directly estimated country-level gaps constitutes strong evidence that the synthetic control approach is capturing a spatially coherent economic signal rather than a measurement artefact confined to any single geographic tier. It rules out the hypothesis that the apparent Brexit cost is an artefact of noise in a particular data series, and it validates the integrity of the data at each spatial scale. The independence of the estimation exercises makes this consistency more, not less, compelling as validation: any alignment across spatial tiers must reflect a genuine economic phenomenon rather than an algebraic identity.

Table 2 documents that this consistency holds reasonably well across both outcome series and both treatment windows. The summed ITL1 GVA gaps approximate the directly estimated country-level GVA gap to within a reasonable margin, and similar consistency holds when aggregating from ITL2 and ITL3 to ITL1. The result is reassuring: the synthetic control exercise produces estimates that are internally coherent across spatial scales despite having been constructed without any cross-scale constraint.

3.2 Comparison with Existing Country-Level Estimates

A natural benchmark for our estimates is the existing body of country-level evidence on the aggregate economic cost of Brexit. This literature spans a wide range of methodological approaches and data vintages, but converges on a consistent conclusion: Brexit has imposed a substantial and growing cost on the UK economy.

Ex-ante predictions vs. ex-post estimates. It is important to distinguish between estimates produced *before* the referendum or before the UK's legal exit—which necessarily relied on structural or partial-equilibrium models to simulate the consequences of counterfactual trade arrangements—and *ex-post* estimates that compare the UK's actual trajectory against a data-driven counterfactual. Pre-referendum forecasts from HM Treasury (2016) projected long-run GDP losses of 3–6% under a free-trade-agreement scenario and 5–8% under a WTO scenario. The OECD and IMF produced comparable estimates. Crucially, these were predictions for scenarios that had not yet materialised.

The actual Brexit implementation—the EU-UK Trade and Cooperation Agreement (TCA) signed in December 2020—represents a *very hard Brexit* relative to the alternatives originally envisaged. Tariffs were avoided on goods, but the TCA imposed significant non-tariff barriers on services, ended free movement of people and capital, and excluded the UK from single-market benefits in financial and professional services. This is notably harder than even the "WTO scenario" in some early forecasts, which did not anticipate the depth of services disruption. Indeed, Dhingra and Sampson (2022) find that the terms of the TCA align more closely with a hard Brexit than the moderate versions that featured in pre-referendum analyses.

A growing consensus on large costs. Early ex-post synthetic control estimates—using data only through 2017 or 2018—necessarily captured only the referendum-shock and uncertainty channel. Born et al. (2019) estimated a loss of approximately 1.7–2.5% of GDP by 2018 using OECD donor economies. Hassan et al. (2024) exploit variation in firm-level trade exposure and arrive at an aggregate cost of around 2% of GDP by 2019. These earlier estimates reflected a period when the bulk of the loss was still unfolding.

As more post-Brexit data have accumulated and the trade barriers of the TCA have taken effect, the estimated costs have widened substantially. The Centre for European Reform's rolling doppelganger estimates, which adopt a method analogous to Born et al. (2019) but with updated data, place the output gap at 5–6% as of 2023 (Springford, 2024). Goldman Sachs's proprietary estimates similarly converge on approximately 5% (2024). Bloom et al. (2026), exploiting the discontinuity created by the TCA for firms with varying trade exposures, estimate a loss of approximately 7% of output, consistent with the upper range of ex-post synthetic control estimates.

Optimistic pre-referendum voices—notably proponents of the view that deregulation and global free trade would generate large growth dividends (Minford, 2019)—have not been validated by the data. Those projections, which in some variants implied GDP *gains* from Brexit, have been comprehensively refuted by the subsequent evidence.

The compounding of evidence over time. Figure 2 illustrates this evolution of the evidence. The left panel organises estimates by year of publication: earlier studies cluster around lower (less negative) estimates, while more recent work — benefiting from longer post-treatment panels and the full materialisation of TCA trade barriers — converges on larger losses. This pattern is not a sign of publication bias; rather, it reflects the fact that the economic cost of Brexit was concentrated in the uncertainty period after the referendum and then deepened as the hard terms of exit were implemented. The right panel documents the remarkable consensus across methodological families: synthetic control, general equilibrium, event-study, and reduced-form approaches all point in the same direction, differing primarily in their quantitative magnitudes.

The headline placebo-weighted ensemble estimates of 3–5 percentage points of quarterly real GDP by 2025 are broadly consistent with the central range of extant ex-post estimates. The annual GVA evidence, which covers the full post-TCA period through 2023, yields a UK-level gap of approximately 7–8%, toward the upper end of the range, reflecting both the longer sample period and the fact that annual GVA captures sectors—agriculture, services, regional production—that may not be fully represented in aggregate GDP series.

A particularly important distinction concerns *hard* versus *soft* Brexit. Many early estimates were predicated on a softer form of exit—remaining in the European Economic Area or a comprehensive customs union—that would have preserved much of the single-market access. The actual TCA is considerably more disruptive: it imposes full third-country status on the UK for services, introduces customs checks and rules-of-origin requirements for goods, and eliminates the mutual recognition of professional qualifications. Had a softer Brexit been implemented, the output cost would plausibly have been at the lower end of the pre-referendum forecasts. The large ex-post estimates we and others document are therefore partly attributable to the specific form of Brexit that was chosen.

Figure 2 contextualises our aggregate UK quarterly estimates relative to this literature by presenting two complementary views of the existing evidence. The left panel organises country-level Brexit cost estimates by the year of publication. The right panel classifies estimates by methodological family—synthetic control, general equilibrium, event study, and reduced-form—illustrating the degree of consensus across approaches. Together, the two panels establish that, while point estimates vary with data vintage, methodology, and donor-pool composition, the broad direction of the evidence is unambiguous.

3.3 Distribution of Costs across Districts

We next turn to the district-level estimates, which form the main empirical focus of this paper. Averaging the annual output or income gap for each district over the post-treatment window yields the empirical distribution shown in Figure 3. Panels A and B report the LAD-level distributions for GVA and GDHI-B6 respectively; Panel C plots their joint density. Negative values indicate Brexit-related losses relative to the placebo-weighted ensemble counterfactual. Both distributions are skewed to the left, confirming that the typical district experiences an output or income shortfall. Figure 4 maps the geographic distribution of these losses using hex cartograms at ITL2, ITL3, and LAD resolution, classifying each unit as a loser, winner, or unclear based on the direction of its post-2016 GVA gap across donor-pool specifications.

The cross-sectional pattern of losses is not randomly distributed. The sections that follow first interpret the joint GVA–GDHI-B6 distribution as evidence on latent regional-flow incidence, then explore three systematic correlates of the cost distribution: trade exposure, the 2016 Leave vote, and the political economy narrative of austerity and immigration concerns that shaped the Brexit campaign.

The GVA–GDHI-B6 wedge and latent regional-flow incidence

Panel C of Figure 3 provides the first diagnostic of the relationship between the production and income gaps. If regional economies were self-contained and production losses passed through one-for-one into resident incomes, the two measures would lie close to a common diagonal. Deviations from that benchmark are informative: they reveal that production, residence, and household welfare are connected through spatial flows not directly observed in standard regional accounts.

We interpret the decoupling statistic $\bar{D}_i$ in (2.5) as a reduced-form measure of this latent flow incidence. Regions with negative GVA gaps but comparatively muted GDHI-B6 gaps are places where the local production shock is partly exported to the commuter belt, absorbed by fiscal redistribution, or offset by residents' claims on production elsewhere. Regions where GDHI-B6 falls more than GVA are places where resident households are exposed to income losses generated outside the local production account — for example through commuting links to affected production centres, self-employment and property-income channels, or the weakening of claims on externally integrated sectors.

This is precisely the class of object that a full Disaggregated Economic Accounts system would measure directly. DEA would observe the bilateral matrices connecting consumer cells, producer cells, government, and the rest of the world. Our synthetic-control design observes only the marginal regional accounts and their counterfactual gaps. Brexit nevertheless acts as a revealing shock: by changing the external integration regime, it makes visible the internal architecture through which production losses are transmitted into residence-based income losses. The GVA–GDHI-B6 wedge is therefore not only a measurement artefact but a source of information about the openness, interdependence, and spatial unbundling of regional economies.

Trade Exposure and the Cost of Brexit

The first natural correlate of Brexit-induced losses is baseline exposure to trade with the EU. Brexit erected a series of non-tariff barriers, regulatory frictions, and customs checks that directly raised the cost of goods and services trade with the UK's largest trading partner. In turn, places more integrated into EU supply chains and export markets should have experienced larger GVA losses.

Panel B of Figure 5 plots the average post-2016 GVA output loss against overall trade openness — the ratio of total goods and services exports plus imports to GVA — at the ITL3 level. The relationship is positive and statistically significant: more trade-open places experienced larger realised losses. This is documented more formally in the regression analysis in Appendix Section B. The results are consistent with Brexit operating primarily through the trade channel, as standard open-economy models would predict.

The Brexit Vote and Realised Costs

A second systematic dimension is the correlation between an area's 2016 Leave vote share and its subsequent realised economic cost. The Remain campaign argued that costs would be widely distributed; the Leave campaign argued that any costs would fall primarily on the London-centric establishments that had benefited most from EU integration and least on the "left-behind" communities that had borne the costs of globalisation and austerity. The data permit a direct test: if Leave voters in economically declining areas were correct, we would expect areas with higher Leave support to show smaller (or positive) GVA gaps relative to the synthetic control.

Figure 5 documents both cross-sectional mechanisms in a single exhibit. Panel A plots the 2016 Leave vote share (horizontal axis) against the average post-2016 placebo-weighted ensemble GVA output loss (vertical axis) across 169 ITL3 subregions. Panel B replicates the overall trade openness scatter at the ITL3 level for comparison.

The two panels reveal a striking asymmetry between the political economy of the Brexit vote and the distribution of its economic consequences. Panel A shows a weakly *negative* correlation ($r \approx -0.1$ to -0.2): areas that voted more heavily for Leave subsequently experienced *smaller* output losses on average, while areas that voted to Remain—notably London, the university cities, and most of Scotland—experienced the largest gaps. This pattern is the empirical foundation of the "levelling up by levelling down" interpretation. It also implies that the economic cost of Brexit was not, in the first instance, borne by the communities that generated the political mandate for it.

Panel B, by contrast, shows a positive and statistically significant correlation between EU trade exposure and GVA losses: more trade-open places—those with deeper integration into EU supply chains and export markets—experienced systematically larger output shortfalls. Naturally, the degree of integration of UK areas with EU supply chain is only a crude measure that is likely to understate the true realized level of integration due to the onward linkages and UK domestic supply chain linkages with firms that are trading with the EU themselves trading with other firms inside the UK market. The direct trade exposure measure thus understates the likely true and indirect effect, yet, together, the two panels trace the fault line between political agency and economic incidence: support for Brexit was concentrated in economically marginalised, less trade-exposed communities that ultimately absorbed relatively smaller losses, while the most internationally integrated areas—which voted against Brexit—bore the brunt of the trade and regulatory disruption.

3.3.1 The Brexit Narrative: Austerity, Immigration, and Realised Costs

The third dimension concerns the political economy context in which the Brexit vote took place. Support for Leave was not uniformly distributed across the economically marginalised, nor was it purely a response to long-term economic decline. Becker et al. (2016) document that the Leave vote was strongly predicted by a combination of socio-economic deprivation (associated with austerity) and cultural anxieties about immigration. This section asks directly: do areas characterised by these specific pre-referendum narratives show a distinctive post-referendum economic cost signature?

Figure 6 compresses the austerity–immigration narrative into a single 2×2 classification at the local authority district level. Areas are classified by whether their pre-referendum austerity exposure and immigration-narrative salience are above or below the national median. If the realised economic costs of Brexit were concentrated in the places whose political mobilisation drove the Leave vote, the high-austerity/high-immigration quadrant should display the largest negative GVA gaps.

The 2×2 yields a striking null result. The high-austerity/high-immigration quadrant — the archetypal Leave-voting heartland — does not display the largest realised GVA losses. Instead, the deepest negative gaps are concentrated in the low-austerity/low-immigration quadrant, which corresponds broadly to London, the South East, and Scottish cities. The implication is direct: the economic cost of Brexit was primarily borne by relatively prosperous, EU-integrated areas that were politically opposed to Brexit, not by the communities whose narrative of decline fuelled the Leave campaign. This pattern fully corroborates the levelling-up-by-levelling-down finding established in Section 3.1.

4 Conclusion

This paper has estimated the regional economic cost of Brexit across the full distribution of UK local authorities, ITL regions, and constituent countries, exploiting a synthetic control approach applied to two complementary outcome series — real gross value added (GVA) and nominal gross disposable household income (GDHI-B6) — under both a post-2016 and a post-2020 treatment window. The primary counterfactual is a placebo-weighted ensemble that scores candidate donor-pool specifications on out-of-sample predictive performance, guarding against in-sample overfitting.

Three findings stand out. First, the economic cost of Brexit is large and near-universal. The vast majority of UK districts show a negative gap between actual and synthetic series for both GVA and GDHI-B6 in both treatment windows. The aggregate UK output loss is broadly consistent with existing country-level estimates (Born et al., 2019; Hassan et al., 2024), and GVA losses accumulate consistently from 2017 onwards, supporting the post-2016 coding as the relevant starting point for capturing the full cost.

Second, the distribution of losses is strongly heterogeneous. Northern Ireland stands apart as the only constituent country with near-zero or positive gaps, reflecting its effective continuation within the EU customs union under the EU-UK Trade and Cooperation Agreement. Among the remaining regions, losses are disproportionately concentrated in initially prosperous areas: London, the South East, and Scotland experience the largest gaps, while areas in the Midlands and the North — which were economically relatively weaker before 2016 — tend to see smaller losses. This spatial pattern implies that Brexit has compressed the UK's regional inequality distribution, constituting a form of "levelling up by levelling down": regional gaps narrow not because lagging areas improve but because leading areas fall further.

Third, comparing GVA and GDHI-B6 gaps reveals that the two series do not move in lock-step. In regions with large net commuting inflows, the production loss (GVA) can outrun the local income loss (GDHI-B6); in regions more dependent on property income, self-employment income, or commuting links to affected production centres, the relationship reverses. This divergence confirms that a single-series analysis would present an incomplete picture of the welfare cost of Brexit at the subnational level. More broadly, the GVA–GDHI-B6 wedge can be read as a reduced-form trace of the latent flow network connecting production, residence, income, and redistribution across regions.

The cross-sectional pattern of output losses is not systematically explained by the socio-economic characteristics that predicted support for Leave in 2016, by post-referendum levelling-up transfers, or by COVID-19 mortality. These null results support the validity of the synthetic control design. The finding that places that voted most strongly for Leave are not the places that have suffered most economically from Brexit reinforces the "levelling up by levelling down" interpretation—the regions that bore the economic cost are not, by and large, the regions that generated the political momentum for Brexit.

References

Abadie, A., A. Diamond, Hainmueller, and Jens (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.

Abadie, A. and J. Gardeazabal (2003). The economic costs of conflict: A case study of the Basque Country. *American Economic Review 93*(1), 113–132.

Andersen, A. L., K. Huber, N. Johannesen, L. Straub, and E. T. Vestergaard (2026). Disaggregated economic accounts. *The Quarterly Journal of Economics 141*(2), 1005–1075.

Athey, S., M. Bayati, G. Imbens, and Z. Qu (2019). Ensemble Methods for Causal Effects in Panel Data Settings. *AEA Papers and Proceedings 109*, 65–70.

Autor, D. (2024). Trade shocks and labor markets: New evidence. Working Paper.

Baker, S. R., R. A. Farrokhnia, S. Meyer, M. Pagel, and C. Yannelis (2020). How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic. *NBER Workigng Paper 26949*.

Bakker, J. D., N. Datta, and J. De Lyon (2023). Non-tariff barriers and consumer prices: Evidence from Brexit. CEP Discussion Paper 1888, Centre for Economic Performance, London School of Economics.

Becker, S. O., T. Fetzer, and D. Novy (2016). Who Voted for Brexit? A Comprehensive District-Level Analysis. CAGE Online Working Paper Series 305, Competitive Advantage in the Global Economy (CAGE).

Bloom, N., P. Bunn, S. Chen, P. Mizen, P. Smietanka, and G. Thwaites (2026). The Economic Impact of Brexit. Technical Report 34459, National Bureau of Economic Research. Exploits TCA discontinuity using firm-level trade exposure variation; estimates approximately 7% output loss.

Bloom, N., P. Bunn, P. Mizen, P. Smietanka, and G. Thwaites (2025). The economic impact of brexit. NBER Working Paper 34459, National Bureau of Economic

Research.

Bloom, N., J. Van Reenen, and H. Williams (2019). A Toolkit of Policies to Promote Innovation. *Journal of Economic Perspectives 33*(3), 163–184.

Born, B., G. J. Müller, M. Schularick, and P. Sedláček (2019). The Costs of Economic Nationalism: Evidence from the Brexit Experiment. *Economic Journal 129*(10), 2722–2744.

Breinlich, H., E. Leromain, D. Novy, and T. Sampson (2022). The Brexit vote, inflation and UK living standards. *International Economic Review 63*(1), 63–93.

Carreras, M., Y. I. Carreras, and S. Bowler (2019). Long-Term Economic Distress, Cultural Backlash, and Support for Brexit.

Dhingra, S. and T. Sampson (2022). Expecting Brexit. *Annual Review of Economics 14*, 495–519.

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2020). The Return to Protectionism. *Quarterly Journal of Economics 135*(1), 1–55.

Fetzer, T. and S. Wang (2020). Measuring the regional economic cost of brexit: Evidence up to 2019. *CEPR Discussion Paper No. DP15051*.

Goldman Sachs (2024). Brexit Macro Estimate: UK Real GDP Lagged Peers by Around 5%. Technical report, Goldman Sachs Research. Research note reported in press; estimates UK real GDP approximately 5% below counterfactual trajectory due to Brexit.

Grassi, B. (2024). Brexit and regional economic divergence. Working Paper.

Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun (2024). The global impact of Brexit uncertainty. *Journal of Finance 79*(1), 413–458.

HM Treasury (2016). The Long-Term Economic Impact of EU Membership and the Alternatives. Technical Report 9250, HM Government, London. Pre-referendum structural model forecasts; projects long-run GDP loss of 3–6% under FTA and 5–8% under WTO scenario.

Javorcik, B., B. Kett, K. Stapleton, and L. O'Kane (2026). Unravelling deep integration: Local labour market effects of the brexit vote. *Journal of the European Economic Association 24*(2), 429–475. First published online 29 September 2025.

Minford, P. (2019). The Effects of Brexit on the UK Economy. *Cardiff Economics Working Papers E2019*(1). Pro-Brexit estimate; projects economic gains from deregulation and global free trade following EU exit.

Novy, D. (2024). Brexit and trade disintegration: Updated estimates. Working Paper.

Pickett, R., J. Hill, and S. Cowan (2026). Synthetic control misconceptions: Recommendations for practice. *arXiv preprint arXiv:2603.19211*.

Rodríguez-Pose, A. (2018). The revenge of the places that don't matter (and what to do about it). *Cambridge Journal of Regions, Economy and Society 11*(1), 189–209.

Sampson, T. (2017). Brexit: The economics of international disintegration. *Journal of Economic Perspectives 31*(4), 163–184.

Springford, J. (2024). The Cost of Brexit: Rolling Doppelgänger Estimates. *Centre for European Reform*. Ongoing series of doppelgänger synthetic control estimates; places UK output gap at 5–6% below counterfactual as of 2023.

Stansbury, A., D. Turner, and E. Balls (2023). Tackling the UK's regional economic inequality: Binding constraints and avenues for policy intervention. *Contemporary Social Science 18*(3–4), 318–356.

Steinberg, J. B. (2019). Brexit and the macroeconomic impact of trade policy uncertainty. *Journal of International Economics 117*, 175–195.

Wiedemann, A. (2024). Redistributive politics under spatial inequality. *Journal of Politics 86*(3), 1013–1030.

YouGov (2023). Brexit: How do you feel about Brexit? opinion tracker. YouGov, `https://yougov.co.uk/topics/politics/trackers/brexit`.

YouGov (2026). What do britons think of brexit, 10 years since the referendum? Survey of 2,114 GB adults; fieldwork 2–3 June 2026.

Figures

Figure 1: UK quarterly GDP: synthetic control estimates across donor pools, post-2016

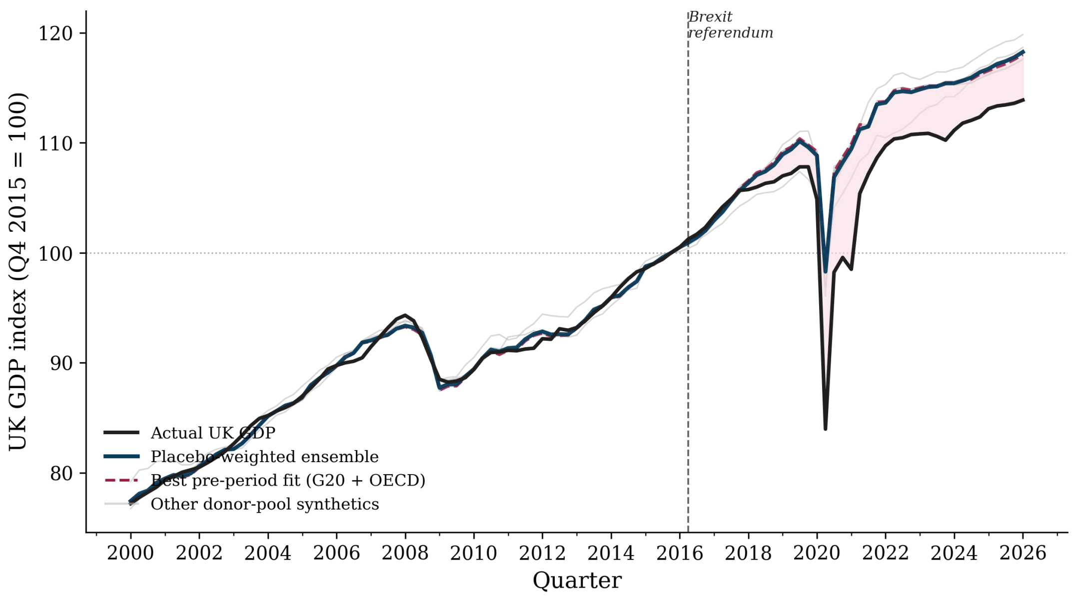

Note: The figure plots quarterly UK real GDP (inflation-adjusted, 2019 prices) against synthetic control counterfactuals from seven donor-pool specifications, expressed as an index normalised to 100 in 2015 Q4 (the final pre-referendum quarter). Faint grey lines are the candidate synthetic series for individual donor-pool combinations (EU-27, G20, OECD, and all pairwise and triple unions of these sources). The dark navy line is the placebo-weighted ensemble counterfactual (primary estimate); the dashed maroon line is the in-sample best-fitting specification (G20 + OECD, lowest pre-period RMSPE). The solid black line is actual UK real GDP. The pink shaded area marks quarters where actual GDP falls below the ensemble counterfactual. The vertical dashed line marks 2016 Q2 (the referendum quarter). Note that the normalisation used for estimation (index = 100 in 2000) differs from the display normalisation used here (index = 100 in 2015 Q4); the choice of display base year does not affect the gap estimates.

Figure 2: Brexit impact estimates in the literature: by publication year and by methodological family

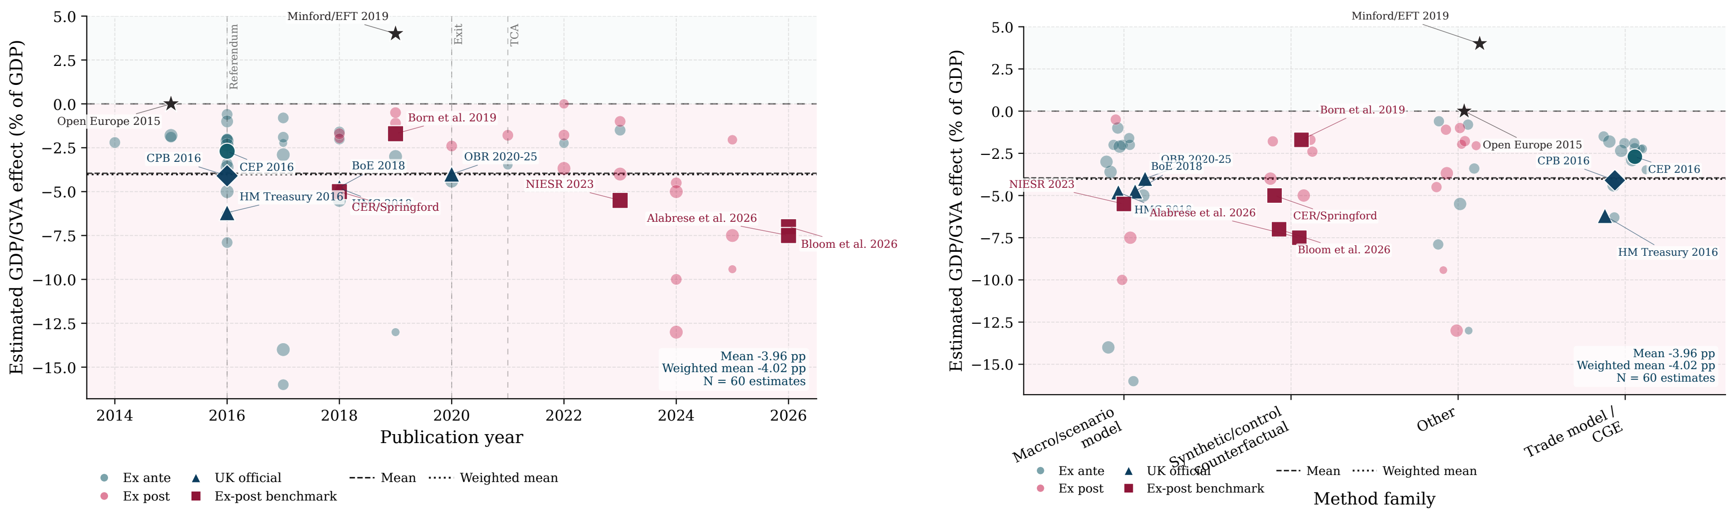

Note: Each point is a published estimate of the aggregate Brexit cost, expressed as a percentage shortfall of UK GDP or GVA relative to a counterfactual without Brexit; where a study reports a range, the midpoint is plotted. The *left panel* organises estimates by year of publication; labelled points identify prominent studies. The *right panel* organises the same estimates by methodological family: *synthetic control* constructs a data-driven counterfactual from a weighted average of comparator economies; *general equilibrium* uses structural trade models to simulate counterfactual integration scenarios; *reduced-form* exploits firm- or sector-level variation in trade exposure; *event study* uses high-frequency market or announcement data around Brexit-related events. Sources: authors' compilation; see text for citations.

Figure 3: Empirical distribution of the average Brexit cost across districts: LAD coding, post-2016

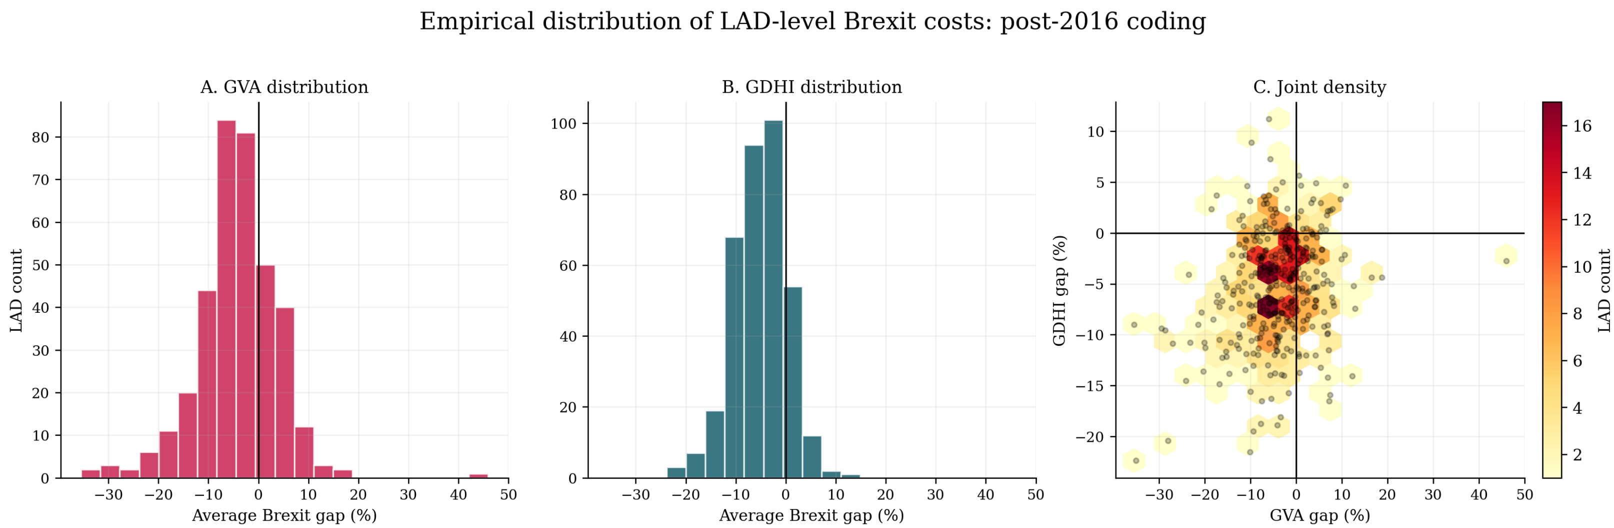

Note: Panel A reports the LAD-level empirical distribution of average annual GVA gaps under the post-2016 coding; Panel B reports the corresponding GDHI-B6 distribution; Panel C plots the joint density of the two measures. GVA is a real chained-volume production measure and GDHI-B6 (balance of secondary distribution of income) is nominal. Negative values indicate that the actual series lies below the placebo-weighted ensemble counterfactual and therefore correspond to Brexit-related losses.

Figure 4: Geographic distribution of Brexit GVA losses: ITL2, ITL3, and LAD hex cartograms

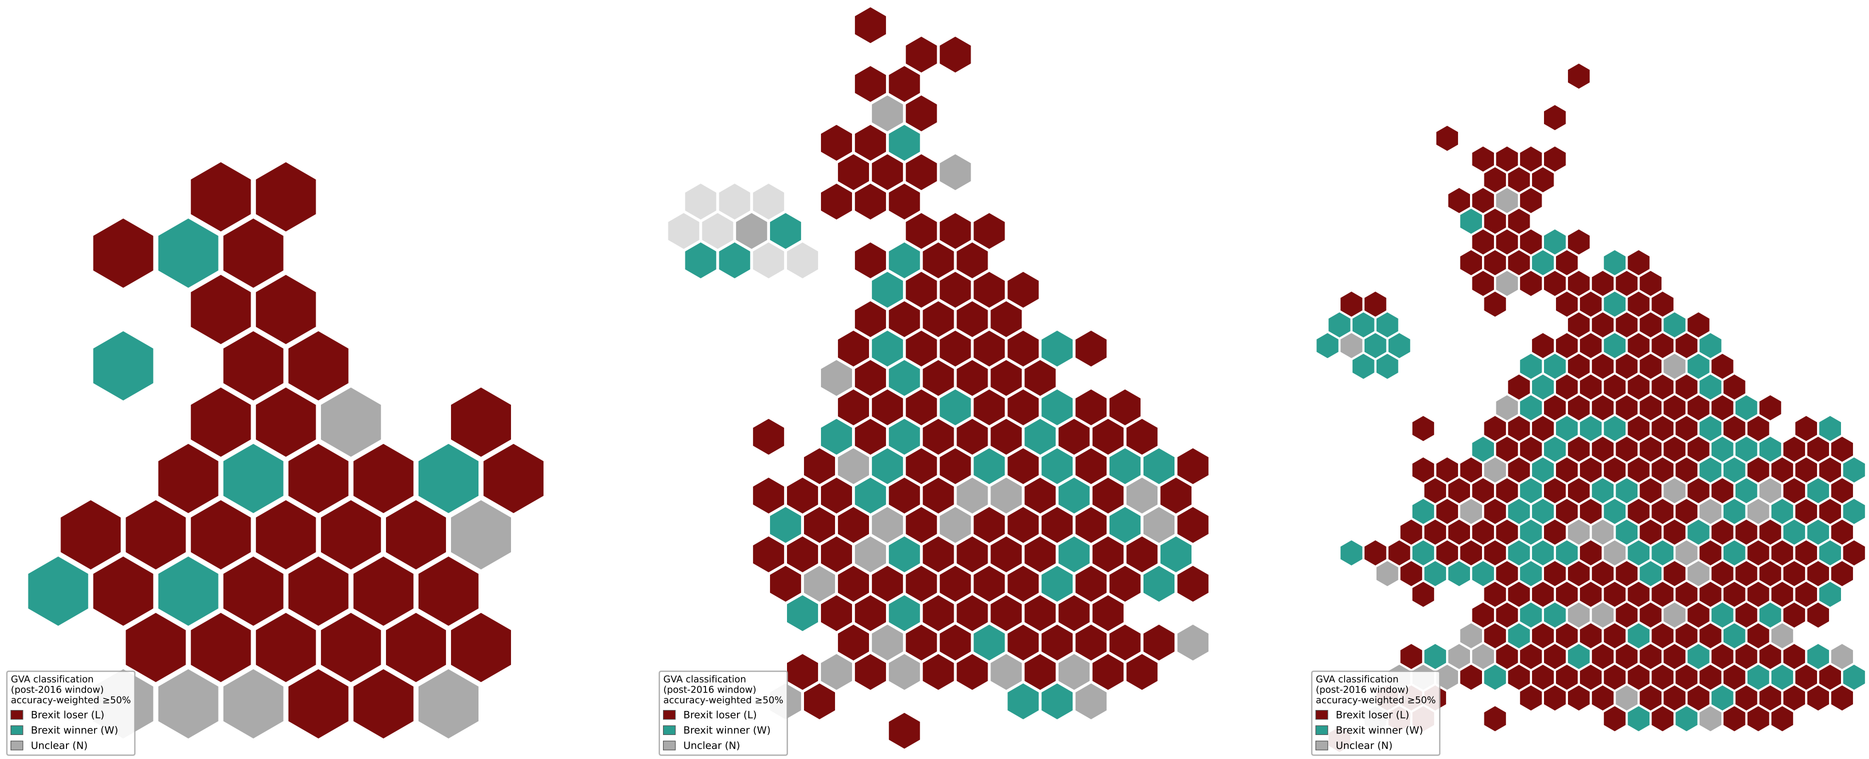

ITL2 regions ($n = 46$)

ITL3 areas ($n = 175$)

Local authority districts ($n = 361$)

Note: Each hex represents one geographic unit, coloured by its post-2016 GVA classification: dark crimson = Brexit loser (L), teal = Brexit winner (W), grey = unclear (N). A unit is classified as a loser (winner) if a simple majority (≥50%) of donor-pool specifications agree on a negative (positive) gap direction and the selected-pool ensemble confirms the same direction. The three panels show progressively finer spatial resolution: ITL2 (34L/4W/8N), ITL3 (112L/29W/34N), and LAD (216L/68W/77N).

Figure 5: 2016 Leave vote, overall trade openness, and realised Brexit GVA losses

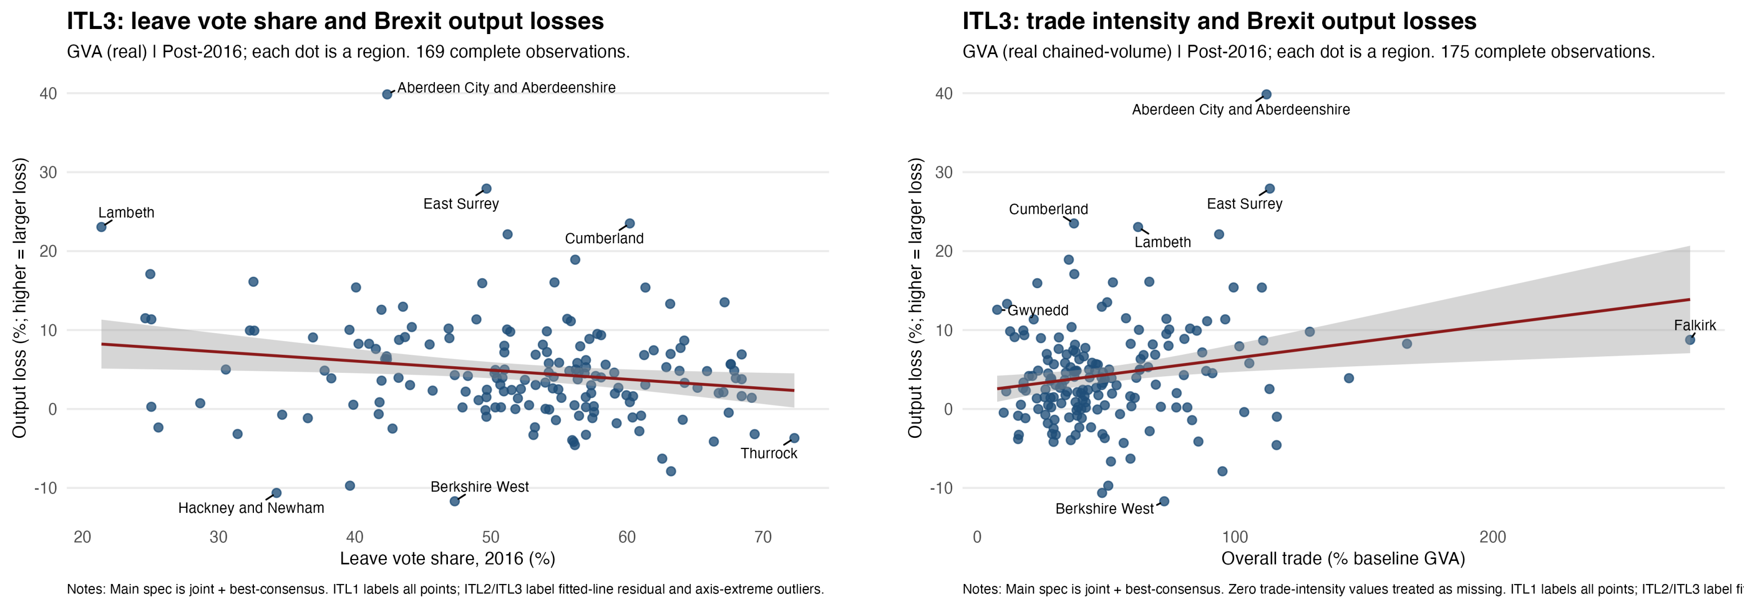

Note: *Panel A* plots the 2016 Leave vote share against the average post-2016 GVA output loss from the placebo-weighted ensemble at the ITL3 level ($n = 169$). Each point is an ITL3 sub-region; the navy regression line is from a bivariate OLS; r denotes the Pearson correlation coefficient. *Panel B* plots overall trade openness (total goods and services exports plus imports as % of regional GVA, 2019 baseline) against the average post-2016 GVA output loss at the ITL3 level. Source: referendum results from Electoral Commission; trade data from ONS and HMRC; SCM gaps from authors' estimates.

Figure 6: Brexit narrative exposure and realised GVA losses

Note: The figure classifies local authority districts into four groups according to whether their pre-referendum austerity exposure and immigration-narrative salience were above or below the national median. Each cell reports the median post-2016 GVA gap relative to the placebo-weighted synthetic-control counterfactual, the interquartile range, the number of districts, and the share classified as Brexit losers. Negative values indicate that observed GVA fell below the synthetic counterfactual. The high-austerity/high-immigration quadrant corresponds to the archetypal Leave-narrative heartland.

Tables

Table 1: Brexit GVA and income losses in 2023 for the United Kingdom, constituent countries, and ITL1 regions

	Post-2016 (loss as of 2023)										Post-2020 (loss as of 2023)									
	GVA (real)					GDHI overall (nom.)					GVA (real)					GDHI overall (nom.)				
Region	bn	%	pc	L/W	Rob.	bn	%	pc	L/W	Rob.	bn	%	pc	L/W	Rob.	bn	%	pc	L/W	Rob.
United Kingdom	-188.76	-7.7	-2835	L	0.87	-186.93	-8.5	-2807	L	0.88	-84.41	-3.6	-1257	L	1.00	-27.50	-1.3	-410	N	0.50
Scotland	-15.84	-8.6	-2908	L	1.00	-14.16	-8.9	-2599	L	1.00	-7.57	-4.3	-1386	L	1.00	-0.50	-0.3	-92	N	0.00
Wales	-3.41	-4.4	-1101	N	0.00	-1.91	-2.7	-617	N	0.60	-7.28	-8.9	-2338	L	0.98	-1.83	-2.6	-588	N	0.00
Northern Ireland	0.28	0.5	147	N	0.00	-0.66	-1.5	-349	N	0.00	-0.93	-1.8	-490	N	0.00	2.12	5.1	1115	N	0.00
England																				
North East	-5.34	-7.8	-2024	L	1.00	-2.05	-3.4	-778	L	1.00	-2.12	-3.2	-797	L	0.98	0.87	1.5	328	N	0.67
North West	1.10	0.5	160	N	0.00	-5.24	-2.8	-764	N	0.00	4.72	2.2	681	L	0.67	0.41	0.2	58	N	0.00
Yorkshire and The Humber	-8.50	-5.3	-1752	N	0.00	-7.60	-5.5	-1568	N	0.00	-5.14	-3.3	-1053	N	0.00	1.32	1.0	271	N	0.00
East Midlands	-6.77	-5.0	-1667	L	0.94	-10.57	-7.9	-2603	L	0.67	-4.04	-3.0	-984	L	0.87	0.41	0.3	100	N	0.00
West Midlands	-11.82	-6.8	-1999	N	0.00	-11.06	-7.1	-1871	N	0.00	-12.60	-7.3	-2110	L	0.98	-0.14	-0.1	-24	N	0.00
East	-8.39	-4.2	-1337	N	0.00	-16.56	-7.6	-2639	L	0.67	-9.24	-4.7	-1456	L	0.98	-3.37	-1.6	-531	L	0.93
London	-69.96	-11.6	-7927	L	1.00	-63.47	-13.1	-7191	L	0.87	-34.64	-6.1	-3916	L	1.00	-18.87	-4.3	-2134	L	1.00
South East	-45.97	-12.4	-4997	L	0.98	-36.65	-10.1	-3983	L	0.87	-19.82	-5.7	-2131	L	0.98	-3.72	-1.1	-400	L	0.87
South West	-12.11	-6.6	-2385	N	0.00	-5.31	-3.2	-1046	L	0.67	-4.47	-2.5	-869	L	0.67	3.89	2.5	757	N	0.40

Notes: All figures refer to the 2023 point estimate. 'bn' is the 2023 GVA/GDHI loss in GBP billions, computed as $\text{actual}_{2023} \times (1 - e^{-\hat{g}_{2023}})$ where $\hat{g}_{2023}$ is the 2023 log-gap from the placebo-weighted ensemble. '%' is $(e^{\hat{g}_{2023}} - 1) \times 100$. 'pc' is the 2023 per-capita loss in GBP, using post-window average population. GVA is real chained-volume; GDHI is nominal. L/W: loser (L) or winner (W) classification. Rob.: share of 31 donor-pool specifications consistent with the ensemble direction. Negative values indicate observed below the synthetic control. See Appendix Table 3 for average post-window values.

Table 2: Country-level aggregation consistency: direct vs. sub-national aggregate 2023 Brexit loss estimates

Country	Source	GVA post-2016		GDHI post-2016		GVA post-2020		GDHI post-2020	
		Real £bn	%	Nom. £bn	%	Real £bn	%	Nom. £bn	%
United Kingdom									
	Direct UK	-128.84	-7.7	-170.78	-8.5	-85.66	-3.6	-10.28	-1.3
	ITL1 Agg.	-160.48	-5.6	-113.19	-6.3	-114.73	-3.8	-1.36	-0.1
	ITL2 Agg.	-198.12	-7.4	-144.18	-6.6	-153.04	-5.6	-15.00	-0.8
	ITL3 Agg.	-210.54	-7.7	-141.01	-5.5	-177.36	-7.0	-24.38	-1.2
	LAD Agg.	-208.03	-7.8	-138.51	-6.1	-192.11	-7.2	-23.54	-1.2
England									
	ITL2 Agg.	-171.71	-6.8	-132.26	-6.2	-131.90	-4.9	-11.09	-0.6
	ITL3 Agg.	-178.76	-7.5	-130.28	-5.4	-153.22	-6.7	-25.10	-1.5
	LAD Agg.	-179.78	-7.7	-125.99	-6.1	-166.09	-6.9	-21.51	-1.3
Scotland									
	ITL2 Agg.	-21.94	-12.1	-10.35	-11.4	-14.85	-8.6	-1.96	-2.2
	ITL3 Agg.	-24.16	-10.7	-11.04	-8.6	-16.26	-8.4	-1.17	-0.9
	LAD Agg.	-23.39	-12.0	-11.01	-9.2	-17.27	-9.2	-2.05	-2.0
Wales									
	ITL2 Agg.	-4.22	-6.9	-1.22	-3.2	-7.03	-8.7	-0.60	-1.9
	ITL3 Agg.	-4.05	-8.0	-0.80	-4.2	-7.48	-9.8	-0.72	-1.8
	LAD Agg.	-2.93	-6.1	-0.83	-3.7	-6.95	-8.3	-0.65	-1.5
Northern Ireland									
	ITL2 Agg.	0.42	0.5	-0.39	-1.5	-1.80	-1.8	1.73	5.1
	ITL3 Agg.	0.37	1.4	-1.15	-3.0	-1.37	-5.0	1.41	2.6
	LAD Agg.	-1.24	-1.5	-1.17	-3.0	-3.61	-7.1	1.10	2.6

Notes: All entries refer to 2023 point estimates. % is $(e^{\hat{g}_{2023}} - 1) \times 100$ from the placebo-weighted ensemble. GBP bn is the 2023 loss, obtained by scaling the average annual loss by the ratio of the 2023 log-gap to the average post-window log-gap. GVA is real chained-volume; GDHI is nominal. Aggregate rows sum independently estimated sub-national series to the country level; agreement across rows confirms that the national estimate is not an artefact of geographic aggregation.

Table 3: Brexit GVA and income losses: average annual values post-2016 and post-2020

	Post-2016 (average over 2017–2023)										Post-2020 (average over 2021–2023)									
	GVA (real)					GDHI overall (nom.)					GVA (real)					GDHI overall (nom.)				
Region	bn	%	pc	L/W	Rob.	bn	%	pc	L/W	Rob.	bn	%	pc	L/W	Rob.	bn	%	pc	L/W	Rob.
United Kingdom	-132.61	-5.8	-1992	L	0.87	-134.47	-6.7	-2020	L	0.88	-91.20	-4.0	-1358	L	1.00	-68.56	-3.4	-1021	N	0.50
Scotland	-8.87	-5.3	-1628	L	1.00	-12.23	-8.4	-2246	L	1.00	-6.11	-3.6	-1117	L	1.00	-4.45	-3.1	-815	N	0.00
Wales	-0.12	-0.2	-39	N	0.00	-1.32	-1.9	-425	N	0.60	-5.67	-7.6	-1820	L	0.98	-1.77	-2.6	-568	N	0.00
Northern Ireland	0.53	1.0	281	N	0.00	-1.57	-3.6	-830	N	0.00	-0.55	-1.1	-291	N	0.00	1.01	2.3	529	N	0.00
England																				
North East	-3.62	-5.7	-1371	L	1.00	-2.96	-5.0	-1123	L	1.00	-2.18	-3.4	-822	L	0.98	-0.53	-0.9	-198	N	0.67
North West	-0.94	-0.4	-138	N	0.00	-5.09	-2.8	-742	N	0.00	3.83	1.7	552	L	0.67	-3.32	-1.8	-478	N	0.00
Yorkshire and The Humber	-3.73	-2.5	-769	N	0.00	-6.43	-4.9	-1326	N	0.00	-1.84	-1.2	-377	N	0.00	-1.22	-0.9	-250	N	0.00
East Midlands	-3.41	-2.7	-840	L	0.94	-8.52	-6.9	-2097	L	0.67	-2.10	-1.6	-511	L	0.87	-2.63	-2.1	-641	N	0.00
West Midlands	-3.63	-2.3	-615	N	0.00	-8.77	-6.0	-1484	N	0.00	-9.81	-6.1	-1643	L	0.98	-4.25	-2.9	-712	N	0.00
East	-1.94	-1.0	-310	N	0.00	-11.76	-5.9	-1874	L	0.67	-7.00	-3.7	-1103	L	0.98	-7.27	-3.6	-1145	L	0.93
London	-44.91	-8.4	-5089	L	1.00	-37.59	-8.9	-4259	L	0.87	-33.61	-6.3	-3799	L	1.00	-22.90	-5.4	-2588	L	1.00
South East	-21.62	-6.6	-2350	L	0.98	-24.27	-7.4	-2638	L	0.87	-11.32	-3.5	-1218	L	0.98	-10.43	-3.2	-1121	L	0.87
South West	-4.95	-2.9	-975	N	0.00	-6.42	-4.0	-1264	L	0.67	-1.70	-1.0	-332	L	0.67	-0.03	-0.0	-6	N	0.40

Notes: All figures are averages over the post-window period, computed from the placebo-weighted ensemble. 'bn' is the average annual loss in GBP billions, $\text{actual}_{2023} \times \overline{(\exp(\hat{g}_t) - 1)}$. '%' is $\overline{(\exp(\hat{g}_t) - 1) \times 100}$, averaged over post-treatment years. 'pc' is the implied average annual per-capita loss in GBP. GVA is real chained-volume; GDHI is nominal. Negative values indicate observed below the synthetic control. See Table 1 for 2023 point estimates.

Online Appendix

Measuring the Regional Economic Cost of Brexit: Evidence as of 2026

Alabrese Eleonora, Jacob Edenhofer, Thiemo Fetzer and Shizhuo Wang

June 16, 2026

A Additional Brexit-loss exhibits

The detailed regional tables for ITL2, ITL3, and local authority district levels—including money-metric GVA and GDHI-B6 losses under both treatment windows—are available in the online data supplement at https://www.brexitcost.org. Throughout, the money-metric columns report average annual output loss, not cumulative loss, and we keep the production and overall primary-income concepts separate.

B Exploring variation in Brexit-cost estimates

We examine whether cross-sectional variation in realised Brexit losses lines up with a small set of plausible mechanism variables. The forensic exercise focuses on three candidates: baseline trade openness, levelling-up funding, and COVID-19 mortality. We also compare the ex-ante local Brexit-cost predictions from Dhingra and NIESR at the ITL2 level against the realised synthetic-control gaps, to assess whether forecasted exposures align with outcomes.

Appendix Table A1 reports bivariate cross-sectional regressions for both mechanism variables at the ITL2 and ITL3 levels, with and without ITL1 parent-region fixed effects. Panel A uses overall trade openness (total exports plus imports as a share of GVA) as the explanatory variable; Panel B uses the 2016 Leave vote share. Across ITL2 and ITL3 geographies, the coefficient on baseline trade openness is positive, implying that more trade-open places subsequently experienced larger realised Brexit losses. The pattern is strongest and most stable at the ITL3 level, where the coefficient is statistically significant both with and without ITL1 fixed effects. For leave vote share, the raw correlation is negative (Remain-leaning areas experienced larger losses), but once within-ITL1 variation is isolated, the coefficient reverses sign, reflecting that even within broad regions more Leave-voting places experienced above-average losses relative to their ITL1 peers. The R^2 correctly rises from the OLS to the fixed-effects columns, reflecting the substantial share of cross-sectional variance in Brexit losses that is explained by broad regional location.

Table A1: Trade openness and leave vote share as predictors of Brexit output losses (GVA, post-2016)

	(1) ITL2	(2) ITL2 + ITL1 FE	(3) ITL3	(4) ITL3 + ITL1 FE
Panel A: Overall trade openness (exports + imports, % GVA)				
Trade openness	0.158	0.153	0.042**	0.035**
	(0.100)	(0.120)	(0.018)	(0.014)
Panel B: Leave vote share (%)				
Leave vote share	-0.167	0.209	-0.115**	0.083*
	(0.109)	(0.155)	(0.055)	(0.045)
Observations (A)	46	46	175	175
R^2 (A)	0.214	0.419	0.043	0.224
Observations (B)	45	45	169	169
R^2 (B)	0.056	0.278	0.033	0.176

Notes: Dependent variable is the average post-2016 Brexit output loss in percent (positive = larger loss; joint estimator, best-consensus selector, GVA, post-2016). Panel A uses overall trade openness (exports plus imports in % of 2019 baseline GVA). Panel B uses the 2016 Leave vote share. Columns (1) and (3) are bivariate OLS with HC1-robust standard errors. Columns (2) and (4) add parent-ITL1 fixed effects estimated with `fixest`; standard errors are clustered by ITL1. R^2 for FE columns is the total model R^2 including the variance explained by the fixed effects (always weakly greater than the corresponding OLS R^2). Significance coding: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C Appendix Tables

Table A2: Total Sets of Combinations of Donor Pools

Pool 1	Size 1	Pool 2	Size 2	Pool 3	Size 3
EU-NUTS2	175	EU-NUTS2 US-STATES	226	EU-NUTS2 US-STATES G20	241
US-STATES	51	EU-NUTS2 G20	191	EU-NUTS2 US-STATES OECD	242
G20	18	EU-NUTS2 OECD	192	EU-NUTS2 US-STATES EU	233
OECD	33	EU-NUTS2 EU	182	EU-NUTS2 G20 OECD	200
EU	27	US-STATES G20	68	EU-NUTS2 G20 EU	197
		US-STATES OECD	83	EU-NUTS2 OECD EU	195
		US-STATES EU	78	US-STATES G20 OECD	91
		G20 OECD	41	US-STATES G20 EU	92
		G20 EU	42	US-STATES OECD EU	90
		OECD EU	40	G20 OECD EU	48

Pool 4	Size 4	Pool 5	Size 5
EU-NUTS2 US-STATES G20 OECD	250	EU-NUTS2 US-STATES G20 OECD EU	253
EU-NUTS2 US-STATES G20 EU	247		
EU-NUTS2 US-STATES OECD EU	245		
EU-NUTS2 G20 OECD EU	203		
US-STATES G20 OECD EU	98		

Notes: The Table presents full set of potential combinations of donor pools drawn from the set of five potential donor sets. Cells coloured light blue include donor pools only constructed using subnational data. Cells coloured light red include only country-level donors; non-coloured cells capture a donor pool set comprised of a mix of country-level and subnational data. The counts indicated in the columns with the respective sizes represent the maximum number of spatial units included in the respective donor pool.