Applications of Data Science: Policy Evaluation and Stock Market Forecasting

Author:

Moatazbellah Farid

Supervisors:

Professor Hans-Martin Krolzig
Dr Guy Tchuente Nguembu

A thesis submitted in fulfillment of the requirements
for the degree of Doctor of Philosophy in Economics.
School of Economics, University of Kent

Date of submission: June 2023

Declaration: The work done in this thesis comprises my original work towards obtaining the PhD in Economics at the University of Kent. This thesis contains approximately 28,000 words exclusive of footnotes, tables, figures and bibliography.

Applications of Data Science: Policy Evaluation and Stock Market Forecasting

Author:

Moatazbellah Farid

Supervisors:

Professor Hans-Martin Krolzig

Dr Guy Tchuente Nguembu

University of Kent

Dedication

This thesis is dedicated to the loving memory of Mamy Mona.

Acknowledgements

The completion of this thesis would not have been possible without the guidance of my academic supervisors Professor Hans-Martin Krolzig, Dr Guy Tchuente Nguembe and Dr Mehdi Hamidi Sahneh. The support I received from Dr Mehdi Hamidi Sahneh between 2018 to 2021 was a tremendous help that I am indebted to. As well as the gratitude to my supervisors, I am grateful for all the comments and advice I received from Dr Alfred Duncan and other members of the Macroeconomic Growth History Centre at Kent. I am also thankful for all the friendly chats I had with Professor Nizar Allouch that made me feel that I am close to home.

I am grateful for the academic and the professional staff at the School of Economics, University of Kent for awarding me a scholarship to pursue this PhD and giving me all the support whenever needed. My gratitude is also extended to Dr Maria Garcia-Alonso for her continous support and understanding and for giving me the chance to teach Economics and Econometrics at different levels. I am also indebted to my PhD colleagues for all the chats and stimulating discussions in the PhD office and on Zoom.

I am extremely thankful to my parents, Shehab Farid, Dalia Hafez, my siblings Kamal Farid and Farida Farid for their continuous love, mental and financial support and encouragement, without them I would have never been standing here. A special thanks goes to my wife Ayda Farid, who never failed to push me, for all the support, encouragement, love and more importantly dealing with me when my codes did not run. To my newly born daughter, Daylia, I hope you are proud.

It ain't what you don't know that gets you into trouble.

It's what you know for sure that just ain't so.

- Mark Twain

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Over the last decade the application of data science in economics and finance has increased considerably. Data science including causal inference methodologies and machine learning techniques for economic analysis, forecasting and nowcasting economic and financial variables has shown great success. In this thesis, I will be exploiting data sciences techniques to evaluate the impact of policy changes and perform forecasting for the stock market. This thesis consists of three other chapters. Chapter two looks at evaluating the economic impacts of the Brexit vote on the UK manufacturing sector employment. Chapter three evaluates the success of the Stamp Duty Land Tax (SDLT) holiday on the UK housing market. In these two chapters, I use the "State of Art" synthetic control method to perform this evaluation. The choice to use the synthetic control method to investigate the impacts of the Brexit vote is motivated from the approach that the method builds on an improvement from traditional methods of policy evaluation such as the Difference in Difference (DiD) method. In contrast to the DiD which is designed to be applied to pooled cross sectional data, the design of the synthetic control method makes it more applicable for comparative case studies where there is only one treated unit, such as in the examples of Brexit referendum and the SDLT holiday. Instead of using one single control unit or simple average of control units as in the DiD, the synthetic control method uses a data driven approach to match the characteristics of the treated unit to a selected group of countries in the donor pool. This data driven approach generate weighted average of a set of potential control units to provide a synthetic control unit that closely resembles the treated unit in the pre-treatment period.[1] Lastly, chapter four applies unsupervised machine learning techniques along with dictionary based method to forecast the volatility of the stock market.

Chapter two of the thesis is titled "***The Winners and Losers of Brexit: UK Manufacturing Sector Analysis***" and looks at the impact the Brexit vote had on the UK economy. More specifically, I look at the UK manufacturing sector to investigate how the Brexit vote affected the labour market in this sector. In answering this question, I use the Synthetic Control method of Abadie et al. (2010). This method allows the performing of causal inference analysis when there is as little as one treated unit and a pool of control units, by using a data driven approach to assign weights to units in the control group so they mimic the behaviour of the treated unit as closely as possible. The findings of this paper were that in the three years following the Brexit referendum, the UK's total manufacturing sector's employment decreased by an average of 3.6% in the post Brexit vote period, equivalent to loss of 91,522 jobs. As the degree of reliance on the European Union (EU) varies across manufacturing divisions, I narrow down my analysis to study the impact that the referendum had on seventeen manufacturing divisions. I found that three manufacturing divisions benefited from the Brexit vote, six manufacturing divisions were negatively affected, while seven manufacturing divisions were unaffected. Earlier version of this paper is available on SSRN as Farid (2021b) and was presented at the 14th FIW Research Conference "International Economics" at Vienna University of Economics and Business (Wirtschaftsuniversität Wien).

Chapter three of the thesis entitled "***Land Tax Holiday and House Price: Evidence from the UK***", investigates the impact that the Stamp Duty Land Tax (SDLT) holiday that was introduced in July 2020 had on the UK housing market. The SDLT was introduced following the ease of the COVID-19 lockdown restrictions in an attempt to boost the UK housing market. I perform a counterfactual analysis using Synthetic Control Method to answer the question "How would the UK housing market perform

[1] Athey and Imbens (2017) describe the synthetic control method as "arguably the most important innovation in the policy evaluation literature in the last 15 years"

in the absence of the tax holiday?" The findings of this paper illustrate that the SDLT holiday resulted in an increase of 2.11% and 2.96% on annual rates in the house prices in the two quarters following the tax holiday. As Ashworth and Parker (1997) point out it would be misleading to assume that the housing markets in all the UK regions operate in the same way. In this paper, I document the heterogeneity of the housing market performance across the 9 Nomenclature of Territorial Units for Statistics 1 regions (henceforth NUTS 1). My analysis was motivated by the findings of Ramani and Bloom (2021) that COVID-19 and the work from home mandate boosted the demand for houses outside central cities. My findings show that Greater London did not benefit from the SDLT holiday in the same way as the other regions as the average house price falls outside the taxable exempt amount. On the other hand, the regions Wales, North West of England and North East of England exhibited the greatest benefits, explained by the difference of the UK and its synthetic counterpart. Earlier version of this paper is available on SSRN as Farid (2021a) and was presented at the ESCoE Conference on Economics Measurement at University of Strathclyde.

Finally, chapter four of the thesis entitled "***Predicting the Volatility of the S&P 500 Stock Index: The Role of News Sentiment***" I deviate from the causal inference analysis performed in the previous two chapters. In this chapter I use unsupervised machine learning models to decompose the US newspaper articles into different economic topics. The unsupervised machine learning approach is called Latent Dirichlet Allocation (LDA) of Blei et al. (2003). The LDA method illustrates that the US newspaper corpus can be decomposed into 17 topics. I then apply dictionary based method to measure the tonality of the news topics over time and perform sentiment analysis using the state of the art Harvard General Inquirer (GI) dictionary. The analysis to investigate which news topics can explain the stock returns and volatilities was motivated by Peng and Xiong (2006) and Da et al. (2011). They show that investors are selective with regard to the information they process as the wealth of information creates a poverty of investor attention as investors may have restrictions on following all the information that can be relevant to their decision making. Although the topical news sentiment indices are found to be noisy, they tend to have extreme observations during the time of economic recessions. I then use the news sentiment indices to predict the volatility of stock returns with the Harvard GI dictionary, my findings show that incorporating topical news sentiment as exogenous regressors in GARCH(1,1)-X and GJR-GARCH(1,1)-X models explains the stock returns volatility. More precisely, the news sentiment indices for topics related to currency, monetary policy, manufacturing, trade and trade policy can explain the variation in the volatility of stock returns. I then perform a forecasting exercise to find that when combining the individual forecasts using regression based methods, the combined forecasts provide accuracy in predicting the volatility of stock returns.

Combining the three chapters mentioned above together in a PhD thesis provides a comprehensive approach to analysing economic data for the use of economic analysis and forecasting. The application of data science and big data as demonstrated in this thesis enables economists and researchers to analyse causal relationships between variables, which provides an understanding of the underlying mechanisms driving economic phenomena. This allows for more accurate policy analysis, economic forecasting, which can ultimately lead to more efficient and effective economic outcomes.

Chapter 2

The Winners and Losers of Brexit: UK Manufacturing Sector Analysis

Abstract

This paper analyses the effect of the Brexit vote on UK manufacturing divisions using a synthetic control methodology. More specifically, this paper answers the question "What if the Brexit vote never happened?" with the focus on the employment of the manufacturing sector. The findings show that in the three years following the Brexit referendum, the UK total manufacturing sector's employment decreased by an average of 3.6%. When investigating the impact on individual manufacturing divisions, I found that three manufacturing divisions benefited from the Brexit vote, six manufacturing divisions were negatively affected, while seven manufacturing divisions were unaffected. These results are of particular importance to policymakers who will support sectors affected as a result of the Brexit referendum.

2.1 Introduction

On 23rd of June 2016, 52% of the British voters have voted to leave the European Union (EU). Having been part of the EU, the UK had access to the single market and the customs union which in turn made all the trade between the UK and the EU tariff free, subject to low non-tariff barriers and allowed for free movement of EU labour across the

EU member states. As a result of the Brexit vote results, the UK government triggered article 50 on 29th of March 2017, which is the plan that a member state has to follow in the event of leaving the EU. There are several ways that the UK manufacturing sector might be affected by Brexit. Firstly, several UK sectors depend on the EU migrant labour, restrictions regarding the free movement of labourers from the EU may lead to further decline in the number of skilled labours and increase in wages and thus the increase in the costs of production (European-Commission 2016). Secondly, the depreciation of the sterling pound caused an increase in the input costs due to raising imported raw materials (Wells and Stephens 2017). Thirdly, UK-EU Trade deal negotiations caused an increase in business and economic uncertainty. The uncertainty associated with the Brexit vote affected the UK manufacturing sector through the option-value channel of (Bloom 2009; Bloom et al. 2019).[1]

Manufacturing accounts for around one tenth of UK economic output, 8% of employment and makes up 80% of UK goods exports and 44% of all UK exports (Tetlow and Stojanovic 2018). In this paper, I assess the impact of the Brexit vote on the UK manufacturing sector employment by constructing an appropriate counterfactual for each manufacturing division to investigate how the employment of manufacturing divisions would have been performing in the absence of the Brexit vote. As displayed in figure 2.1, the manufacturing sector confidence had a dip at the time around the Brexit referendum and another big decline at the end of 2019, when there was concerns that the UK will leave the EU without a deal. This paper is of particular importance for policymakers who aim to support British industries and businesses that are affected by Brexit. The British government is to provide training and financial support to businesses and industries affected by Brexit.[2] The need for financial support for each manufacturing division depends on its degree of internalisation. Divisions that rely on European skilled labour as well as trade with the EU are likely to be affected the most. My results show policymakers offering support, the divisions that were affected the most in terms of employment in the three year transition period, as the number of European skilled labour registered to work in the

[1] According to (Bloom 2009), during uncertain times, hiring and investment rates fall due to the option-value effect, as businesses delay their plans until uncertainty subside.

[2] At the point of writing this paper, the British government announced the launch of operation Kingfisher, Brexit Readiness Fund and the availability of 1.3 billion pounds to lenders to support businesses affected from Brexit

UK fell sharply following the Brexit referendum. The support that the government can offer can be through campaigns that aim to promote employment opportunities, attract domestic workers, and develop the skills required in the manufacturing industry such as The National Manufacturing Competitiveness Levels program. This program offers tailored support and funding to help businesses identify and address skills gaps, improve processes and the sector-specific apprenticeship programs that provide structured training and on-the-job learning opportunities for individuals interested in pursuing a career in the manufacturing sector.

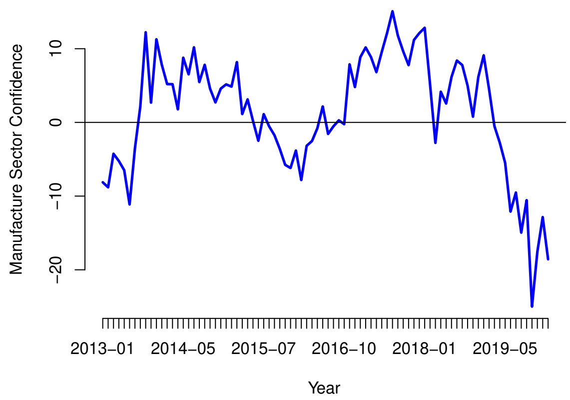

Figure 2.1: Manufacturing Sector Confidence.

OECD Manufacturing - Confidence indicator, Quarterly Seasonally Adjusted. Source: Quandl.

When studying the impact of the Brexit vote on the UK manufacturing sector, I use synthetic control method (Abadie and Gardeazabal 2003; Abadie et al. 2010; Abadie et al. 2015) on the aggregate manufacturing sector employment and sixteen 2-digit manufacturing divisions to construct a counterfactual for what the labour market of each manufacturing division would have been in the absence of the leave vote. As the Brexit vote was unanticipated and unrelated to macroeconomic performance, the synthetic UK should carry to evolve in the same direction the UK would have done in the absence of the referendum. Comparing the evolution of the UK to its synthetic counterpart shows the causal impact of the Brexit referendum in the post-Brexit period. The use of synthetic control method in the context of comparative case studies have grown considerably in the last decade since the seminal paper of Abadie and Gardeazabal (2003). This method has several advantages such as the ability to analyse policy change when there is as few as one treated unit, the method constructs a counterfactual scenario by estimating what would have happened in the absence of the treatment, the method takes into account time-varying confounding factors that may affect both the treated and control units. By matching the pre-treatment covariates, the method creates a synthetic control unit that closely resembles the treated unit's pre-treatment characteristics, lastly, the method allows for the performance of inference through the application of placebo tests; this enables the robust evaluation for the treatment effect.

This paper contributes to the research on the impact of the Brexit vote on the UK economy using the synthetic control methodology. Born et al. (2019) found that the Brexit vote caused the UK real GDP to decrease by 2% by 2018. Breinlich et al. (2020) found that the Brexit vote caused the UK outward FDI to increase by 17% and inward FDI to decrease by 9% to and from the EU, as firms move their operations from the UK to the EU to benefit from the access to the single market. Opatrny (2020) investigates how the UK financial markets would have performed if the UK was to remain in the EU; his results found that as a result of the leave vote, the 10 year UK bond yield increased by 1.2% and real effective exchange rate decreased by 10.8 index points. Fetzer and Wang (2020) study the impact of the Brexit vote on a regional level, they show that the Brexit vote caused 168 local district to lose on average 8.54 percentage points of their output. They show that regions that highly rely on manufacturing are affected the most. Papyrakis et al. (2023) show that the Brexit vote had a little impact on the UK employment, causing an increase in the unemployment rate. In contrast to these papers, I use dis-aggregated data of the UK manufacturing divisions to investigate the economic impact of the Brexit vote. Such that I narrow my analysis by looking at the manufacturing sector only and the impact Brexit had on the labour of this sector this is due to the fact that this sector relies heavily on international trade and skilled EU labour.

This paper also contributes to the literature on forecasting the impact that Brexit can have on the economy and on the manufacturing sector in particular. Gasiorek et al. (2019) investigate the impact of five different scenarios of trade relationships between the UK and the EU, ranging from soft Brexit to a hard Brexit scenarios, on the output, international trade and prices of 122 manufacturing industries using partial equilibrium model.[3] Their results found that all the scenarios yield a negative impact on the UK economy, with hard Brexit having the greatest negative impact. They also found that textiles, clothing and footwear industries will see the largest negative impact, but food processing may see an increase in production. Dhingra et al. (2017) perform counterfactual analysis on the impact of the Brexit vote on welfare using general equilibrium trade model, covering 31 sectors and 35 regions. Their results found that a hard Brexit will reduce the welfare by 2.7%. Lawless and Morgenroth (2019) use data for 6-digit products from United Nations ComTrade for 2015 to assess the impact of the UK reverting back to WTO tariffs with the EU countries. As different sectors have different tariffs, their results show that the exports from EU to UK will be reduced by 30%, while exports from UK to EU will be reduced by 22%. The approach that is taken in this paper is different from the previously highlighted literature, my paper is mainly motivated by the degree of heterogeneity of UK manufacturing sector divisions in respect to trade and reliance on the EU labour. By focusing on the main manufacturing divisions and investigate how the Brexit referendum affected their employment, this paper sheds light on the nuanced effects of the Brexit referendum up to 2019. The counterfactual analysis performed in this paper looks at answering the question "What if the Brexit vote never happened" by looking at the manufacturing sector's employment which is an aspect that has not been investigated before. This allows the estimating of the average impact of the Brexit vote across different manufacturing divisions. More precisely, the focus of this paper is on how the manufacturing sector's employment would have evolved in the absence of the Brexit vote. This question represents a vital argument in the Brexit effect debate as by looking at the changes in the employment in the manufacturing sector, one can understand the real-world visible consequences of the Brexit vote. This sheds the light on the businesses responses to the referendum outcome, whether businesses reduced their workforce, shifted their business operations outside of the UK in response to the Brexit uncertainty or did no actions. This also answers the question of how certain divisions are more affected

[3] Hard Brexit is defined as reverting back to World Trade Organization terms (WTO).

than others? By addressing these questions in the context of counterfactual analysis, one can provide valuable insights to academics and researchers seeking to understand the economic effects of Brexit. This can also support policymakers to develop strategies aimed to mitigate any negative impacts on the manufacturing sector employment.

The results of this paper demonstrate that within the sixteen divisions analysed, the manufacture of wearing apparel and manufacture of motor vehicles were negatively affected the most due to their reliance on the EU. The manufacture of wearing apparel is heavily reliant on EU trade and EU migrant skilled labour and imposition of trade tariffs is likely to hit this division hard, the results show that it had an average decline of 28% of its employment. Similarly, the manufacture of motor vehicles division, that exports over 50% of its output to the EU, showed great negative responses during the transition period as big manufacturers such as Jaguar Land Rover and Nissan decided to cut their labour force or shift their operations outside the UK, my results find that employment declined by an average of 9.3%. Gasiorek et al. (2019) postulate that Stratford-on-Avon, and Leicester are at risk of the greatest loss of their employment due to their heavily reliance on manufacturing of motor vehicles and manufacturing of wearing apparel respectively. Unsurprisingly, there are some divisions that benefited during the transition period, these divisions are manufacturing of non-metallic mineral products and manufacture of chemical products. The former had an average increase in employment by 9.1% and the latter had an average gain of 5.8%. The gain that these manufacture divisions experienced can be attributed to the some factors such as the increase in the domestic demand for their products which could offset the negative effect of the Brexit referendum, increase in exports as a result of the sterling depreciation or less reliance on European skilled labour.

This paper is organised as follows. In section 2.2, I describe the data used and the synthetic control method. In section 2.3, I present my empirical results of the impact of the Brexit on manufacturing divisions employment. In section 2.4, I perform robustness check. In section 2.5, I discuss the main findings, implications and provide recommendations. In section 2.6, I conclude.

2.2 Data and Model

2.2.1 Data

To investigate the impact of the Brexit vote on UK employment, I use the number of persons employed index, for the UK and 16 EU countries from Eurostat database Short Term Business Activities from 2013 Q1 to 2019 Q4.[4] The choice to work with this data was mainly based on the advantage that Eurostat database is the only database that provides long enough data on the manufacturing sector employment and its individual divisions across different countries.[5]

Table 2.1 below shows descriptive statistics and a breakdown of each manufacturing division that is used in my analysis.[6] The table shows that the manufacture of food products division has the greatest share of the total manufacturing sector in terms of output and employment. In terms of international trade, manufacture of other transport equipment and manufacture of machinery and equipment are the greatest net exporters divisions, while manufacture of food products and manufacture of paper and paper products are the greatest net importers.

Figure 2.2 below shows that following the Brexit referendum the sterling pound dropped to a 31 year low, causing the input prices to increase for almost all the manufacturing divisions. As displayed, the manufacture of leather and manufacture of wearing apparel had the highest PPI input price by the end of the sample. This depreciation in the sterling has a negative impact on the manufacturing sector production as it puts pressure on the cost of imported products. Breinlich et al. (2019) found that the depreciation of the sterling following the referendum led to higher inflation by increasing the cost of imports, causing product groups with high import shares to experience significantly higher inflation.

[4]Eurostat defines the total number of persons employed as total number of persons who work in the observation unit as well as persons who work outside the unit but belong to it and paid by it.

[5]The United Nations Industrial Development Organization (UNIDO) provides annual data on employment. Working with annual data is not appropriate in the context of synthetic control method due to the limited data available in the post treatment period.

[6]All data used to calculate manufacturing divisions share of total manufacturing sector, and international trade imports and exports is available upon request.

Division (1)	Output (2)	Establishment (3)	Employee (4)	Net Exports (5)
Food products	16.77%	5.49%	13.59%	-5058.31
Textiles	1.77%	3.49%	3.29%	335.29
Wearing apparel	1.08%	3.86%	2.51%	-188.31
Leather	0.30%	-	-	-
Wood	1.40%	5.38%	2.27%	-622.71
Paper	2.58%	1.61%	2.56%	-705.56
Printing	5.56%	15.22%	7.45%	173.36
Chemicals	10.37%	2.42%	6.26%	3865.70
Rubber and plastic	4.33%	4.29%	6.07%	234.37
Non metalic	2.83%	3.05%	3.67%	95.47
Basic metals	3.81%	1.51%	3.04%	2135.97
Fabricated materials	6.01%	17.51%	10.23%	1350.07
Machinery	8.38%	10.92%	10.23%	7806.70
Electrical equipment	3.54%	3.59%	4.76%	473.92
Motor vehicles	9.28%	2.04%	6.36%	3420.87
Other transport	5.03%	1.65%	4.73%	11161
Furniture	3.01%	11.55%	5.39%	-345.69

Table 2.1: UK manufacturing sector breakdown by division

Columns 2 to 4 for each manufacturing division are the average between 1990 and 2018 of the share of each variable of the manufacturing division of the total manufacturing sector, variables are obtained from UNIDO IND STAT, 2020 Revision 3 with own share calculations. Column 5 I deflate the ONS trade in goods by industry, country and commodity for exports and imports using ONS CPI Index and calculate the net exports as the difference between real exports and imports of the UK with the worldwide countries.

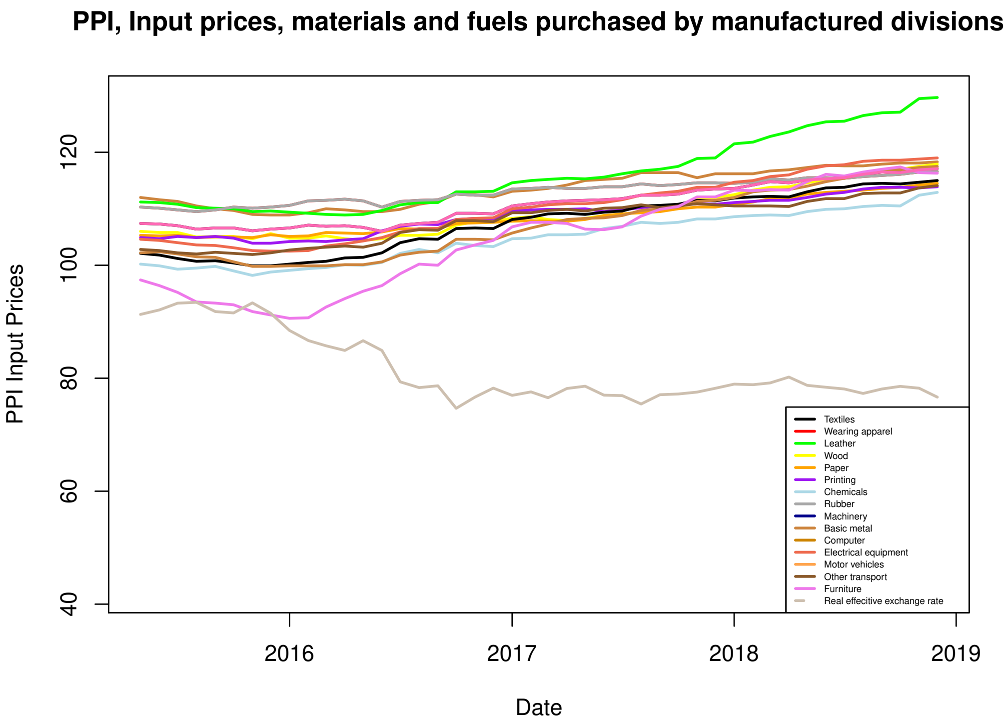

Figure 2.2: Producer Price Index (PPI)

PPI for 15 manufacturing divisions on a monthly basis between May 2015 to December 2018. Source: Office of National Statistics.

2.2.2 Model: Synthetic control method

The choice to use the synthetic control method to investigate the impacts of the Brexit vote is motivated from the approach that the method builds on an improvement from the traditional panel data models such as Difference in Difference (DiD) or Dynamic Panel Data method. This method was once referred to as "the most important innovation in the policy evaluation literature in the last 15 years" (Athey and Imbens 2017). The advantages of the synthetic control method stem from the ability to model the treatment effect when there is as little as one treated unit, something that can sometimes be unfeasible in other panel data methods. In addition to this, synthetic control method is superior with respect that instead of using one single control unit or simple average of control units as in the DiD, the method uses a weighted average of a set of potential control units to provide a synthetic control unit that closely resembles the affected unit in terms of predictors. That is to say, instead of using just one European country or a simple average of all the European countries in the donor pool, the weighted average control group has better match than any single country would be. This allows for the estimation of the counterfactual outcome that would have occurred for the treated unit in the absence of the treatment. By comparing the treated unit to its counterfactual, policymakers can understand the effectiveness of the treatment on the treated unit. Lastly, the method opens the door to performing several placebo tests in order to assess the validity of the estimated treatment effect.

In the context of the UK Brexit scenario, suppose that there are $J+1$ units (here: country pairs) in periods $t = 1, 2 \ldots, T$ (here: quarters), where unit $j = 1$, is the treated unit (here:UK) and units $j = 2, \ldots J+1$ are the untreated units used in the control group. Unit $j = 1$ is exposed to the intervention (here: outcome of the Brexit referendum) of at periods $T_{0+1}, \ldots, T$ (2016 Q3 to 2019 Q4), which is 15 observations, while being unaffected during the periods $1, \ldots, T_0$ (2013 Q1 to 2016 Q1). The data ends at 2019 Q4 so the COVID-19 pandemic shock does not influence the results.[7] Y_{it}^N is the outcome of unit i in the absence of intervention for units $i = 1, \ldots, J+1$. Y_{it}^I is the outcome of unit i at time t if the unit is exposed to intervention in periods $T_{0+1}, \ldots, T$. The intervention

[7] Stopping the exercise at 2019 Q4 prevented the analysis from being expanded to look at the long term effects of Brexit especially after the Brexit deal came in place.

has no impact on the pre-intervention periods. $Y_{it}^N = Y_{it}^I$ for $t = 1, \ldots, T_0$. However, in practice, interventions may have an impact prior to their implementation (e.g., via anticipation effects) (Abadie et al. 2010). For the treated unit, Y_{it}^N can not be observed in the post-intervention period $t = T_0 + 1, \ldots T$ and can best be modelled by a weighted average combination of untreated units in the donor pool. The synthetic control can be represented by a $(J\text{x}1)$ vector of weights $W = (w_2, \ldots, w_{j+1})'$, where $0 \leq w_j \leq 1$ for $j = 2 \ldots, J+1$ and $w_2 + \ldots + w_{j+1} = 1$, where

$$\hat{Y}_{1t}^N = \sum_{j=2}^{J+1} w_j^* Y_{jt}. \tag{2.2.1}$$

The constructed synthetic control unit $\hat{Y}_{1t}^N$ is a weighted average of the untreated units of the control group. An important assumption of the synthetic control method is that manufacturing employment of countries in the control group is not directly affected by the results of the Brexit referendum, nor any contemporaneous events.[8] The assumption of conditional independence helps ensure that the control units selected for the synthetic control group have similar pre-treatment trends, making them plausible substitutes for the treated unit in estimating the counterfactual outcome. That is, the assumption suggests that the treatment is unrelated to the historical outcomes and covariates, except through their common dependence on observed pre-treatment characteristics. By matching the pre-treatment characteristics of the treated unit with appropriate control units, the synthetic control method aims to create a weighted average of the control units that closely approximates the treated unit's outcomes in the absence of the treatment. When this assumption holds, it implies that the treated unit and the control units used to construct the synthetic control group should have similar trends in past outcomes and covariates. This similarity is crucial because it allows the synthetic control unit to serve as a meaningful representation of the unobserved counterfactual scenario for the treated unit (Gilchrist et al. 2023).

The choice of W such that the characteristics of the treated unit are best resembled by the characteristics of the synthetic control. X_1 is a $(k \times 1)$ vector of pre-intervention

[8] I am aware of the weakness of this assumption, however, this assumption is maintained in (Born et al. 2019) (Breinlich et al. 2020) and (Papyrakis et al. 2023) that countries in the donor pool are not affected by Brexit vote. In my robustness check, I show that the UK had the greatest impact of the Brexit in comparison to the countries in the control group.

characteristics of the treated unit. X_0 is $(k \times j)$ matrix of the same characteristics of the donor pool. There is no one way that is always accepted in the literature for the choice of covariates. Abadie et al. (2010) advises on using economic predictors in addition to lags of the treated variables in the pre-intervention period, while Abadie et al. (2015) advises on using economic predictors in addition to the average of the treated unit in the pre-intervention period. However, Ferman et al. (2020) finds that specification searching is reduced when the number of pre-intervention period goes to infinity and all pre-intrevention lags of the variable of interest of the treated unit are used as predictors.The vector $\mathbf{W}^*$ is chosen to minimise this difference between the pre-intervention characteristics of the treated unit and the synthetic control unit $||X_1 - X_0 W||$ subject to the constraints of w_j mentioned above. Abadie et al. (2010) measure the discrepancy between the characteristics of the treated and the synthetic control unit by introducing vector $\mathbf{V}$, which is a symmetric and positive semidefinite $(k \times k)$ matrix. An optimal choice of $\mathbf{V}$ assigns weights to the linear combinations of the variables in $\mathbf{X}_0$ and $\mathbf{X}_1$ to minimize the mean square error of the synthetic control estimator. Typically, $\mathbf{V}$ is selected to weight the predictors in accordance to their predictive power on the outcome, this makes the minimization problem as follows:

$$\|\mathbf{X}_1 - \mathbf{X}_0\|_{\mathbf{v}} = \sqrt{(\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})'\mathbf{V}(\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})}.$$

As suggested by Abadie and Gardeazabal (2003) and Abadie et al. (2010), I set the V to be determined using a data driven approach such that the MSPE of the outcome variable is minimised for the pre-treatment period. Abadie et al. (2010) show that if the weighted average of the synthetic control characteristics X_0 can match the treated unit characteristics X_1, it provides a valid counterfactual for Y_{jt}^N in the sense that $\hat{Y}_{1t}^N - Y_{jt}^N$ is close to 0 in the pre-intervention period, $t = 1, ..., T_0$. In the case of evaluating the impact of the Brexit vote on the UK manufacture sector employment, I use the variables Gross Fixed Capital Formation (GFCF) share of the GDP, manufacturing sector employment share index (2015=100), Manufacturing Value Added (MVA) and the dependent variable (Log of employment index). Including the pre-Brexit values of the dependent variable helps to limit the sensitivity of the weighting procedure to any unobserved characteristics which could result in differences between the actual UK and its synthetic counterpart and ensures a better fitting pre-intervention model. These observed covariates are not affected by the treatment in the pre-treatment period. The synthetic control is primarily designed to use any covariates that help explain the outcome variable as predictors, and not only pre-treatment values of the outcome variable (Gardeazabal and Vega-Bayo 2017). The choice of these data is motivated by the fact that these variables are available for the manufacturing sector in general and each manufacturing division in particular, hence these covariates are division specific economic covariates when it comes to estimating the synthetic control method for individual manufacturing divisions. I evaluate their economic predictive power by running a regression of the dependent variable, Log manufacturing sector employment on several covariates (Shimeng Liu 2015; Breinlich et al. 2020).

$$Y_{it} = \beta_0 + \beta_1 GFCFShare_{it} + \beta_2 EstablishmentsShare_{it} + \beta_3 EmploymentShare_{it} + \\ \beta_4 MVA_{it} + u_{it}, \;\; i = 1, \ldots, J, \; t = 1, \ldots, T.$$

As suggested by Botosaru and Ferman (2019), the inclusion of covariates as predictors when the outcome equation is nonlinear in covariates should not necessarily be a problem if one uses the data-driven approach to determine the matrix V. As long as the pre-treatment outcome lags are also included as predictors, then the importance given to such covariates would be relatively small and the estimation procedure would not attempt to match on covariates that should not be matched on. As the output of the regression shows in table 2.2, there exists linear and significant relationship between all of the variables, with exception of GFCFShare. Therefore, the inclusion of variables as economic predictors will be based of being significant determinants for the employment in the manufacturing sector. The variables with significant coefficients along with average of the pre-treatment lags of the dependent variables will be used as covariates in the data driven approach that estimates the vector of weights assigned to each country in the donor pool as in Abadie et al. (2015). As a robustness check, I also consider alternative specifications such as using the pre-treatment period lags only as economic covariates (Ferman et al. 2020; Breinlich et al. 2020; Botosaru and Ferman 2019) and using economic covariates along with all the lags of the dependent variable in the pre-treatment period (Billmeier and Nannicini 2013; Bilgel and Galle 2015; Stearns 2015; Villar and Papyrakis 2017; Pellegrini et al. 2021; Papyrakis et al. 2023).

For the post-intervention period, where $t \geq T_0$, the treatment effect is the difference between the realized outcome and the synthetic control outcome.

	Dependent variable:	
	Log(EmploymentIndex)	
	(1)	(2)
EstablishmentsShare	0.005^{***}	0.005^{***}
	(0.001)	(0.001)
GFCFShare	-0.0002	-0.0005
	(0.0002)	(0.0004)
EmploymentShare	0.003^{***}	0.003^{***}
	(0.001)	(0.001)
MVA	-0.0004	0.013^{*}
	(0.0004)	(0.007)
Constant	3.873^{***}	
	(0.097)	
Observations	169	169
R^2	0.312	0.364
Adjusted R^2	0.296	0.237
F Statistic	18.631^{***} (df = 4; 164)	20.013^{***} (df = 4; 140)
FE:	(NO)	(YES)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 2.2: Regression estimates for the Manufacturing sector employment and covariates

Estimation of the relationship between the economic covariates and the dependent variables. The dependent variable is the natural logarithm of the Employment Index 2015=100. Column (1) is the estimation with Pooled OLS. Column (2) is the estimation with country and time fixed effects.

$$Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt} = Y_{1t} - \hat{Y}_{1t}^N \ , \forall t = T_{0+1}, \ldots, T \tag{2.2.2}$$

Throughout this paper, the main identifying assumption is that the UK total manufacturing sector and its individual employment division indices would have developed as the synthetic unit had it not been for the Brexit vote. To quantify the short run effects of the Brexit vote, I estimate the following regression equation:

$$Y_{1t} - \hat{Y}^{N}_{1t} = \mu + \lambda d_t + \epsilon_t, \tag{2.2.3}$$

$$d_t = \begin{cases} 0, & \text{for } t < T_0 \\ 1, & \text{for } t \geq T_0 \end{cases}$$

The step dummy variable, d_t, models the average impact in the post intervention period, where the coefficient λ will be referred to as the "synthetic control regression estimator". When subtracting the synthetic control unit from the treated unit, this estimates the synthetic control effect.

As a robustness check, I will compare the estimates of equation 2.2.3 to a DiD regression model. The synthetic control method is more appropriate method to be used in comparative case study and in the context of Brexit referendum evaluation in particular than the DiD as there is one treated unit, where the synthetic control relies on the assumption that the synthetic control captures the unobserved factors affecting the treated unit, enabling causal inference. On the other hand, DiD is commonly employed in pooled cross-sectional analyses involving multiple treated and control units. The difference between the synthetic control method and the DiD method is that the synthetic control method creates a weighted average of the countries in the donor pool that matches the characteristics and the trend of the treated unit. The synthetic control method shows a plot of how the treated unit would have been looking in the absent of the treatment. The synthetic control method allows an algorithm to determine which combination of countries match the evolution of UK manufacturing sector employment before the Brexit vote with highest possible accuracy. On the other hand, DiD is based on the conditional expectation assumption i.e. the expected outcome of the treatment group (or the treated units) given a specific set of covariates or observed characteristics. It represents the average treatment effect for the treated group, taking into account the control group's outcomes and the changes over time. Although the synthetic control method compares the change in outcomes over time between a treated unit and a control unit or group of units that are assumed to be similar in all relevant respects except for the treatment or policy intervention. The DiD does not take in consideration matching the treated group and their characteristics, instead it is based on the common trend assumption. That is, the treatment and control group follow a similar pattern in the pre-treatment period, they would also follow a similar pattern in the post-treatment period in the absence of the treatment. The DiD model can be written as:

$$Y_{it} = \beta_0 + \beta_1 treated + \beta_2 treatment + \delta(treated \times treatment) + Z_{it}{}'\theta + \epsilon_{it}.$$

For the scenario of the Brexit, the DiD model will be written as:

$$Y_{it} = \beta_0 + \beta_1 UK + \beta_2 d_t + \delta(UK \times d_t) + Z_{it}{}'\theta + \epsilon_{it}, \tag{2.2.4}$$

Where Z_{it} is the set of economics predictors which may be included in the regression. The coefficient β_1 illustrates the effect of the UK on the dependent variable, the coefficient β_2 shows the effect of the post-2016Q2 on the dependent variable. Lastly, the coefficient δ is the DiD estimator which represents the average treatment effect (ATE) of Brexit on the treated unit. This specification estimated with and without Fixed Effects (FE).
The difference between equations 2.2.3 and 2.2.4 can be explained as the former looks at estimating the effect of the Brexit referendum on the gap between the treated unit and its synthetic counterpart, where the synthetic counterpart is assumed to be the same as the treated UK in the absence of the Brexit referendum. The coefficient λ captures the change in this gap over time, indicating the evolving effect of Brexit on the treated unit relative to the synthetic control unit. However, in the latter, the coefficient δ is the DiD estimator which captures the difference in the average outcomes between the UK and the countries in the donor pool after accounting for time and group-specific effects. Having said that, the coefficients λ and δ serve different purposes and capture different aspects of the treatment effect.

To perform inference on the results of the synthetic control method in the context of Brexit analysis, the following steps are undertaken:

1. Construction of Confidence Intervals: Confidence intervals are constructed for each trajectory of the treated unit in comparison to its synthetic counterpart. The con-

fidence interval is calculated based on 2 standard deviations of the pre-Brexit gap between the UK and the synthetic UK in order to have a 95% confidence interval. This approach allows for assessing the significance of the differences between the UK and its synthetic counterpart in the post-Brexit period. This inference approach is based on the work of Born et al. (2019).

2. In-Time Placebo Test: A placebo test is conducted by assigning the Brexit vote to a false year. The purpose of this test is to evaluate whether there would be any significant divergence between the UK and its synthetic counterpart after this false year. If the false assignment of the Brexit vote leads to a notable and significant divergence, it raises doubts about the main findings.

3. Country Placebo Test: Another placebo test is performed by assigning the Brexit vote to each of the countries in the control group. The gap between the treated unit (UK) and its synthetic counterpart is calculated based on this assignment. The objective of this test is to assess whether the observed gap between the UK and its synthetic counterpart is driven by the Brexit vote or by other random fluctuations within the control group.

4. Calculations of Fisher p-value based on country placebo effects: Abadie et al. (2010) propose an approach to assign the treatment to each country in the control donor pool and estimating the “country placebo effect”. For each control unit, the ratio of post-treatment MSE to the pre-treatment MSE is computed as the following:

$$R_j = \frac{MSE_{j,post-Brexit}}{MSE_{j,pre-Brexit}} = \frac{\left(\sum_{t=T_{0+1}}^{T}\left(Y_{jt}-\hat{Y}_{jt}^{N}\right)^2\right)/\left(T-T_0\right)}{\left(\sum_{t=1}^{T_0}\left(Y_{jt}-\hat{Y}_{jt}^{N}\right)^2\right)/T_0},$$

Where $\left(Y_{jt}-\hat{Y}_{jt}^{N}\right)$ is the treatment effect. Having large value of the MSE in the post-Brexit period indicates the Brexit effect on the treated unit. To validate that the gap in the post-Brexit period is the Brexit effect for the UK, the p-value is then calculated by comparing the value of R_1, the MSE ratio for the UK, to that of

$R_j \, \forall j = 2, \ldots, J+1$ control units, :

$$p_1 = \frac{\sum_{j=1}^{J+1} I_+ \left(R_j - R_1\right)}{J+1}, \tag{2.2.5}$$

Where $I_+(.)$ is an indicator function that takes a value of one for nonnegative arguments and zero otherwise. That is to say, the indicator function can be written as $1\left(R_j \geq R_1\right)$ Comparing the value of this statistic for the UK, R_1, to that of all other units, the p-value measures the chances of obtaining a ratio as high as R_1 if one is to pick a country randomly from the sample of countries in the donor pool. If the p-value is less than some pre-specified significance level, such as the the 0.1 traditional value then the null hypothesis that the treatment has no effect can be rejected implying that there is some regions with a non-zero effect for some period.

These steps help provide additional evidence and address potential concerns regarding the reliability of the main findings derived from the synthetic control method in the context of Brexit analysis.

2.3 Results and Discussion

2.3.1 Results

Total Manufacturing:

The employment of the manufacturing sector is particularly vulnerable to the impacts of Brexit referendum due to several key factors such as that this sector relies heavily on international trade, with a significant portion of its output exported to EU27 markets. The potential disruptions in trade relationships, such as the imposition of tariffs or non-tariff barriers, could lead to decrease in demand for UK-manufactured goods in the EU, thereby impacting employment levels. In addition to this, the manufacturing sector has historically benefited from the availability of European skilled labour to fill positions requiring specialised skills and expertise. The uncertainty associated with the future work arrangements for EU workers in the UK following Brexit referendum has created concerns about potential labour shortages. Any changes to the restrictions on the movement of labour may disrupt supply chains and create uncertainties for business that are not limited to closing of businesses premises, shifting business operations outside of the UK and most importantly disruptions to the value added within the manufacturing sector.

Total manufacturing sector employment is defined as the sum of the employment of all the manufacturing divisions. I construct synthetic UK manufacturing employment to measure how the Brexit vote affected the UK's manufacturing sector. As in equation 2.2.1, the synthetic UK is constructed with a weighted average of countries in the donor pool. The weights assigned to each country in the donor pool is based on matching the vector of pre-intervention characteristics of the treated unit $\mathbf{X}_0$ to that of the control group $\mathbf{X}_1$. These characteristics are manufacturing sector employment share index (2015=100), manufacturing value added, manufacturing sector establishments share and all the lags of dependent variable which is the log of employment index. These variables are averaged for the pre-intervention period. The pre-intervention period takes place from 2013 Q1 to 2016 Q1. Table 2.3 shows a comparison between the pre-intervention characteristics of the actual UK, synthetic UK and a population weighted average of all the 16 countries in the donor pool. As illustrated, the synthetic UK is closely matched to the actual UK than the weighted average of all the countries in the donor pool.

Economic predictor	UK	Synthetic	Sample Mean
Establishments Share	95.982	101.043	101.946
Employment Share	100.079	100.223	98.424
Manufacturing Value added	10.524	10.531	18.344
Log Employment Index (Dependent Variable)	4.591	4.593	4.622

Table 2.3: Employment Index Predictor Means Matching.

Manufacturing sector employment share index (Eurostat), Manufacturing Value Added (MVA) (UNIDO), the log of employment index (dependent variable) are averaged for the pre-intervention period 2013 Q1 to 2016 Q1.

Table 2.4 below shows the weights assigned to each country in the donor pool to construct the synthetic UK total manufacturing sector's employment. Almost 80% of the Synthetic UK total manufacturing sector employment is formed by countries such as Hungary, Germany, Norway, Denmark and Spain. This weighted combination makes sense as countries such as Hungary and Norway which account for almost 50% of the control group are similar to the UK in the aspect that Hungary is not part of the Eurozone, while Norway is not an EU member, but part of the European Economic Area (EEA).

It is important to note that the weights assigned to countries in the control group are based on matching the covariates between the treated unit and the control group and they refrain from any structural interpretations.

Country	Synthetic control weight
Hungary	42.4%
Germany	16.3%
Norway	6.5%
Denmark	6.4%
Spain	6.3%
France	4.4%
Belgium	3.3%
Italy	3.3%
Finland	2.9%
Greece	2.6%
Luxembourg	0%
Czech Republic	0%
Ireland	0%
Netherlands	0%
Austria	0%
Portugal	0%

Table 2.4: Synthetic Control weights for UK total manufacturing index.

Country weights to construct synthetic UK total manufacturing index. The weights illustrate the importance of country j in the donor pool to approximate the UK.

Figure 2.3 displays the plot of the natural logarithm of the total manufacturing sectors employment index for the UK and synthetic UK. The synthetic UK closely and consistently mimics the evolution of the UK series during the pre-Brexit vote period. The figure suggests the negative impact of the Brexit vote on the employment of the total manufacturing sector can be visible following the Brexit referendum till the end of the sample. The advantages of using the synthetic control method is that during the pre-Brexit vote period, the performance of the synthetic UK is almost identical to that of the UK. Therefore, one can conclude that the total manufacturing employment significantly underperformed its synthetic counterpart following 2016 Q2 until the end of the sample, capturing the negative impact of the Brexit vote. This significance is described by the shaded area which is 2 standard deviations of the difference between the UK and the synthetic UK in the pre-Brexit period.

Table 2.5 shows the estimation of the OLS regression of equation 2.2.3 and the DiD equation 2.2.4. As in Zhuang et al. (2023), DiD can be used as a robustness check for the synthetic control method to measure the average treatment effect on the treated group.

In this table I estimate the average impact of the Brexit vote on the total manufacturing employment in percentage points. Column (1) shows the estimation of the coefficient λ which shows the Brexit vote effect as the difference between the UK and its synthetic counterpart. While columns (2) and (3) show the estimation of the coefficient δ which is the DiD estimator, estimating the effect of the Brexit vote on the employment. As illustrated, the results show that in the post-Brexit period, the average loss in the manufacturing sector employment 3.6%, according to the synthetic control specification. For the DiD specification, the coefficient δ is very small and not statistically significant. This loss in the manufacture sector employment is equivalent to a loss of 91,522 jobs. This decrease in the employment in the manufacturing sector can be attributed to the drop of EU migrant skilled labour and the shift of UK business operations to EU destinations (Breinlich et al. 2020). This finding is consistent with the results of Born et al. (2019) and Fetzer and Wang (2020), who find that Brexit had a negative impact on the aggregate and regional UK economic activity using synthetic control methodology.

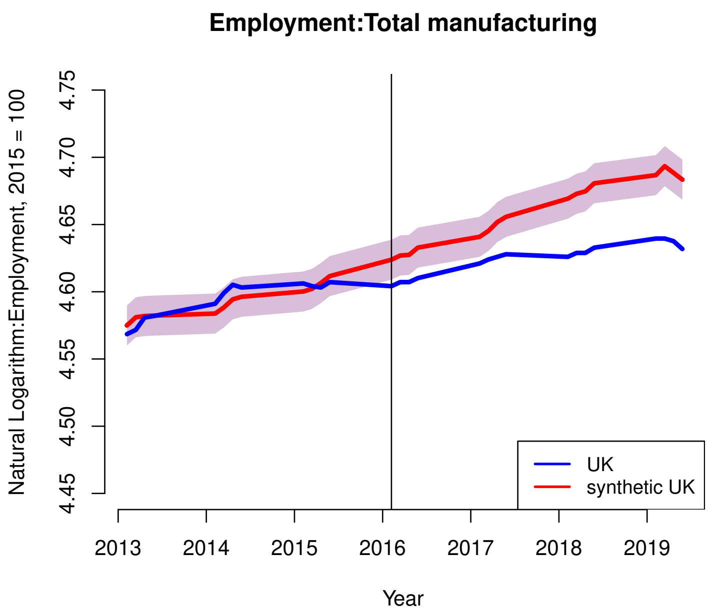

(a) Employment

Figure 2.3: Impact of the Brexit vote on the total manufacturing sector.

I follow Born et al. (2019) in constructing the shaded area as 2 standard deviations of the difference between the UK and the synthetic UK in the pre-Brexit period.

Given the complexity associated with the nature of the impact of the Brexit vote on employment, it is challenging to identify the causal effects of the Brexit vote using

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^N$	Y_{it}	
	(1)	(2)	(3)
d_t	-0.036^{***}	-0.006	
	(0.004)	(0.007)	
UK		-0.011	
		(0.008)	
$d_t \times UK$		-0.019	-0.005
		(0.013)	(0.012)
Constant	-0.001	3.984^{***}	
	(0.003)	(0.063)	
EstablishmentsShare		0.004^{***}	0.003^{***}
		(0.0005)	(0.0005)
EmploymentShare		0.003^{***}	0.005^{***}
		(0.0004)	(0.0005)
GFCFShare		0.0002	-0.00002
		(0.0001)	(0.0002)
MVA		0.001^{*}	0.023^{*}
		(0.0003)	(0.005)
Fixed Effects		No	Yes
Observations	28	273	273
R^2	0.714	0.403	0.462
Adjusted R^2	0.703	0.387	0.378
Residual Std. Error		0.027 (df = 265)	
F Statistic		25.503^{***} (df = 7; 265)	40.409^{***} (df = 5; 235)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 2.5: Effect of the Brexit vote on the UK total manufacturing sector employment.

Estimating the effect of the Brexit vote on the manufacturing sector employment. Column (1) shows the estimation of equation 2.2.3, where the gap between the UK and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

standard statistical methods. The use of synthetic control method becomes beneficial here to generate the counterfactual outcomes for the sectors having been affected from the referendum. The individual divisions in the manufacturing sector experience heterogeneous impact from the Brexit referendum, due to their different exposure to the EU and international trade. In the following sections, I will discuss the main findings of this paper. I will show the estimates of the impact of each division in the manufacturing sector, the divisions that gained from the Brexit referendum will be referred to as "Brexit vote winners" and the divisions that were affected from the Brexit referendum will be referred to as "Brexit vote losers". Some divisions did not display significant differences from their synthetic counterparts, I label them as unaffected divisions and their plots are in figure A.1 in the appendix.

As the synthetic control unit for each manufacturing division is constructed using different combination of control units, table 2.6 shows the highest non zero weights assigned to each country to construct the synthetic control units for each manufacturing division.

Winning manufacturing divisions

The manufacture of Non-metallic mineral products division is a winner in the post-Brexit period. Figure 2.4 illustrates that this division's employment significantly overperformed its synthetic counterpart in the post-Brexit period. Table 2.7 shows that the employment increased in the post-Brexit period by an average of 9.4%, which is equivalent to 7,242 extra jobs.

The manufacture of chemical products division is another winner in the post-Brexit period. Table 2.1 shows that this division is the second largest net exporter in the UK manufacturing sector, 60% of these exports go to the EU and 75% of the chemical imports, including essential raw materials, come from the EU (Chemical Industries Association 2018). Figure 2.4 displays that this division's employment overperformed its synthetic counterpart in the post-Brexit period. Table 2.7 shows that on average in the post-Brexit period, this division's employment increased by 5.8%, equivalent to an increase of 6,939 jobs in the post-Brexit period.

Lastly, manufacture of furniture had an increase in the employment during the post-

Manufacture Division	Synthetic control country weights
Wearing Apparel:	Belgium = 46.2%, Czech = 29.6%, Denmark = 23.9%
Chemical Products	Netherlands = 56%, Belgium 41.8%
Non Metallic Mineral Products	Norway = 55.2%, Denmark = 13.4%, France = 5.5%, Netherlands = 4.3%
Other Transport	Netherlands = 34.9%, Czech = 30.3%, Belgium = 22.2%, Denmark = 12.4%
Computer	Norway = 47.4%, Spain = 28.1%, Netherlands = 4.2%
Motor Vehicles	Czech = 41.9%, Netherlands = 34.8%, Hungary = 22.5%
Food products	Netherlands = 45.9%, Germany = 27.4%, Hungary = 13.3%, Spain = 10.3%
Electric Equipment	Hungary = 62.9%, Germany = 15.6%, Czechia = 6.2%, Norway = 4.6%
Printing	Spain = 71.2%, Hungary = 15.2%, Belgium = 8.3%, Norway = 2.0%
Paper and paper products	Czech = 61.7%, Belgium = 33.6%%, Netherlands = 1.7%
Machine and Equipment	Belgium = 70.5%, Greece = 13.0%, Italy = 3.9%, Spain = 3.6%
Fabricated Metal Products	Czech = 41.20%, Greece = 17.5%, Hungary = 16.6%, Germany= 7.4%
Rubber	Denmark = 53.6%, Hungary = 32.6%, Greece = 2.7%, Norway = 2.0%
Textiles	Greece = 76.4%, Hungary = 18.0%, France =1.3%, Belgium = 1.2%
Wood	Hungary = 56.0%, Norway = 34.9%, France = 2.6%, Belgium = 2.4%

Table 2.6: Synthetic Control weights for each UK manufacturing division index of employment.

Country weights to construct synthetic UK for each manufacturing division. Weights are estimated based on matching the vector of pre-intervention characteristics of the treated unit $\mathbf{X}_0$ with that of the control group $\mathbf{X}_1$. The weights illustrate the importance of country j in the donor pool to approximate the UK.

Brexit period, this division's employment increased by 5.5%, equivalent to an increase of 8,684 jobs in the post-Brexit period. As illustrated in figure 2.4, the employment index series for the UK overperformed its synthetic counterpart in the post-Brexit period.

Losing manufacture divisions

The greatest losing manufacturing division is the manufacture of wearing apparel division. This division relies heavily on foreign trade, where the UK imports material from Asia and low cost countries, and exports high quality heritage products that have strong brand recognition. Table 2.1 illustrates that this division is a net importer division. Figure 2.5 displays the significant contemporaneous negative impact on this division's employment following the Brexit vote. Table 2.7 illustrates that on average, this sector's employment

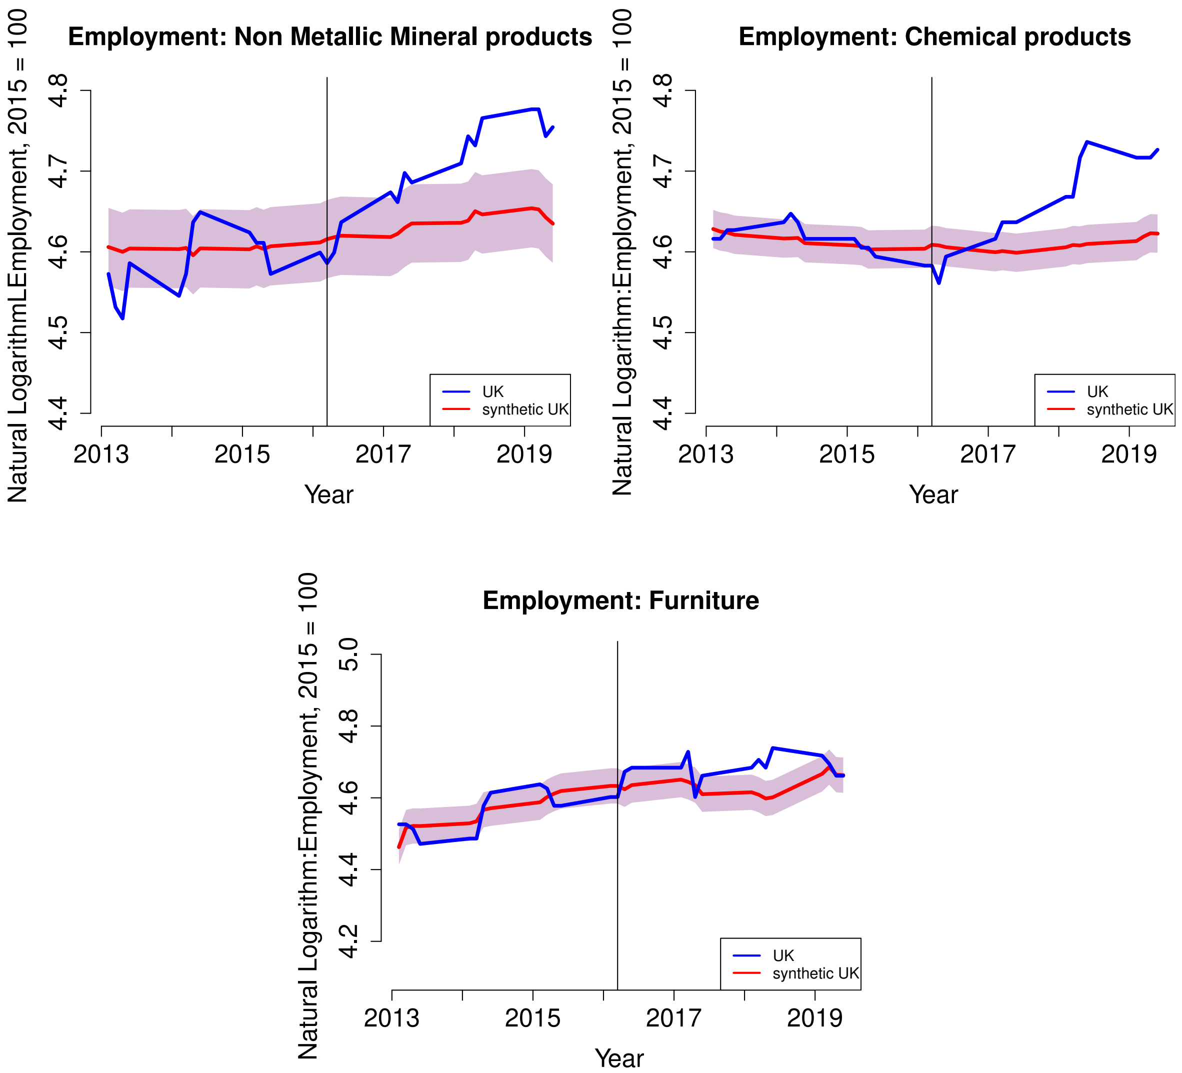

Figure 2.4: The winning manufacturing divisions.

I follow Born et al. (2019) in constructing the shaded area as 2 standard deviations of the difference between the UK and the synthetic UK in the pre-Brexit period.

dropped by 27.6% in the post-Brexit period, which is equivalent to a loss of 7,486 jobs.

The manufacture of other transport equipment includes industries such as manufacturing of ships, boats and yachts, the manufacture of railway rolling stock and locomotives, air and spacecraft. Figure 2.5 displays that the UK's employment underperformed its synthetic counterpart in the post-Brexit period. On average in the post-Brexit period, employment of this division decreased by 4.8%, which is equivalent to a loss of 7,583 jobs.

The manufacture of motor vehicles division's employment is strongly negatively affected in the post-Brexit period. Figure 2.5 displays that in the post-Brexit period this division's employment significantly underperformed its synthetic counterpart. Table 2.7

shows that this division's employment declined by an average of 9.3%, equivalent to 14,878 jobs lost.

The manufacture of computer, electronic and optical products division relies heavily on the international trade for imports of material inputs and exports of final products. As figure 2.2 shows, following the Brexit vote, the input prices increased in a similar manner to most of the manufacturing division, this put a pressure on the division's profitability. Figure 2.5 displays that following the Brexit vote this division's employment underperformed its synthetic counterpart significantly. Table 2.7 shows that this division's employment declined by an average of 5% in the post-Brexit period. In addition to this, the division was also negatively affected due to the reduction in both foreign and domestic capital investment following the Brexit referendum and the lack of business confidence that resulted in a decrease and postponement in investments in technological innovations (Breinlich et al. 2020).

The manufacture of electrical equipment division is also negatively affected in the post-Brexit period. Figure 2.5 displays that the UK and synthetic UK are on the same trend and have a very close fit in the pre-Brexit period. Following the Brexit vote, the UK series significantly underperformed its synthetic counterpart until the end of the sample. Table 2.7 illustrates the negative impact this division encountered as the employment decreased by an average of 18.3% in the post-Brexit period, equivalent to a loss of 14,210 jobs.

The manufacture of printing products division is the second largest division in terms of establishments as illustrated in table 2.1. Figure 2.5 displays that this division's employment decreased significantly following the Brexit vote. Table 2.7 shows that in the post-Brexit period, the UK decreased by an average of 13.9%.

Column (2) of table 2.7 shows the DiD estimator of equation 2.2.4 with two ways Fixed Effects (FE) regression. Tables A.1 to A.11 show the detailed estimation of these regressions.[9] For most of the manufacturing divisions, the coefficient λ of equation 2.2.3 and the coefficient δ have the same sign which is consistent on the effect of the Brexit referendum on different manufacturing divisions. However, the magnitude of the coeffi-

[9] I estimate the regression without controlling for covariates due to missing observations in the division specific economic covariates after 2017Q1 for most of the manufacturing divisions.

cient is different as each of the equation models the effect of the Brexit referendum in a different way. The former looks at the effect of the Brexit referendum in the context of the gap between the UK and its synthetic counterpart i.e. the effect of the referendum on the divergence between the two series which are assumed to be the same in the absence of the Brexit vote. While the latter looks at the effect of the treatment on the dependent variable, the log employment index.

In figure A.1 in the appendix, I show the unaffected manufacturing divisions, i.e. the manufacturing divisions that did not exhibit significant change in the post-Brexit period as well as they have insignificant regression coefficients in table 2.7. These manufacturing divisions are manufacture of machinery and equipment, manufacture of fabricated materials, manufacture of rubber products, manufacture of textiles, manufacture of wood and products of wood and cork, manufacture of food products and manufacture of paper and paper products.

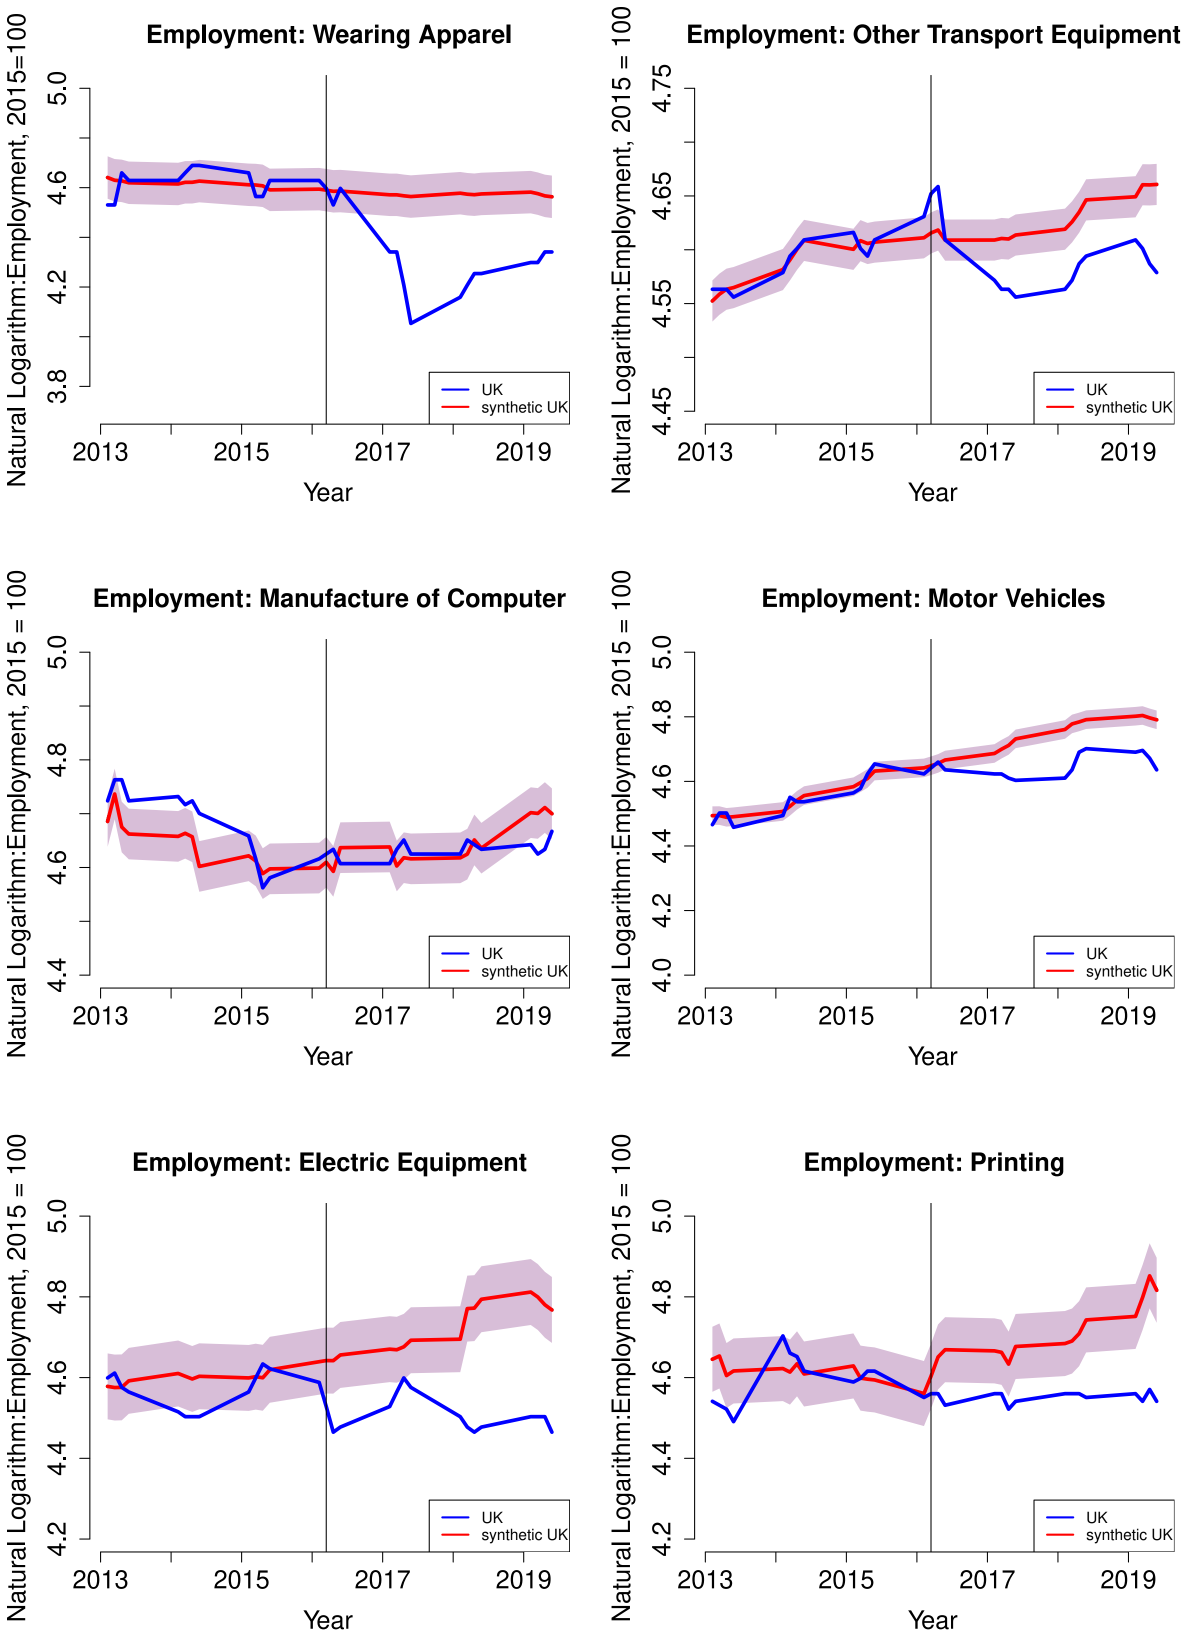

Figure 2.5: The losing manufacturing divisions.

I follow Born et al. (2019) in constructing the shaded area as 2 standard deviations of the difference between the UK and the synthetic UK in the pre-Brexit period.

Affected Division	(1) Impact (λ)	(2) DiD (δ)	(3) Number of employees	(4) Unaffected Division	(5) Impact (λ)	(7) DiD (δ)
Wearing apparel	-0.276^{***} (0.039)	-0.341^{***} (0.031)	-7,486	Machinery and Equipment	0.0001 (0.015)	-0.018 (0.032)
Other transport equipment	-0.048^{***} (0.009)	-0.043 (0.031)	-7,583	Fabricated metals	0.001 (0.011)	0.025 (0.021)
Motor vehicles	-0.093^{***} (0.013)	0.093^{***} (0.004)	-14,878	Wood and products of wood and cork	0.010 (0.026)	0.024 (0.025)
Chemical products	0.058^{***} (0.014)	0.002 (0.017)	6,939	Rubber and plastic	-0.033 (0.030)	-0.027 (0.023)
Computer	-0.050^{***} (0.015)	-0.078^{***} (0.024)		Food products	-0.017 (0.017)	0.005 (0.012)
Printing	-0.139^{***} (0.025)	0.002 (0.006)		Textiles	0.043 (0.027)	-0.004 (0.039)
Non metallic mineral products	0.094^{***} (0.016)	-0.090^{***} (0.019)	7,242	Paper and paper products	-0.019 (0.013)	-0.023 (0.035)
Furniture	0.055^{***} (0.016)	0.096^{***} (0.025)	8,684			
Electrical equipment	-0.183^{***} (0.028)	-0.071^{***} (0.022)	-13,511			

Table 2.7: Regression estimates of the impact of Brexit vote.

Column 1 reports the OLS estimates of equation 2.2.3, of log employment for each manufacturing division using HAC-SE regression. The standard errors are reported in parentheses. The significance codes are *p<0.1; **p<0.05; ***p<0.01.
Columns (2) shows the estimation of the DiD equation: $Y_{it} = \beta_0 + \beta_1 UK + \beta_2 d_t + \delta(UK \times d_t) + \epsilon_{it}$
Column (3) is the multiplication of the column 1 and the average number of employment per manufacturing division in the post-Brexit period obtained from the ONS.

2.4 Robustness Check

In this section, I perform a series of robustness tests to check for the validity of the results I obtained for the UK total manufacturing sector employment. I perform time placebo test, country placebo test and Mean Square Error (MSE) test and leave one control unit out test. My results show that the variation between the UK and the synthetic UK series after 2016 Q2 can be attributed solely to the impact of the Brexit vote.

2.4.1 Time placebo test

I assign the Brexit vote to a different period, 2014 Q1 and investigate if the UK had any significant changes in comparison to its synthetic counterpart. Figure 2.6 displays a plot of the UK and synthetic UK total manufacturing sector employment, with a placebo Brexit taking place in 2014 Q1. The plots show that there is no significant divergence between the UK and its synthetic counterpart after 2014 Q1. This provides evidence that the divergence in figure 2.3 is due to the impact of the Brexit vote. I also perform a time placebo test for each of the manufacturing divisions. Figure A.2 shows that when assigning the Brexit vote to a different year, the manufacturing divisions did not exhibit any significant change. Therefore, the impact illustrated in section 2.3.1 can be attributed as a causal impact of the Brexit vote on the manufacturing divisions employment.

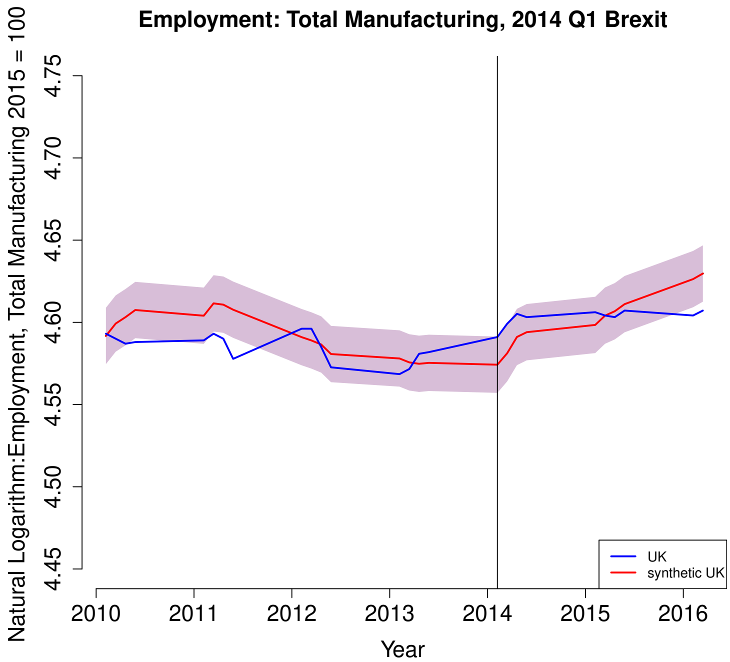

Figure 2.6: In time placebo test.

Assigning the Brexit vote to a different period, 2014 Q1. The in-sample period used in estimation of the weights is from 2010 Q1 to 2013Q4. The predictors used for matching are GFCF share, employment share, manufacturing value added, the employment index. The predictors are averaged for the period 2010 Q1 to 2013 Q4. I follow (Born et al. 2019) in constructing the shaded area as 2 standard deviation of the difference between the UK and the synthetic UK in the pre-Brexit period.

2.4.2 Country placebo test

To assess the significance of my estimates, I conduct permutation test by applying the synthetic control method to the countries in the control group. As the countries Hungary, Germany, Norway, Denmark and Spain account for almost 80% of the synthetic UK, in each iteration, I assign the Brexit vote to one of these countries, shifting the UK to the control group. That is, I am interested in the gap, $Y_{1t} - \hat{Y}_{1t}^{N}$, between each country of these countries and its synthetic counterpart. This gap is an indicative of the treatment effect. The main purpose of this exercise is to understand whether or not the estimated effect of the Brexit vote discussed earlier is large relative to other countries in the control group that were not affected by the vote. The coloured lines in figure 2.7 show the difference between each country in the control group and its synthetic counterpart, while the black line is the gap between the UK and its synthetic counterpart. As shown in figure 2.7 the gap for the UK's employment is more negative than the gaps for the countries in the control group. Although the gap between Norway and its synthetic counterpart is negative after 2016 Q2, this country has a fluctuating series in the whole sample, with a negative gap in the pre-Brexit period as well. That is to say, no other countries in the donor pool were affected to the same extent as the UK.

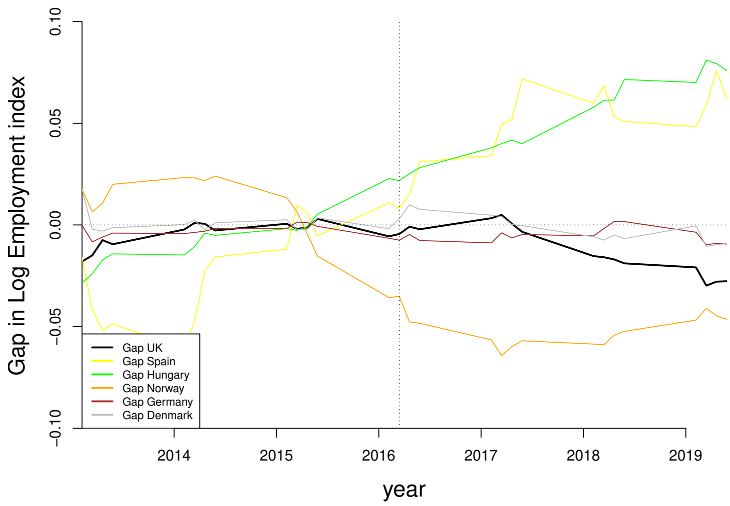

Figure 2.7: Country placebo test.

Employment gaps in the UK and placebo gaps for the main countries with non-zero weights. For each iteration, the UK is removed from the estimations of the control group country weights.

To verify my results that none of the countries in the control group were affected in a similar way to the UK, I calculate the ratio of the post-Brexit vote MSE to the pre-Brexit MSE in the next section and the p-value to test the null hypothesis that the Brexit vote had no effect of on the UK manufacturing sector employment.

2.4.3 Mean Square Error test

Another way to measure the robustness of the results, is to compute the ratio of the post-Brexit MSE to the pre-Brexit MSE for the United Kingdom and the countries in the control group. The MSE measures the lack of fit between the treated unit and its synthetic counterpart. This ratio can be calculated as:

$$\frac{MSE_{j,post-Brexit}}{MSE_{j,pre-Brexit}} = \frac{\left(\sum_{t=T_0+1}^{T}\left(Y_{jt} - \hat{Y}_{jt}^{N}\right)^2\right)/\left(T - T_0\right)}{\left(\sum_{t=1}^{T_0}\left(Y_{jt} - \hat{Y}_{jt}^{N}\right)^2\right)/T_0}, \tag{2.4.1}$$

Where $\left(Y_{jt} - \hat{Y}_{jt}^{N}\right)$ is the treatment effect. This ratio of the MSE is estimated for each country in the control group. Having big value for the ratio indicates that there is a divergence between the trajectories of the treated unit and the synthetic control unit. Since the pre-intervention period had a close fit between the UK and its synthetic counterpart, thus the pre-brexit vote MSE is small. The ratio of the post-Brexit vote MSE to the pre-Brexit vote MSE indicates the effect of the Brexit vote in the post 2016 Q2 period. As illustrated in the previous section, I assign the treatment, Brexit vote, to all the countries in the control group. I then estimate the ratio of the post-Brexit vote MSE to the pre-Brexit vote MSE for each of the control units as illustrated in 2.4.1. This algorithm calculates the ratio for all the countries in the control group including the ones with very small weights. I compare the ratio of the post-Brexit vote MSE to pre-Brexit vote MSE of these countries to that of the United Kingdom (UK). Having a large ratio indicates that the gap between the treated unit and its synthetic counterpart in the post-Brexit vote period is greater than the gap in the pre-Brexit vote period. Figure 2.8 displays that the UK stands out in comparison to the other donor countries. As displayed, the UK has the largest ratio, whereas the post-Brexit gap is 21 times the pre-Brexit gap. Therefore, this illustrates the causal impact of the Brexit vote on the aggregate manufacturing sector employment.

To calculate the p-value as illustrated in 2.4.1, if one were to pick a country at random from the sample the chances of obtaining a ratio as high as that of the UK is $\frac{1}{12} \approx 0.08$. Therefore a rejection of the null hypothesis that the Brexit vote had no impact of the UK total manufacturing sector employment. When doing the same robustness exercise for the winning and losing manufacturing divisions figures A.3 display the results. As displayed that the UK outperforms all other control units for all the divisions with exception manufacture of non metallic mineral products, manufacture of computer and manufacture

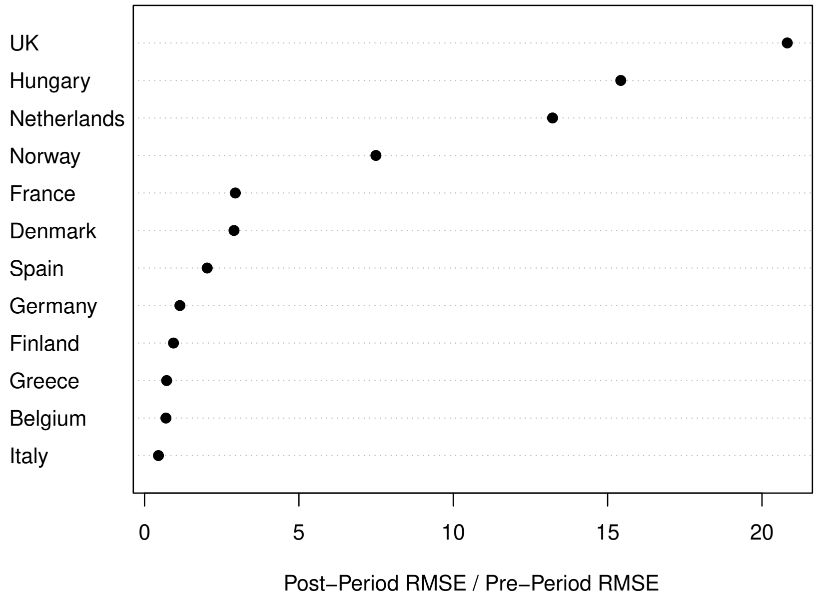

Figure 2.8: Ratio of post-Bexit MSE to pre-Brexit MSE.

Synthetic control algorithm is estimated for the United Kingdom and the 12 countries of the control group. Having biggest value of MSE ratio for the UK indicates the causal effect of the Brexit vote on manufacturing sector employment. The Pre-Period is from 2013 Q1 to 2016 Q2. The Post-Period is from 2016 Q3 to 2019 Q4.

of electric equipment. That is to say, the p-value for the divisions where the UK post-Brexit MSE ratio is the greatest is in the range of 0.08 to 0.091, hence the null hypothesis that the Brexit vote had no effect on these sectors employment can be rejected.

2.4.4 Leave one control unit out test

Another robustness check that can be made is to test the sensitivity of our main results to changes in the country weights W^*. If one to choose three countries to construct the synthetic UK total manufacturing sector employment index, these countries will be Hungary, Norway and Spain. I iteratively re-estimate the synthetic UK total manufacturing sector employment by omitting in each iteration one of the countries from the three countries of the control group and recalculate the weights. Alternatively, when re-estimating the model, synthetic UK can be formed of Hungary and Norway only. Figure 3.13 displays that my results of the total manufacturing sector employment is robust to the exclusion of any particular country from our control group. Although there can be small variation between the synthetic UK (in red) and the individual iterations (in grey), and synthetic

UK leave Hungary out. The conclusion remains the same, that the Brexit vote caused a decrease in employment in the UK manufacturing sector even when taking out the country with the highest weight.

Synthetic control combination	Countries and W-Weights		
Three units control group	Hungary 46%	Norway 35%	Spain 22%
Two units control group	Hungary 52%	Norway 48%	
One unit control group	Hungary 100%		

Table 2.8: Leave one control out exercise.

The baseline model is re-estimated iteratively to show the countries that contribute the most to the sparse version of the synthetic UK.

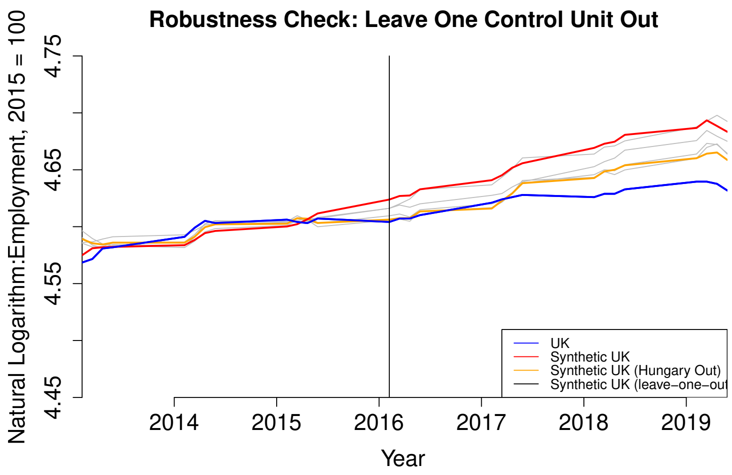

Figure 2.9: Leave one control unit out

2.4.5 Change in Covariates

I change the specification of the synthetic control estimation such that the synthetic control unit is estimated using two approaches. The first approach is based on economic covariates along with all the pre-intervention lags of the dependent variable (Billmeier and Nannicini 2013; Bilgel and Galle 2015; Stearns 2015; Villar and Papyrakis 2017; Pellegrini et al. 2021; Papyrakis et al. 2023). The second approach is based on using all the pre-intervention lags only, without the inclusion of economic covariates (Ferman et al. 2020; Breinlich et al. 2020; Botosaru and Ferman 2019). As displayed, across all the specifications and the original estimation of the synthetic control in figure 2.3 the results are unchanged, the UK significantly under perform its synthetic counterpart following the Brexit referendum till the end of the sample.

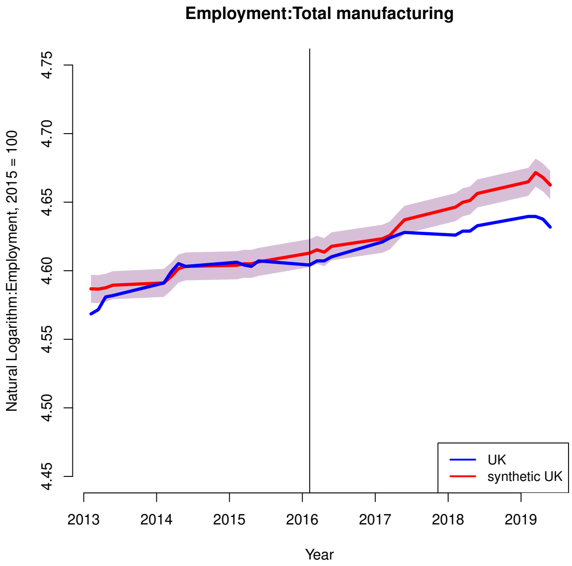

(a) Synthetic UK, with economic covariates and all lags of dependent variable

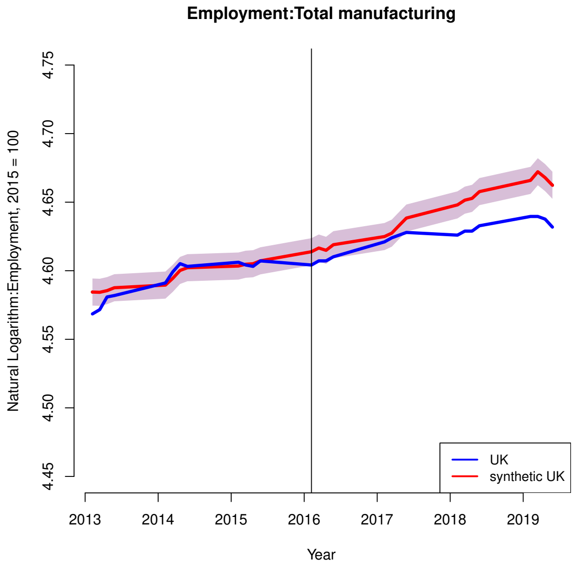

(b) Synthetic UK, with all lags of dependent variable as predictors and no economic covariates

Figure 2.10: Covariates Robustness Check

2.5 Discussion

Having shown the results and performed a series of robustness checks to validate the main results of the paper, I now present the discussion to explain these results and illustrate how practitioners and policymakers can benefit from the use of these results. As highlighted above, the main question that this paper aims to answer is "What if the Brexit vote never happened?" In another words, How will the manufacturing sector as a whole and key manufacturing division's employment be looking like in the absence of the Brexit vote. As estimated above, the referendum resulted in a 3.6% decrease in employment in the manufacturing sector. This can be explained by factors such as losing contracts, increasing imports costs and economic uncertainty associated with the Brexit deal negotiations. Using the NiGEM model of the NIESR, Erken et al. (2018) shows that Brexit could result in unemployment raising from 4.6% in 2018 to 6.2% in 2020. Fetzer and Wang (2020) find that local areas that depend highly on manufacturing were affected the most from the Brexit vote. Similarly, Davenport and Levell (2022) show that men are more exposed than women to Brexit effect as they are more likely to work in manufacturing industries that are affected by the Brexit trade barriers. More than one in five British manufacturing companies are expected to lay off workers as a result of the Brexit referendum (Politico 2018). A report by the Chartered Institute of Procurement and Supply data found that 11% of manufacturers say that they have already lost contracts as a result of the Brexit referendum (The Guardian 2018). As illustrated by Breinlich et al. (2020) that as a result of the Brexit vote, outward FDI from the UK to the EU has increased as firms shift their operations outside the UK. The vulnerability of each manufacturing division varies based on their share of trade with the EU as well as the share of EU labour in their labour force. The EU labour constitutes significant number of the British labour market. For instance, out of around 2.5 million jobs that were added to the UK in 2005-15, 2.2 million were supplied by immigrants, with nearly 60% originating from the EU (Kierzenkowski et al. 2016). A survey conducted with EU nationals working in the UK found that 52% of those earning £50,000-£100,000 said they would leave the UK or were thinking about it (KPMG 2017b). Minford (2015) postulates that as a result of Brexit, manufacturing sector will be eliminated with regard to the service sector. The chair of the CBI's manufacturing council, indicated that some manufacturing companies are putting their investment and hiring decisions on hold given the uncertainty over Brexit (Financial Times 2018a). This involves delay in capital investments for the UK manufacturing sector, this delay in the investment is similar to the option value of (Bloom 2009) that during uncertain times, firms reduce irreversible investment decisions. As highlighted by Born et al. (2019) and Bloom et al. (2019) the Brexit referendum constitutes an uncertainty shock to the economy as the details of the Brexit are unclear to the businesses. In a survey to business executives, Bloom et al. (2017) find that employment growth is likely to slow down with a chance of 30% reduction in employment in 2017Q4. Similarly, as illustrated in figure 2.11 below the proportion of the respondents who indicated that Brexit is either the largest source of uncertainty or among the top three sources of uncertainty increased overtime. This uncertainty weigh on investments and business hiring decisions for the manufacturing sector in general and different manufacturing divisions based on their exposure to the EU. It is possible to argue that some of the manufacturing divisions could have benefited from increase in domestic projects especially the projects that are related to government infrastructure projects during the post-Brexit period. These infrastructure investments such as the development of new roads, airports and rail networks require the engagement of various manufacturing divisions to supply the necessary materials, equipment and components. As the government prioritised infrastructure development in the years following the Brexit referendum, this is likely led to offsetting some of the negative impacts of the Brexit referendum on different sectors in the economy.

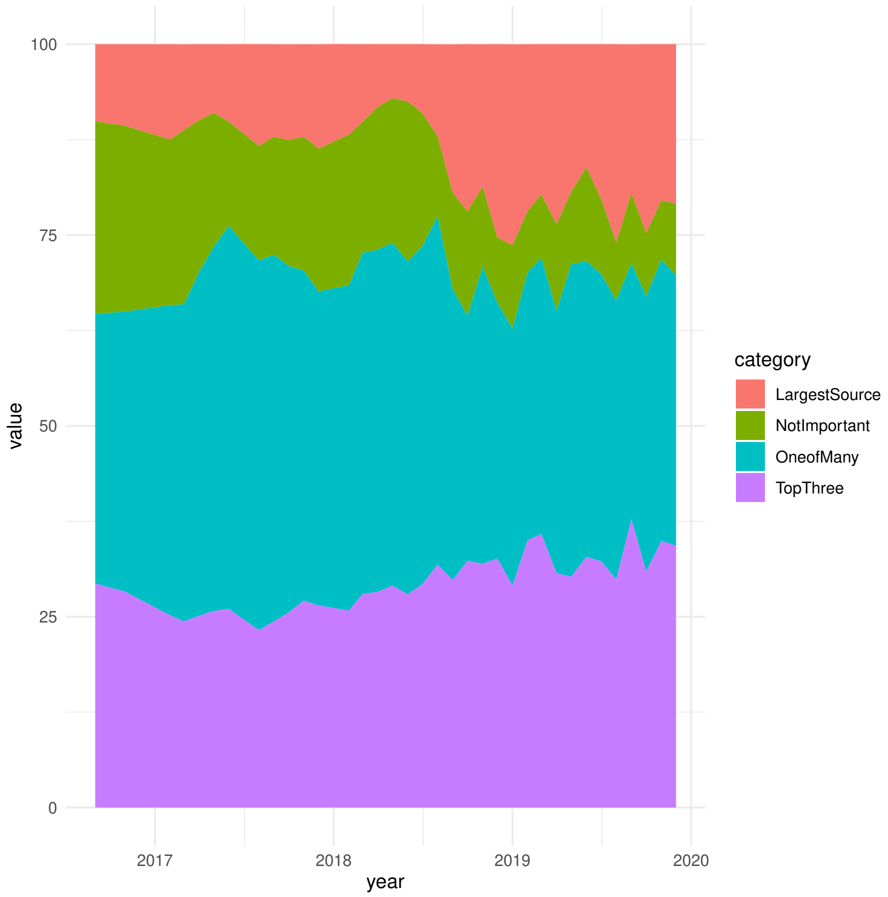

Figure 2.11: Brexit as a source of uncertainty for own business

Brexit as a source of uncertainty. Data are based on responses to the question: “How much has the result of the EU referendum affected the level of uncertainty affecting your business?”. The four response categories are: i) not important, ii) one of many drivers of uncertainty, iii) one of the top two or three drivers of uncertainty for our business, iv) the largest current source of uncertainty for our business. Source: Decision Making Panel (DMP), Bloom et al. (2019) .

Although the counterfactual analysis is conducted for the sixteen manufacturing divisions, I will be focusing my analysis on the two divisions that are the biggest winners, the manufacturing of non metallic mineral products and and the manufacture of chemical products and the two divisions that are the biggest losers, manufacture of wearing apparel and manufacture of motor vehicles. The four manufacturing divisions to be investigated have significant share of economic activity such that they account for almost 20% of the manufacturing sector labour force and almost 25% of the manufacturing sector output. Therefore, the effect of the Brexit referendum on these divisions needs to be investigated to understand why they are winners or losers.

In a survey by the Society for Motor Manufacturers and Traders (SMMT) that was conducted before the Brexit referendum, 70% of its members expect Brexit to have a negative or very negative medium to long-term impact on their business operations, SMMT (2014). The UK house of commons discussed the negative impact the Brexit vote may have on the manufacture of motor vehicles in general and the Automotive industry in particular, such that the Brexit vote will affect the supply of highly skilled labour and engineers from the European Union. The government acknowledged that in the short run, the current government engineering apprenticeships are unlikely to compensate for the shortfall resulted from the leakage of skilled European labour (House of Commons Business, Energy and Industrial Strategy Committe 2018). Around 10% of the workforce in the motor vehicles division is from EU27 countries, this number can increase to 30% in some parts of the industry (The UK in a Changing Europe 2020). That is to say, the Brexit referendum could have severe negative effect on the motor vehicles division, this corroborates with the findings illustrated in the previous section. In a similar manner, 58% of the British-made vehicles exported go to the EU (KPMG 2017a). These reports also highlight that automotive companies rely on the movements of skilled labour from European countries to the UK to support projects such as new vehicles launches. Restrictions on the free movement of labour can negatively affect this division. Financial Times (2019) reported that the UK car production almost halved in April 2019, after car companies shut down factories during the month in the expectation that the UK would have left the EU as planned on March 29. Honda and Nissan started to move some of their plants outside of the UK and some of the key players in the industry such as Toyota, BMW and Jaguar Land Rover have trade with the EU, and stress the importance of the free trade with the EU. Jaguar Land Rover confirmed to cut 4,500 jobs of its UK plant in 2019, while Nissan closed part of its Sunderland plant resulting in 2800 jobs lost (IBIS World 2020b). 90% of firms in the car manufacturing industry reported to the Society of Motor Manufacturers and Traders that their operating costs has increased as a result of leaving the EU, (The Guardian 2021). Hancké (2018) discussed that for the automobile industry, parts in an average car cross the English Channel several times before they end up in the final product. Introduction of tariffs as a result of Brexit will make car manufacturing more expensive, which could result in manufacturers to shift their production outside the UK.[10] This can be illustrated by figure 2.12, where as following the Brexit referendum, there has been a persistent decline in the output in the motor vehicle manufacturing industry measured by the SMMT volume percentage growth in motor vehicle manufacturing. As this division drives innovation and technological advancements, R&D in complex automotive engineering, design and sustainability, any changes in the structure of this division or the withdraw of big players could affect the competitiveness of the wider manufacturing sector. The results highlighted by the application of the synthetic control method to this division, which demonstrates that the employment in this division was severely affected, confirm the actual findings of the negative effect this division suffered from as a result of Brexit.

Manufacturing of wearing apparel division includes industries such as clothing and textiles. This division is highly reliant on trading with the EU. The industry worth £32 billion and employs more than 890,000 person. The EU-UK textile trade was worth £14 billion in the period 2014-2016, on average the UK imported £8.7 billion and exported £5.3 billion worth of products and 74% of the exports of the fashion and textile were to the EU (BFC 2019; Casadei and Iammarino 2021). In terms of employment, this division depends on EU talented labour such as fashion designers and sewing machinists and 70% of the London clothing manufacturing firms labour force comes from EU countries (UKFT 2019). Tariffs and duties related to clothing and textiles with the EU are high such that EU duties are 4-5% for yarns, 8% for fabrics and 12% for clothes. Therefore, any trade barriers imposed as a result of the Brexit referendum can affect this industry severely. A recent survey by the UK Fashion & Textile association found after the UK formally left the EU, 98% of UK fashion and textile companies have suffered difficulties

[10]The Guardian (2017), tracked the assembly of Mini Cooper in the UK, to find that the parts used in this car crossed the English Channel three times.

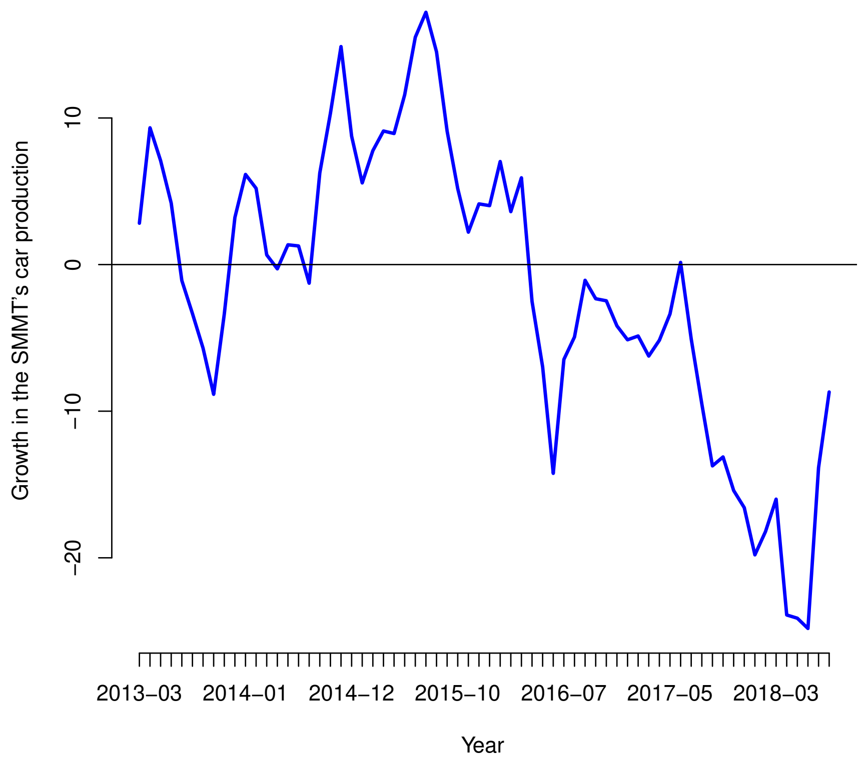

Figure 2.12: Growth in the SMMT's car production.

Manufacture of motor vehicles, trailers and semi-trailers. Source ONS.

during their trade with the EU, these difficulties are in terms of increase of trade tariffs, freight costs and customs clearing (UKFT 2021). Gasiorek et al. (2019) highlight that textiles, apparel and footwear also appear likely to experience relatively large declines in their output and exports due to their high reliance on the EU trade. The depreciation of the sterling pound caused an increase in the prices of imported materials. The clothing manufacturing industry in the UK had a decrease in its revenues in 2017 by 22.5% and in 2018 by 8.5% in comparison to the previous years. Similarly, the exports and imports of this manufacturing division decreased by their greatest amounts in 2016 and 2017 since 2008, as exports decreased by 23.4% and 37.3%, and imports decreased by 31.2% and 29.9% in 2016 and 2017 (IBIS World 2020a). This slowdown in this division's performance resulted in a decrease in the level of employment. Figure 2.13 shows that the monthly manufacture of wearing apparel output for this manufacturing division has decreased in the post-Brexit referendum period, reaching its lowest value of 159.1 million by December

2018. (Lawless and Morgenroth 2019) calculates the sectoral level elasticity to change in tariffs, they found that the wearing apparel division has a median tariff elasticity of -10.5%. Therefore, this reliance on the EU market and labour and the uncertainty associated with the Brexit referendum explain the synthetic control results of why this manufacturing divisions is the biggest loser.

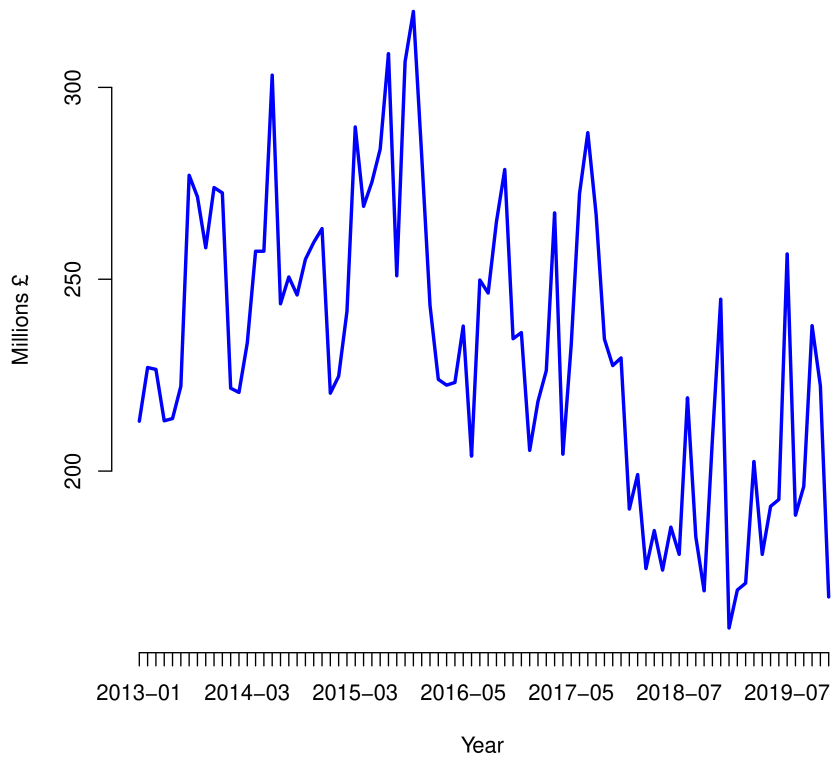

Figure 2.13: Manufacture of Wearing Apparel.

Values are in £Millions. Source ONS.

Manufacture of non-metallic mineral products is a winner manufacturing division in the post-Brexit referendum period. The divisions largest product segment is road construction and resurfacing products, this includes the manufacture of asphalt and road surfacing products, this makes the sector to be heavily reliant on public sector investment in major infrastructure projects. In addition to this, this division's subsectors are highly dependant on the construction market as they include manufacture of insulation materials, roofing and waterproofing materials, and bricks, mortar, glass. The UK government has invested heavily in infrastructure projects since the Brexit referendum. For example, the UK government invested £1.5bn on the upgrade of the A14 Cambridge to Huntingdon motorway as well as another £1.5bn part of constructing smart motorways across England (Highways England 2017). These infrastructure investments are made through British companies in this division such as Tarmac Group. Moreover, the rapid rise in new residential housing construction projects in the UK that was boosted by several government initiatives to encourage first time buyers to get on the property ladder such as the Help to Buy and Shared Ownership Schemes could have benefited some industries in this division. This resulted in an increase in the demand for construction materials, which are produced by some of the industries in this manufacturing divisions (ONS 2021). These developments in the domestic UK economy could explain the behaviour of the employment in this manufacturing division following the Brexit referendum, which shows that this manufacturing division is the biggest winner with an increase in its employment by 9.1%.

Manufacture of chemical products may have certain characteristics or attributes that make it less responsive to the negative effects of the Brexit referendum. For example, the sector relies heavily on exports, being the second largest net exporter manufacturing division (Chemical Industries Association 2018), the devaluation of the Sterling pound makes the exports of this division attractive in international markets. Additionally, this division may have a strong domestic market demand, resilient supply chains, or specialized production capabilities that allow it to adapt more effectively to changing economic conditions (Grylls 2020). These factors can contribute to maintaining or even increasing employment levels within the sector despite the uncertainties brought about by Brexit. This division has shown significant overperforming in comparison to its synthetic counterpart in the post-Brexit period. Although this positive impact could only be short lived, companies in the manufacturing of chemical products sector will face the burden of duplicating and reconstructing data on product safety and compliance. As the UK is departing from the EU, it will no longer follow the European Union's chemical regulation system known as REACH. The creation of UK REACH will require extensive and costly filings, potentially adding up to £1 billion in fees and expenses over several years. This regulatory change not only poses financial challenges but also introduces additional bureaucratic hurdles for businesses operating in the sector (The New York Times 2021).

While many reports published by industry and governmental bodies have looked at the size of the industries and their reliance on the EU, the analysis of this paper goes beyond this by using counterfactual analysis to estimate how employment in each sector could have looked in the absence of the Brexit referendum. This approach allows economists and policymakers to better understand the true extent of the Brexit loss or gain for each divisions to provide more understanding of the impact of the referendum on the manufacturing sector.

Government Economic Policy

The manufacturing sector plays a vital role for the UK economy, it contributes to the UK with a share of 10% of its output. It contributes to 50% of the UK exports and drives business-led research and development (R&D). As the UK departs from the EU and negotiates trade and labour movements arrangements, it will encounter both challenges and opportunities for the manufacturing sector. The UK policy makers should be made aware of the manufacturing divisions that have been affected in the three years following the referendum. HM Government (2017) highlighted in a white paper the need to support domestic manufacturers and businesses as the UK leaves the EU, this includes investment in sectors that were affected from the departure from the EU as well as the new “high value” manufacturing industries. The results obtained in the previous section of the paper provide policy implications for the government to support the affected sectors and to mitigate the potential negative consequences of Brexit. More specifically, the result suggest that the government can offer appropriate training and support to British labour to develop the technical and transferable skills needed to find work in the affected sectors. In addition to this, by understanding the divisions that were hit hard as a result of Brexit, the UK government support should not be limited to financial support it should consider promoting tax incentives, R&D subsidies and R&D tax credits, investing in High Value Manufacturing (HVM) catapult centres and skill development programs. This also involves the consideration of technology development and investing in new areas that are not reliant on EU policies and labour such as AI and robotics and green investments.

The UK government has been prioritising the manufacturing sector in its negotiations with the EU. The government needs to be informed about the divisions that were affected the most during the Brexit transition period so that they can support the affected divisions. Potential changes in immigration policies after Brexit may have impact on the manufacturing sector employment. Figure 2.14 below displays the new national insurance registrations for EU citizens in the UK. The figure shows that after the Brexit referendum, specifically after 2016 Q4, the number of new national insurance registrations dropped to the lowest levels and did not rebound until the end of the sample. This can be explained by factors such as the UK businesses slowed down after the Brexit referendum, UK businesses are shifting their operations outside the UK or to imply that the stock of European workers in the UK is decreasing, causing a shortage in skilled labour. The UK government should ensure that the new arrangements with the EU allow for the flexible movement of workers from the EU to avoid any skills gap in the affected division as training British labour to fill this gap could take time. This decrease of in flow of new EU labour to the UK could result in a weaker economy due to the loss of skilled labour who can move to other EU countries, this could result in reduction of the pool of skills which could efficiently be combined in the production. This will not only decrease productivity, but also could result in structural unemployment over the long run. The government needs to interfere in order to fill the shortage of skilled labour arising. For instance, the government can interfere by providing appropriate training and support to British labour, to develop the technical and transferable skills needed to find work in the affected sectors. In terms of the manufacturing sector, this paper helps the government's policymakers by identifying the manufacturing divisions that need support to hire skilled labour. The UK government announced the "Apprenticeship Levy Scheme", where companies with a pay bill of more than £3 millions to put aside an equivalent of 0.5 per cent of this total for levy-approved training schemes (Financial Times 2018b), the purpose of this campaign is to train labour with the needed skills to boost productivity. By identifying the manufacturing divisions that were affected by the Brexit referendum in terms of employment, the government can support the training for skilled labour. In addition to this campaign, the UK government launched National Manufacturing Competitiveness

Levels (NMCL) program in order to develop and strengthen the overall competitiveness of the UK manufacturing sector through providing targeted support to individual businesses. The results of this paper in terms of addressing the affected divisions can help the government in facilitating improvement across various aspects of operations to enhance the performance and growth potential of manufacturing sector.

The findings of this paper are relevant to the policymakers who aim to support the manufacturing divisions that are affected from the impact of the Brexit vote. The UK contributes roughly £8 billion pounds annually to the EU budget. As result of Brexit, the government will have £8 billion pounds net budgetary savings that could be spent on supporting manufacturing divisions to mitigate the negative impact of Brexit on the UK economy (Erken et al. 2018). The financial support that was introduced by the UK government during the transition period did not take in consideration any additional support to compensate for the loss of the EU skilled labour. For instance, in August 2019, the UK government outlined the introduction of operation Yellowhammer, a 2.1 billion pound contingency plan in response to the anticipated short-term disruption in the event of a no-deal Brexit.[11] On the other hand, the minister responsible for the management of a no-deal Brexit announced the introduction of a financial package known as operation Kingfisher to offer support to businesses and industries affected from the no-deal Brexit.[12] As a part of the "Get Ready for Brexit" campaign, the government announced financial helps for businesses to prepare for Brexit. The government announced the availability of 1.3 billion pounds to lenders in guarantee schemes - via the government-owned British Business Bank to support small and medium-sized businesses. The government also introduced the Brexit Readiness Fund to help business group and trade associations to support business preparing for Brexit. None of the government initiatives outlined extra support for sectors affected from the loss of labour due to Brexit, thus in this paper I shed light on the manufacturing sector to outline the extent of damage it suffered as a result of the Brexit vote. When the government considers an intervention, it should take in consideration the long term vision on how this can affect each division separately,

[11] Information about this plan is available on https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/831199/20190802_Latest_Yellowhammer_Planning_assumptions_CDL.pdf.

[12] The government did not announce the amount of money available in this bail out fund nor the channels on how the funds will be delivered to struggling businesses.

as each division reacts to the events of no deal Brexit heterogeneously. The degree of internationalisation of each division can affect the degree of financial support businesses can receive.

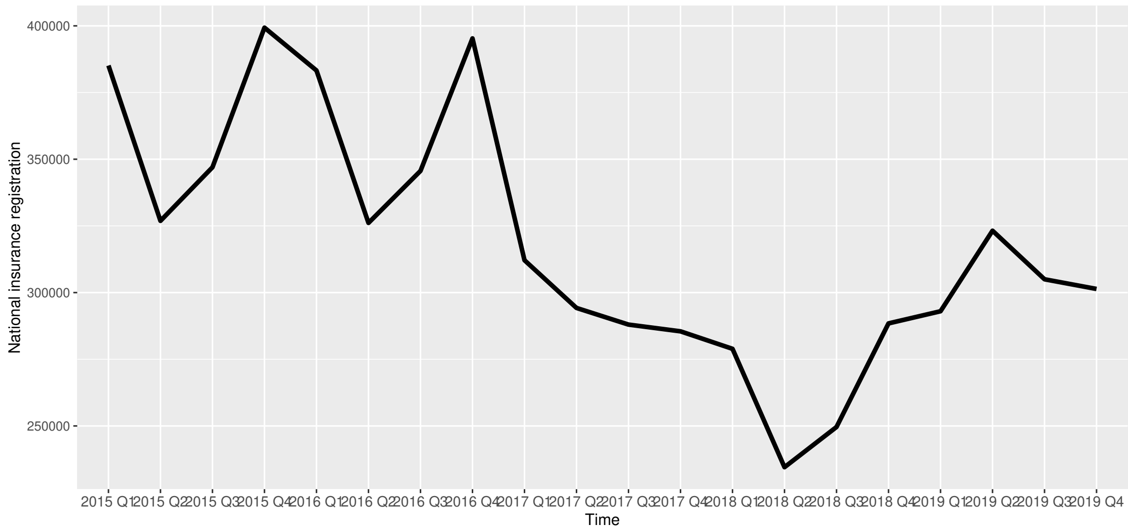

Figure 2.14: National insurance registrations for non-British nationals

Source: Office of National Statistics

2.6 Conclusion

In this paper, I investigated the impact of the outcome of the Brexit referendum on the employment in the aggregate UK manufacturing sector and sixteen dis-aggregate manufacturing divisions using synthetic control method. By providing empirical evidence and policy implications based on this analysis, this paper adds valuable insights to the ongoing debate surrounding Brexit and its impact on the UK economy.

The results show that the UK total manufacturing sector's employment decreased by 3.6%. However, this effect is heterogeneous across manufacturing divisions due to their different exposure to the EU. Six manufacturing divisions were negatively affected, with manufacturing of wearing apparel division the biggest loser. Three manufacturing divisions benefited with manufacture of non metallic mineral products is the biggest winner. These heterogeneous effects were due to several factors such as the slowdown in the manufacturing sector following the referendum, shift of business operations outside the UK to benefit from the access to the EU single market and the shortages of the EU migrant labour.

The results of this paper add to the existing body of knowledge on the impact of Brexit on the manufacturing sector in the UK by providing empirical evidence and policy implications. The paper mainly looks at the short run effects of the Brexit vote with the analysis stopping at 2019 Q4 to avoid the COVID-19 shock. The COVID-19 shock resulted in many businesses laying off their employees as well as employees not coming back to the labour force due to suffering from “long COVID-19”. It is hard to expand the analysis beyond 2020 as disentangling the effect of Brexit only will not be possible. This work is particularly relevant in the current political climate, as the UK continues to negotiate its future relationship with the EU and the manufacturing sector continues to face uncertainty regarding its access to skilled labour and markets. These results provide valuable information to policymakers about manufacturing divisions that require the most support from the impacts of the Brexit vote.

Chapter 3

Land Tax Holiday and House Prices: Evidence from the UK

Abstract

This paper provides evidence on the impact of the Stamp Duty Land Tax holiday on house prices in the UK using synthetic control methodology. I postulate that the tax holiday caused an increase in the UK house prices by annual rates of 2.11% and 2.96% in 2020 Q4 and 2021 Q1. On a regional level, the paper document regional heterogeneity in the impact of the tax holiday on house prices. All the UK large regions exhibited an increase in the house prices except for Greater London. The findings of this paper suggest that the pandemic caused city-dwellers to move out of big cities to the countryside and coastal cities, driving their house prices to increase.

3.1 Introduction

On the 8$^{\text{th}}$ of July 2020, the UK Chancellor of the Exchequer announced an immediate holiday on the Stamp Duty Land Tax (henceforth, SDLT), one of the oldest UK taxes, for England and Northern Ireland for properties sold at a price of up to £500,000. Scotland and Wales applied holiday on the Land and Building Transaction Tax (LBTT) and Land Transaction Tax (LTT) for properties up to £250,000. The tax holidays were scheduled to end on 31$^{\text{st}}$ of March 2021 for all four nations. However, England, Northern Ireland and

Wales decided to extend the tax holiday until 30th of June 2021, with a lowered threshold, £250,000, until the end of September 2021 and a reversion back to normal rates on 1st of October 2021, such that the first £125,000 are taxed at the nil rate, the next £425,000 is taxed at 5 percent and the remainder £575,000 is taxed at 10 percent. The government goal with this tax holiday is to stimulate the housing market and therefore the economic activity. Figure 3.1 displays that by June 2021, the SDLT receipts reached historical high levels.

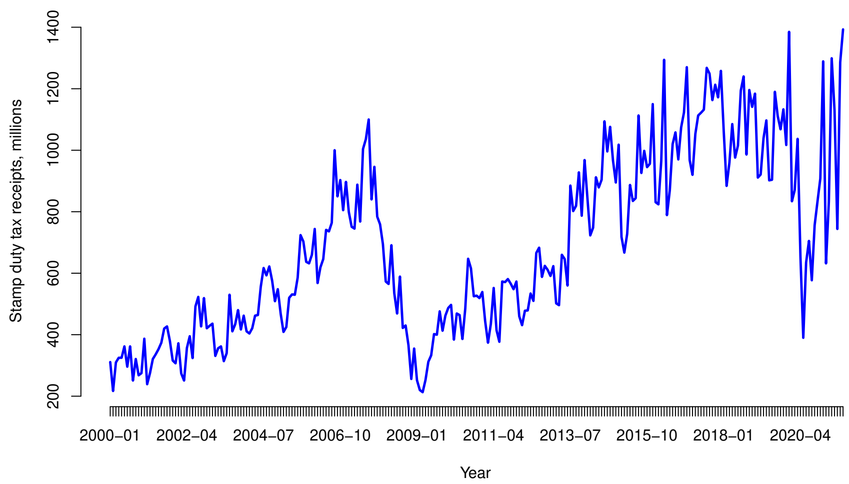

Figure 3.1: SDLT receipts in millions.

In this paper I evaluate the impact of the SDLT holiday by investigating the behaviour of the UK housing market during the period of the SDLT holiday. As the SDLT constitutes a natural experiment, it was unexpected and came in force with an immediate effect. To answer this question, I develop a counterfactual scenario based on the synthetic control method (Abadie and Gardeazabal 2003; Abadie et al. 2010; Abadie et al. 2015). Essentially, I will compare the developments of UK house prices to a donor pool that include OECD countries and regions that did not impose such tax holidays. The main assumption of the synthetic control method is that the synthetic control unit would have continued to behave in the same way as the treated unit in the absence of the SDLT holiday. The synthetic control methodology enables the weights assigned to the control group to be determined using a data driven approach, so that the behavior of the synthetic control unit best resembles the UK economy and as closely as possible before the

SDLT holiday.

I find that the synthetic control units track the behaviour of the UK and its large regions house price index prior to the tax holiday, providing a good benchmark for evaluating the impact of the SDLT holiday. My results show that in the two quarters following the SDLT holiday, the aggregate UK House Price Index (HPI) increased by annual rates of 2.11% and 2.96% relative to its synthetic counterpart. This increase reflects the boost the SDLT holiday caused the UK housing market. The aggregate analysis should be complemented with regional level analysis to take into consideration the regional heterogeneity. Ashworth and Parker (1997) point out that it would be misleading to assume that the housing markets in all the UK regions operate in the same way. Carlton (1981) is one of the earliest papers to address the spatial heterogeneity of property taxes, his theoretical model shows that property taxes disproportionately affect values of housing close to the central business districts more than those which are far.

The regional differences of the housing market indicate that the response to the SDLT holiday will be heterogeneous across regions. The response of different regions to the SDLT holiday can be influenced by a combination of regional housing market characteristics, affordability considerations, and broader economic factors such as local economic performance, regional policies and demographics and housing supply-demand dynamics. Although the SDLT is a tax on mobility that discourages residential moves and labor markets (Hilber and Lyytikäinen 2017b). The working from home mandate that was introduced at the start of the COVID-19 pandemic caused individuals to value bigger houses with spaces to be made into offices. A survey by the UK's leading property portal Zoopla found that lockdown has changed home hunters' priorities with 46% of respondents having requirements that must include more outside space. Moreover, Bank of England suggests that working from home caused houses outside city centres to become more valuable (Bloomberg 2021). When evaluating the impact on a regional level, the results of this paper confirm the hypothesis that the pandemic and working from home caused residents to move out from big cities in favour of the countryside and coastal cities. I answer the question of "What if the SDLT holiday never happened" by looking at the agregate UK house prices and the 12 Nomenclature of Territorial Units for Statistics 1 regions (henceforth NUTS 1). These 12 regions are made of the three nations, Scotland, Wales and Northern Ireland and the nine English regions. I show that all the 12 NUTS 1 regions of the UK, with exception of Greater London, had an increase in their house price indices with Wales, North West of England and North East of England gaining the most in the two quarters of the tax holiday. Although Greater London exhibited an increase of 1.03 percent in its house prices in 2020 Q4, there was a decrease in the house prices by 1.48 percent in 2021 Q1. The effect of the SDLT holiday is measured by the difference between the treated unit and its synthetic counterpart. Sales data on Zoopla show that buyers have moved from London and big cities in England to countryside locations such as Ryedale in Yorkshire, Herefordshire in West Midlands, Sevenoaks in Kent and the Cotswolds (Zoopla 2020b). These results confirm the findings of a growing literature on the impact of COVID-19 work displacement and the housing market in the U.S, as a result of the work from home mandate, the housing demand reallocates away from major city centres towards lower density areas, pushing the house prices in the central business districts to decrease (Ramani and Bloom 2021; Sitian Liu and Su 2021; Rosenthal et al. 2022).

The results of this paper also confirm the economic theory that agents respond to property taxes by "bunching", increasing their property transactions just below or around the tax thresholds. Best and Kleven (2018) show that temporary tax stimulus by reducing the tax rate causes both timing and extensive margin effects. The timing effect indicates that the agents are induced to shift their purchases forward for the period of the tax cuts, while the extensive margin effect indicates that the agents that were indifferent between buying in the period of the tax cuts and never buying are induced to buy in the period of the tax cuts. This has a positive impact on the house prices as it results in a sudden increase in the property demand.

This paper contributes to a growing body of research by policymakers and academics about SDLT and the impact that the SDLT holiday had on the UK economy. The nature of this tax is that it must be paid in full by the purchaser of the property at the time of the purchase and can not be funded by the mortgage, which increases the amount needed for the down payment of the property purchase (Seely and Keep 2017). Oates (1969) was one of the earliest literature to show that higher property tax rates depresses the housing markets as property taxes and public spending are capitalised into the house prices. SDLT has been debated in the housing economics literature as it is viewed as being an inefficient form of tax as it affects employment, productivity and distorts decisions about moving for existing and first time buyers (Van Ommeren and Van Leuvensteijn 2005). Reducing housing transactions due to housing market frictions can affect the economy more widely as housing purchases are associated with a range of economic activities including spending on home improvements, maintenance and purchase of durable goods.[1] The Office for Budget Responsibility (OBR), calculated semi-elasticities for property prices to changes in SDLT, showing that if the SDLT was cut by 1%, properties worth between £250,000 and £1 million will have an increase in their price by 1.5%. Best and Kleven (2018) using the UK consumption survey data, estimate that a house transaction triggers extra spending of about 5% on house prices. Scanlon et al. (2021) using a survey from Family Building Society highlight that the money saved from the SDLT holiday was spent on moving expenses and expenditures on furnishing, decorating and appliances, which could have stimulated the economy by an average of £2.2 billion. Family Building Society surveys show that SDLT holiday caused a 50 percent increase in transactions, leading to better housing stock. Arnold (2020) stipulates that if the SDLT holiday was made permanent, house prices would be an average 1.3% higher than they otherwise would have been. This is because as a result of the removal of SDLT, the demand for properties will increase causing house prices to increase by some proportion of the tax eliminated. In comparison to these papers, I evaluate the impact of SDLT holiday on the aggregate and regional houses prices by performing a counterfactual analysis, identifying heterogeneity in how the housing market responded. The only paper that used a counterfactual analysis in the context of SDLT was Besley et al. (2014), investigating the impact of SDLT holiday of 2008-2009 using Difference in Difference methodology with the control group as the transactions in the next taxable bracket. They show that the tax holiday caused an increase in the number of transactions by 8% and a decrease to the after tax house purchase prices by 0.7%. In contrast, I evaluate the impact of SDLT cuts, using a control

[1] House Builders Federation report in 2017 found that over a third of homeowners surveys spend between £10,000 - £40,000 on upgrading their recently purchased homes.

group of OECD countries that did not have such tax cuts, by asking if the house prices in the UK would have increased in the absence of this tax holiday.

This paper also adds to a growing literature that studied the use of fiscal measures to stimulate the economic activity out of the lockdown freezes. US COVID stimulus on economic activity (Bayer et al. 2020; Kaplan et al. 2020). They show that the stimulus stabilised the economy in the early stages of the recession. Fetzer et al. (2021), looks at the impact of the UK hospitality fiscal stimulus "Eat Out to Help Out" scheme on the COVID infection rates, showing that the scheme caused an increase in COVID infections by up to 17%. In comparison to these papers, this paper is the first to look at the impact of the fiscal stimulus on the UK housing market on both aggregate and regional level.

The remainder of this paper is organised as follows. Section 3.2 discusses the data used in the analysis. Section 3.3 describes the synthetic control methodology. Section 3.4 presents the results on the aggregate UK and across all the UK regions. Section 3.5 presents robustness checks. Section 3.6 presents the discussion of the findings. Section 3.7 concludes.

3.2 Data

In order to construct the synthetic control estimates for the house price index for the UK and across its regions, I rely on the OECD Analytical House Price Indicators dataset. The control group includes OECD countries and large regions. The data is available up to 2021 Q1. Stamp duties on property transactions exist in several countries other than the UK such as Australia, Japan and France. Japan and France had some reliefs in their property tax systems around the same period as the UK. Therefore, these two countries are removed from the control group.[2] In addition to this, since all the countries in the donor pool imposed lockdown restrictions and eased the restrictions around the same time except for Sweden which did not impose any restrictions, I removed Sweden from the control group.[3]

SDLT generates important source for tax receipts for the UK treasury. Residential

[2] Japan introduced deduction measures in property taxes in October 2019 to stimulate its housing market, due to COVID-19, they were extended until 2022. On the other hand, France introduced property tax reliefs for businesses affected by COVID-19.

[3] My results are unchanged when including Sweden in the control group.

SDLT, not including the LBTT of Scotland, accounts for about 1.3% of HMRC receipts in 2015/2016 (Scanlon et al. 2017). Historically, the UK introduced a SDLT holiday by raising the threshold of SDLT from £125,000 to £175,000 in September 2008 until December 2009 to stimulate the housing market during the great recession. The SDLT had a major reform in its operation in 2014, changing from a slab tax, also known as notching structure, where a marginal tax rate is applied to the entire value of the transaction to a slice structure, applying different tax rates to portions of the property price that exceeds £125,000. In 2016 an extra 3% was imposed for residential purchases made by existing homeowners, affecting buy-to-let landlords and buyers of second homes. SDLT generates an important sources for tax receipts for the UK treasury.

Figure 3.2 displays the number of residential property transactions. It shows that as the result of the national lockdown restrictions that were introduced in late March, the number of residential property transactions dropped by more than 56% in April 2020. As a result of the relaxation of the rules, the number of transactions started to increase from June 2020 onwards. According to Besley et al. (2014) the time lag between the date agreeing on the purchase price and the date completing the housing contract is 60 days on average. The SDLT holiday was announced on 8$^{\text{th}}$ of July, this suggests that a large fraction of the contracts taking place in the first few months of the holiday period may have been agreed before the policy change announcement. The positive impact of the stimulus is short-lived and started to become visible from September onwards. SDLT is an efficient tax that causes a mismatch in the housing market between the dwellings and the households that could result in large households living in small houses and small households living in large houses. This leads to waste in the form of misallocation costs (Hilber and Lyytikäinen 2017b). Thus the elimination or the reduction in the SDLT will increase the number of housing transactions. The number of property transactions had two peaks, the first one was in March 2021 as the scheme was scheduled to end on March 2021, while the second peak was in June 2021, when the scheme was due to end after the extension in its deadline, as people were rushing to take advantage of the stimulus before its expiry. Figures 3.1 and 3.2 imply that the housing stimulus resulted in a large boost for the housing market activities during the period of the SDLT holiday. This illustrates the Laffer curve effect as the SDLT holiday encouraged homeowners that are within and outside the exempt threshold to purchase properties, causing an increase in the number of property transactions and the SDLT revenues. In addition to this, figure 3.3 displays the Halifax House Price Index, showing a decrease in the house price index in the three months March to June 2020, where the average house price is £238,295 before it starts to pick up and increase by July 2020. The average house price in the first quarter of 2021 is £251,844. A report by NAEA Propertymark, the membership body for property agents, shows that in 2021 32% of the properties are sold for more than the original asking price.[4] In addition to this, the magnitude of the SDLT holiday on the economy goes beyond the housing market and has a bigger impact on the macroeconomy as a whole. Best and Kleven (2018) show that households spend an additional 1.98% of the house value on estate agents fees and a further 1.24% on other commissions.

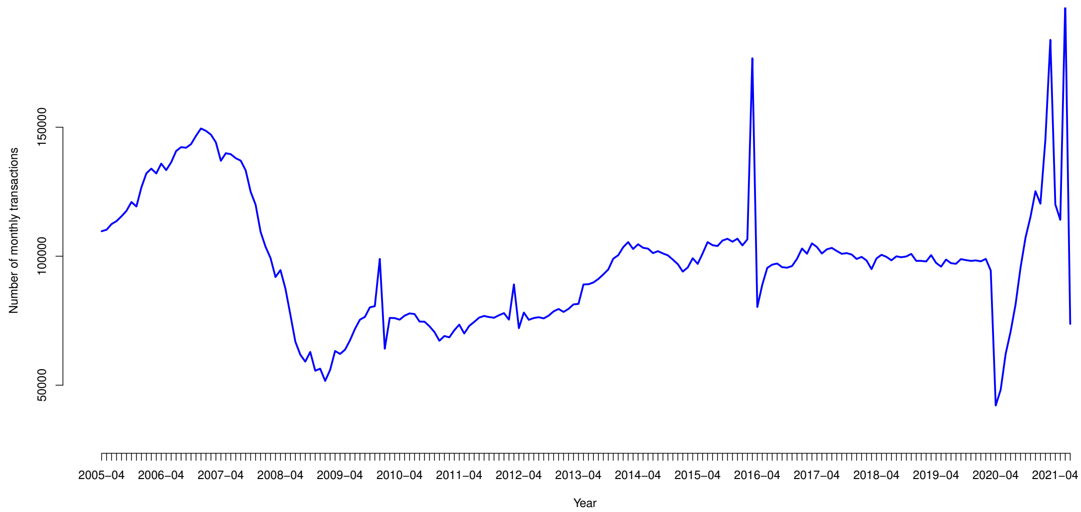

Figure 3.2: Number of residential transactions in the UK.

Monthly data from April 2005 to July 2021. The series is seasonally adjusted. Source: HM Revenue & Customs.

In figure 3.4, I plot the density of property transactions at different prices, ranging from £50,000 to £1.5 million in bins of £5,000. The top panel shows the transactions for England, while the bottom panel shows the transactions for Wales for the period of the SDLT holiday. The vertical red lines are at £500,000 and £250,000 to indicate the cut off amount that is exempt from the SDLT for each region. The figures display that

[4] https://www.propertymark.co.uk/resource/housing-report-apr-2021.html

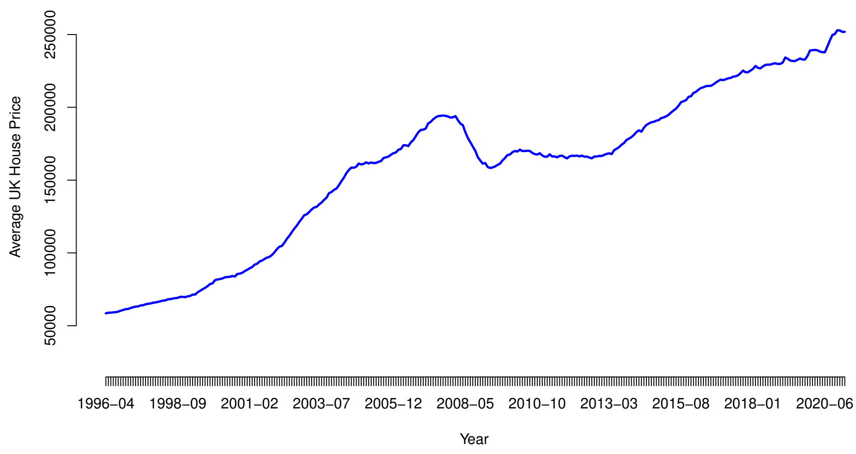

Figure 3.3: Monthly Halifax House Price Index.

Source: Data Stream.

during the period of the SDLT holiday, higher density of the properties are sold for an amount that is included in the exempt tax bracket. As illustrated, the histograms are skewed to the right, this shows that at higher prices beyond the exempt amount, the number of residential transactions made during this period decreases. The presence of the SDLT causes a distortion in the price distribution, resulting in a significant number of property transactions that are happening just below the £500,000 for England and £250,000 for Wales, illustrating the "bunching" phenomena of Best and Kleven (2018). The SDLT holiday means that home buyers could save up to £15,000 on their SDLT tax bill. According to Kopczuk and Munroe (2015), transactions occur when the surplus between what the buyer is willing to pay and the price negotiated and the lump-sum property tax is non-negative. That is, when the surplus is reduced, the uniform lump-sum discourages some sales. Hence, the agents respond to the SDLT by increasing the transaction activities just below the tax thresholds. In figure A.4 in the appendix, I show the density of property transactions for England and Wales in 2019. I show that the number of property transactions executed in 2019 in England beyond the cut off of £500,000, is greater than the number of transactions executed in 2020.

Table 3.1 illustrates descriptive statistics for the housing transactions that took place during the period of the SDLT holiday for England, Wales and Scotland. I remove transactions with values less than £50,000 and more than £1,500,000 to avoid the outliers effect. The table shows that among all the regions, South East of England had the highest number of transactions during that period with almost 123,000 housing transaction made, with an average value of transactions being £405,094, falling within the SDLT holiday bracket. However, North East had the lowest number of transactions. The table also shows that the mean of the housing transactions taking place in Greater London is at almost £640,000, while the median is £499,995 which fall outside the exempt amount of the SDLT holiday. This is evident from the fact that 50% of the housing transactions taking place are above the SDLT holiday threshold of £500,000.

			Percentile			
Region	Volume	Mean value	25th	50th	75th	95th
Scotland	112,205	193,634	154,922	176,486	209,538	284,710
Wales	25,718	209,662	125,000	175,000	254,950	440,000
South East	122,999	405,094	260,000	355,000	505,000	965,000
South West	77,294	354,752	214,000	290,000	410,000	760,000
Greater London	94,896	636,956	375,000	499,995	705,000	1,500,000
East of England	81,148	354,986	220,000	300,000	415,000	725,000
West Midlands	74,655	268,112	155,000	215,000	315,000	595,000
East Midlands	60,173	255,553	155,000	213,750	300,000	540,000
Yorkshire	74,570	231,445	125,000	180,000	275,000	530,000
North West	102,313	235,425	125,000	181,995	277,500	540,000
North East	32,850	185,713	95,000	145,000	222,000	430,000

Table 3.1: Descriptive statistics on the housing transactions.

Housing transactions taking place during the SDLT holiday period. To avoid outliers effect, I removed transactions with values less than £50,000 and more than £5,000,000. Source: HM Land Registry and Registers of Scotland.

3.3 Methodology: Synthetic Control Method

The choice to use the synthetic control method to investigate the impacts of the SDLT holiday is motivated from the approach that the method builds on an improvement from the traditional panel data models such as Difference in Difference (DiD) or Dynamic Panel Data method. This method was once referred to as "the most important innovation in the policy evaluation literature in the last 15 years" (Athey and Imbens 2017). The advantages of the synthetic control method stem from the ability to model the treatment effect when there is as little as one treated unit, something that can sometimes be unfeasible in other panel data methods. In addition to this, synthetic control method is superior with respect that instead of using one single control unit or simple average

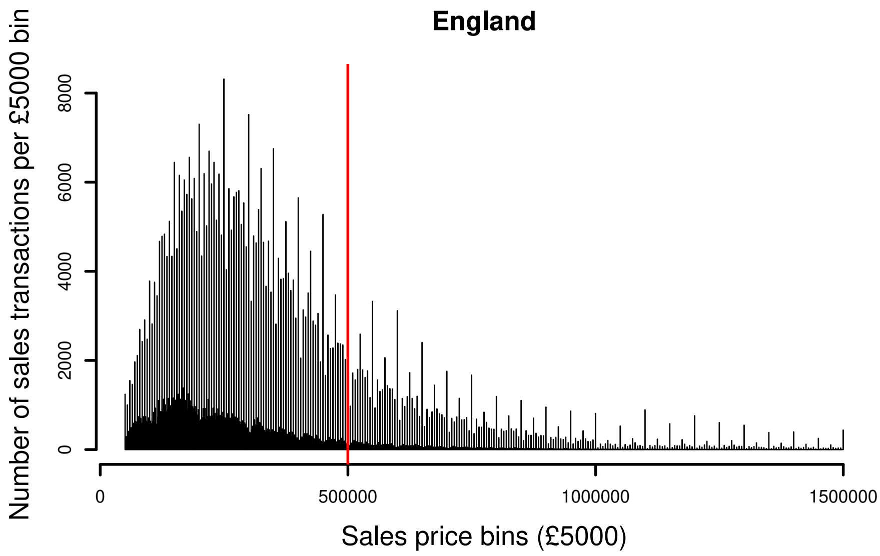

(a) England

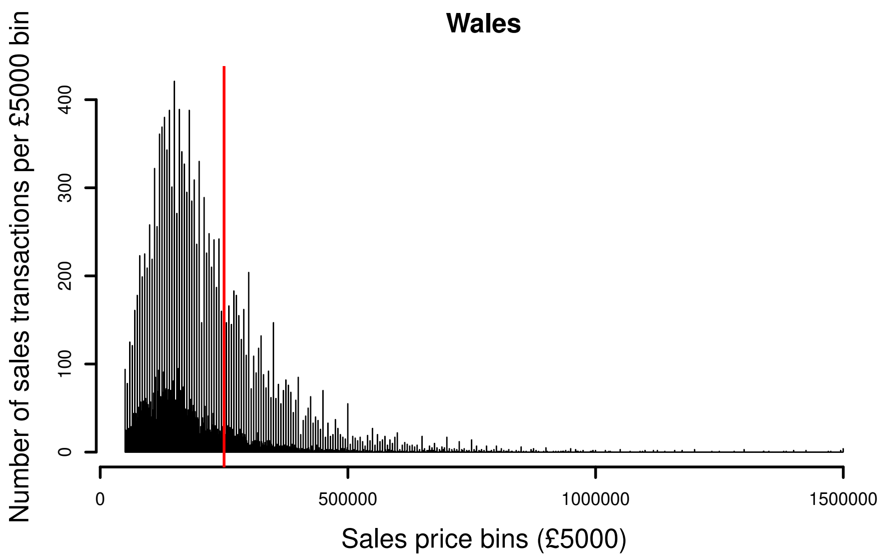

(b) Wales

Figure 3.4: Density of property transactions for England and Wales.

The period of study is July 2020 to June 2021. The vertical red line indicates the cut off amount for SDLT exemption, £500,000 for England and £250,000 for Wales. Source: HM Land Registry.

of control units as in the DiD, the method uses a weighted average of a set of potential control units to provide a synthetic control unit that closely resembles the affected unit in terms of predictors. That is to say, instead of using just one OECD country or a simple average of all the OECD countries in the donor pool, the weighted average control group has better match than any single country would be. This allows for the estimation of the counterfactual outcome that would have occurred for the treated unit in the absence of the treatment. By comparing the treated unit to its counterfactual, policymakers can understand the effectiveness of the treatment on the treated unit. Lastly, the method opens the door to performing several placebo tests in order to assess the validity of the estimated treatment effect.

The synthetic control method uses a weighted average of a set of potential control units to provide synthetic control units that closely resemble the affected unit in terms of predictors. In the context of the SDLT scenario, suppose that there are $J+1$ units (here: country and region pairs) in periods $t = 1, 2, \ldots, T$ (here: quarters), where unit $j = 1$ is the treated unit (here:UK) and units $j = 2, \ldots, J+1$ are the untreated units used in the control group. Unit $j = 1$ is exposed to the intervention (here: outcome of the SDLT holiday) at periods $T_{0+1}, \ldots, T$, (2020 Q4 and 2021 Q1). while being unaffected during the periods $1, \ldots, T_0$ (2013 Q1 to 2020 Q3), the in-sample size is 31 observations. In my analysis, I use 2020 Q4 as the period of intervention. Y_{it}^N is the outcome of unit i in the absence of intervention for units $i = 1, \ldots, J+1$. Y_{it}^I is the outcome of unit i at time t if the unit is exposed to intervention in periods $T_0+1, \ldots, T$. The intervention has no impact on the pre-intervention periods. $Y_{it}^N = Y_{it}^I$ for $t = 1, \ldots, T_0$. For the treated unit, Y_{it}^N can not be observed in the post-intervention period $t = T_0 + 1 \ldots, T$. However, it can best be modelled by a weighted average combination of untreated units in the donor pool. The synthetic control can be represented by a (Jx1) vector of weights $W = (w_2 \ldots, w_{j+1})'$, the synthetic control unit can be expressed as:

$$\hat{Y}_{1t}^N = \sum_{j=2}^{J+1} w_j^* Y_{jt}, \tag{3.3.1}$$

subject to $0 \leq w_j \leq 1$ for $j = 2, \ldots, J+1$ and $w_2 + \ldots + w_{j+1} = 1$. The constructed synthetic control unit is $\hat{Y}_{1t}^N$ is a weighted average of the untreated units in the control group. An important assumption of the synthetic control method is that the housing market in the control group are not directly affected by the results of the SDLT holiday, nor any contemporaneous events and the housing market for the treated unit would have developed in the same way as the synthetic control had it not been for the SDLT holiday.[5] The assumption of conditional independence helps ensure that the control units

[5] In my robustness check of section 5, I show that the UK had the greatest impact of the SDLT in comparison to the

selected for the synthetic control group have similar pre-treatment trends, making them plausible substitutes for the treated unit in estimating the counterfactual outcome. That is, the assumption suggests that the treatment is unrelated to the historical outcomes and covariates, except through their common dependence on observed pre-treatment characteristics. By matching the pre-treatment characteristics of the treated unit with appropriate control units, the synthetic control method aims to create a weighted average of the control units that closely approximates the treated unit's outcomes in the absence of the treatment. When this assumption holds, it implies that the treated unit and the control units used to construct the synthetic control group should have similar trends in past outcomes and covariates. This similarity is crucial because it allows the synthetic control unit to serve as a meaningful representation of the unobserved counterfactual scenario for the treated unit (Gilchrist et al. 2023).

The choice of W is such that the characteristics of the treated unit are best resembled by the characteristics of the synthetic control. X_1 is a $(k \times 1)$ vector of pre-intervention characteristics of the treated unit. X_0 is $(k \times j)$ matrix of the same characteristics of the donor pool. There is no one way that is always followed in the literature for the choice of covariates. Abadie et al. (2010) advises on using economic predictors in addition to lags of the treated variables in the pre-intervention period, while Abadie et al. (2015) advises on using economic predictors in addition to the average of the treated unit in the pre-intervention period. However, Ferman et al. (2020) finds that specification searching is reduced when the number of pre-intervention period goes to infinity and all pre-intrevention lags of the variable of interest of the treated unit are used as predictors.[6] The vector $\mathbf{W}^*$ is chosen to minimise this difference between the pre-intervention characteristics of the treated unit and the synthetic control unit $||X_1 - X_0 W||$ subject to the constraints of w_j mentioned above. Abadie et al. (2010) measure the discrepancy between the characteristics of the treated and the synthetic control unit by introducing vector $\mathbf{V}$, which is a symmetric and positive semidefinite $(k \times k)$ matrix. An optimal choice of $\mathbf{V}$ assigns weights to the linear combinations of the variables in $\mathbf{X}_0$ and $\mathbf{X}_1$ to minimize the mean square error of the synthetic control estimator. Typically, $\mathbf{V}$ is

countries in the control group.

[6] I follow the approach of Ferman et al. (2020) approach of using the synthetic control specification that yields the least Mean Square Predictive Error.

selected to weight the predictors in accordance to their predictive power on the outcome, this makes the minimization problem as follows:

$$\|\mathbf{X}_1 - \mathbf{X}_0\|_{\mathbf{v}} = \sqrt{(\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})'\mathbf{V}(\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})}. \tag{3.3.2}$$

As suggested by Abadie and Gardeazabal (2003) and Abadie et al. (2010), I set the V to be determined using a data driven approach such that the MSPE of the outcome variable is minimised for the pre-treatment period. Abadie et al. (2010) show that if the weighted average of the synthetic control characteristics X_0 can match the treated unit characteristics X_1, it provides a valid counterfactual for Y_{jt}^N in the sense that $\hat{Y}_{1t}^N - Y_{jt}^N$ is close to 0 in the pre-intervention period, $t = 1, ..., T_0$. In the case of evaluating the impact of the SDLT holiday on the British housing market, I use the variables employment rate, inflation CPI, house price to income ratio, interest rate and the dependent variable HPI year on year (YoY) growth rate. This is based on the fact that these are macro and financial factors that could influence the housing market. Including the pre-SDLT holiday values of the dependent variable helps to limit the sensitivity of the weighting procedure to any unobserved characteristics which could result in differences between the actual UK and its synthetic counterpart and ensures a better fitting pre-intervention model. These observed covariates are not affected by the treatment in the pre-treatment period. The synthetic control is primarily designed to use any covariates that help explain the outcome variable as predictors, and not only pre-treatment values of the outcome variable (Gardeazabal and Vega-Bayo 2017). I evaluate the relationship between the dependent variable and the vector of covariates by running the regression as in (Shimeng Liu 2015; Breinlich et al. 2020):

$$\begin{aligned} Y_{it} = \beta_0 + \beta_1 EmploymentRate_{it} + \beta_2 CPI_{it} + \beta_3 PIRatio_{it} + \\ \beta_4 InterestRate_{it} + u_{it},\ i = 1, \ldots,\ t = 1, \ldots, T. \end{aligned}$$

As suggested by Botosaru and Ferman (2019), the inclusion of covariates as predictors when the outcome equation is nonlinear in covariates should not necessarily be a problem if one uses the data-driven to determine the matrix V. As long as the pre-treatment outcome lags are also included as predictors, then the importance given to such covariates would be relatively small and the estimation procedure would not attempt to match on covariates that should not be matched on. As the output of the regression shows in table 3.2, there exists linear and significant relationship between all of the variables and the HPI YoY growth rate. Therefore,all of the variables with significant coefficients can be used as characteristics in the data driven approach that estimates the vector of weights assigned to each country in the donor pool. The variables with significant coefficients along with average of the pre-treatment lags of the dependent variables will be used as covariates in the data driven approach that estimates the vector of weights assigned to each country in the donor pool as in Abadie et al. (2015). As a robustness check, I also consider alternative specifications such as using the pre-treatment period lags only as economic covariates (Ferman et al. 2020; Breinlich et al. 2020; Botosaru and Ferman 2019) and using economic covariates along with all the lags of the dependent variable in the pre-treatment period (Billmeier and Nannicini 2013; Bilgel and Galle 2015; Stearns 2015; Villar and Papyrakis 2017; Pellegrini et al. 2021; Papyrakis et al. 2023).

For the post-intervention period, where $t \geq T_0$, the treatment effect is the difference between the realized outcome and the synthetic control outcome:

$$Y_{1t} - \sum_{j=2}^{J+1} w_j^* Y_{jt} = Y_{1t} - \hat{Y}_{1t}^N, \forall t = T_{0+1}, \ldots, T \tag{3.3.3}$$

Throughout this paper, the main identifying assumption is that the UK aggregate and regional house prices would have developed as the synthetic unit had it not been for the SDLT holiday. I estimate the following regression equation:

$$Y_{1t} - \hat{Y}_{1t}^N = \mu + \lambda d_t + \epsilon_t, \tag{3.3.4}$$

$$d_t = \begin{cases} 0, & \text{for } t < T_0 \\ 1, & \text{for } t \geq T_0 \end{cases}$$

The step dummy variable, d_t, models the average impact in the post intervention period, where the coefficient λ will be referred to as the "synthetic control regression estimator". When subtracting the synthetic control unit from the treated unit, this estimates the

	Dependent variable:	
	HPI YoY Growth Rate	
	(1)	(2)
Employment	-0.001^{***}	0.0002
	(0.0003)	(0.002)
CPI	0.0001	-0.003^{***}
	(0.001)	(0.001)
PI	0.003^{***}	0.003^{***}
	(0.0002)	(0.0003)
Interest Rate	0.0002	-0.032^{***}
	(0.002)	(0.004)
Constant	-0.190^{***}	
	(0.059)	
Observations	425	425
R^2	0.341	0.206
Adjusted R^2	0.335	0.114
F Statistic	54.342^{***} (df = 4; 420)	24.682^{***} (df = 4; 380)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 3.2: Regression estimates for the HPI and economic covariates

Estimation of the relationship between the economic covariates and the dependent variables. The dependent variable is the HPI YoY Growth Rate. Column (1) is the estimation with Pooled OLS. Column (2) is the estimation with country and time fixed effects (two ways FE).

synthetic control effect.

As a robustness check, I will compare the estimates of equation 3.3.4 to a DiD regression model. The synthetic control method is more appropriate method to be used in comparative case study and in the context of SDLT holiday evaluation in particular than the DiD as there is one treated unit, where the synthetic control relies on the assumption that the synthetic control captures the unobserved factors affecting the treated unit, enabling causal inference. On the other hand, DiD is commonly employed in pooled cross-sectional analyses involving multiple treated and control units. The difference between the synthetic control method and the DiD method is that the synthetic control method creates a weighted average of the countries in the donor pool that matches the characteristics and the trend of the treated unit. The synthetic control method shows a plot of how the treated unit would have been looking in the absent of the treatment. The synthetic control method allows an algorithm to determine which combination of countries match the evolution of UK HPI before the SDLT holiday with highest possible accuracy. On the other hand, DiD is based on the conditional expectation assumption i.e. the expected outcome of the treatment group (or the treated units) given a specific set of covariates or observed characteristics. It represents the average treatment effect for the treated group, taking into account the control group's outcomes and the changes over time. Although the synthetic control method compares the change in outcomes over time between a treated unit and a control unit or group of units that are assumed to be similar in all relevant respects except for the treatment or policy intervention. The DiD does not take in consideration matching the treated group and their characteristics, instead it is based on the common trend assumption. That is, the treatment and control group follow a similar pattern in the pre-treatment period, they would also follow a similar pattern in the post-treatment period in the absence of the treatment. The DiD model can be written as:

$$Y_{it} = \beta_0 + \beta_1 treated + \beta_2 treatment + \delta(treated \times treatment) + {Z_{it}}'\theta + \epsilon_{it}.$$

For the context of the SDLT holiday, the DiD model will be written as:

$$Y_{it} = \beta_0 + \beta_1 UK + \beta_2 d_t + \delta(UK \times d_t) + Z_{it}{}'\theta + \epsilon_{it}, \tag{3.3.5}$$

Where Z_{it} is the set of economics predictors which may be included in the regression. The coefficient β_1 illustrates the effect of the UK on the dependent variable, the coefficient β_2 shows the effect of the post-2020Q2 on the dependent variable. Lastly, the coefficient δ is the DiD estimator which represents the average treatment effect (ATE) of SDLT holiday on the treated unit. This specification estimated with and without two ways Fixed Effects (FE).

The difference between equations 3.3.4 and 3.3.5 can be explained as the former looks at estimating the effect of the SDLT holiday on the gap between the treated unit and its synthetic counterpart, where the synthetic counterpart is assumed to be the same as the treated UK in the absence of the SDLT holiday. The coefficient λ captures the change in this gap over time, indicating the evolving effect of SDLT on the treated unit relative to the synthetic control unit. However, in the latter, the coefficient δ is the DiD estimator which captures the difference in the average outcomes between the UK and the countries in the donor pool after accounting for time and group-specific effects. Having said that, the coefficients λ and δ serve different purposes and capture different aspects of the treatment effect.

To perform inference on the results of the synthetic control method in the scenario of SDLT holiday, the following steps are undertaken:

1. Construction of Confidence Intervals: Confidence intervals are constructed for each trajectory of the treated unit in comparison to its synthetic counterpart. The confidence interval is calculated based on 2 standard deviations of the pre-SDLT gap between the UK and the synthetic UK in order to have a 95% confidence interval. This approach allows for assessing the significance of the differences between the UK and its synthetic counterpart. This inference approach is based on the work of Born et al. (2019).

2. In-Time Placebo Test: A placebo test is conducted by assigning the SDLT holiday

to a false year. The purpose of this test is to evaluate whether there would be any significant divergence between the UK and its synthetic counterpart after this false year. If the false assignment of the SDLT holiday leads to a notable and significant divergence, it raises doubts about the main findings.

3. Country Placebo Test: Another placebo test is performed by assigning the SDLT holiday to each of the countries in the control group. The gap between the treated unit (UK) and its synthetic counterpart is calculated based on this assignment. The objective of this test is to assess whether the observed gap between the UK and its synthetic counterpart is driven by the SDLT holiday or by other random fluctuations within the control group.

4. Calculations of Fisher p-value based on country placebo effects: Abadie et al. (2010) propose an approach to assign the treatment to each country in the control donor pool and estimating the "country placebo effect". For each control unit, the ratio of post-treatment MSE to the pre-treatment MSE is computed as the following:

$$R_j = \frac{MSE_{j,post-SDLT-Holiday}}{MSE_{j,pre-SDLT-Holiday}} = \frac{\left(\sum_{t=T_{0+1}}^{T}\left(Y_{jt} - \hat{Y}_{jt}^{N}\right)^2\right) / \left(T - T_0\right)}{\left(\sum_{t=1}^{T_0}\left(Y_{jt} - \hat{Y}_{jt}^{N}\right)^2\right) / T_0},$$

Where $\left(Y_{jt} - \hat{Y}_{jt}^{N}\right)$ is the treatment effect. Having large value of the MSE in the post-SDLT holiday period indicates the tax holiday effect on the treated unit. To validate that the gap in the post-SDLT holiday period is the tax holiday effect for the UK, the p-value is then calculated by comparing the value of R_1, the MSE ratio for the UK, to that of $R_j\, \forall j = 2, \ldots, J+1$ control units:

$$p_1 = \frac{\sum_{j=1}^{J+1} I_{+}\left(R_j - R_1\right)}{J+1}, \tag{3.3.6}$$

Where $I_{+}(.)$ is an indicator function that takes a value of one for nonnegative arguments and zero otherwise. That is to say, the indicator function can be written as $1\left(R_j \geq R_1\right)$ Comparing the value of this statistic for the UK, R_1, to that of all other units, the p-value measures the chances of obtaining a ratio as high as R_1 if one is to pick a country randomly from the sample of countries in the donor pool. If the p-value is less than some pre-specified significance level, such as the the 0.1 traditional value then the null hypothesis that the treatment has no effect can be rejected implying that there is some regions with a non-zero effect for some period.

These steps help provide additional evidence and address potential concerns regarding the reliability of the main findings derived from the synthetic control method in the context of evaluating the impact of the SDLT holiday.

3.4 Empirical results and discussion

Now I present the main results of this paper which estimates the impact of the SDLT holiday on the aggregate house price index of the UK and its regions.

3.4.1 Aggregate UK House Price Index

The identification strategy is that the UK and its synthetic counterpart have the similar performance and trend. The difference between them following 2020 Q3 is solely attributed to the SDLT holiday. This assumption is reasonable as none of the countries in the donor pool had any similar land tax relief or holiday on their properties. Also, another important assumption is that the treatment only affects the UK house prices with no such spillover effect on the countries in the control group.

With these assumptions in mind, I use the data driven approach that solves the minimization problem of equation 3.3.2 to construct synthetic UK house price index year on year growth rate. The synthetic UK is constructed with a weighted average of countries in the donor pool. The weights assigned to each country in the donor pool is based on matching the vector of pre-intervention characteristics of the treated unit $\mathbf{X}_0$ to that of the control group $\mathbf{X}_1$. These characteristics are employment rate, consumer price index (CPI), house price to income ratio and the dependent variable. These variables are averaged for the pre-intervention period. Table 3.3 shows a comparison between the pre-intervention characteristics of the actual UK, synthetic UK and a population weighted average of all the countries in the donor pool. As illustrated, the synthetic UK is closely matched to the actual UK than the weighted average of all the countries in the donor pool.

Economic predictor	UK	Synthetic UK	Sample Mean
Employment rate	73.95%	68.22%	67.63%
Consumer Price Index (Inflation)	102.14	102.09	101.54
House price to income ratio	102.91	100.82	103.45
Interest Rate	0.586	0.581	0.444
HPI annual growth rate (Dependent Variable)	3.9%	3.9%	3.3%

Table 3.3: Matching countries economic predictors.

Employment rate, Consumer Price Index and House price to income ratio and the HPI annual growth rate (dependent variable) are averaged for the pre-intervention period. Data source: OECD.

In table 3.4, I show the countries that are used to construct the synthetic control unit. The synthetic UK house price index year on year growth rate is constructed as a weighted average of Korea, Finland, Ireland, Norway, Canada and Denmark. This indicates that pre-SDLT holiday, the UK series had a similar evolution to these six countries. Although the weights assigned to the countries in the donor pool appear to be reasonable, the data driven approach used to generate these weights refrain from any structural interpretation for the country weights.

As in Cebula (2009) and Coombs et al. (2012) fiscal differentials are capitalised in the house prices, such that an increase in the property taxes reduce the house prices. That is to say, reduction in the SDLT should cause an increase in the house prices. Figure 3.5 displays that the synthetic UK closely mimics the performance of the UK series between the beginning of the sample and 2021 Q1. The shaded area corresponds to two standard deviation of the pre-treatment gap between the UK and its synthetic counterpart house price index year on year growth rate. As illustrated, both series had a jump in their house price index year on year growth rate following 2020 Q2, which indicates the increase in the prices is due to the rising demand following easing of the the lockdown restrictions. However, the UK overperformed its synthetic counterpart by 2.11% in 2020 Q4 and 2.96% in 2021 Q1 from the same period the year before, measured by the difference between the UK and the synthetic UK series in the during the SDLT holiday period. This increase can be attributed to the impact of the SDLT holiday. Similarly, table 3.6 shows the estimation of the SDLT holiday effect, where the coefficient reported in column (1) is λ estimated using equation 3.3.3. Columns(2) and (3) display the DiD estimators of equation 3.3.5, in the absence and presence of time fixed effects. As illustrated by the coefficient of d_t, the average effect of the SDLT holiday on the UK housing market, illustrated by the difference between the UK and its synthetic counterpart, is 1.214 percentage points. While the coefficient of d_t in columns (2) and (3) just show the effect of the post-SDLT period on the HPI on the entire panel data. As there is only one treated unit, UK, there is not much difference between columns (2) and (3) with the exclusion and the inclusion of the Fixed Effect (FE). In fact, the coefficients of the dummy variable d_t and UK for the FE model are absorbed in the estimation due to the presence of only one treatment unit.

Country	Country weights	Country	Country weights
Austria	0	Belgium	0
Canada	6.8%	Czech Republic	0
Denmark	3.2%	Finland	24.4%
France	0	Germany	0
Greece	0	Hungary	0
Ireland	21.2%	Italy	0
Korea	36.1%	Luxembourg	0
Netherlands	0	Norway	8.3%
Poland	0	Portugal	0
Spain	0	Switzerland	0
US	0		

Table 3.4: Country weights to construct synthetic UK house price index year on year growth rate.

The weights illustrate the importance of country j in the donor pool to approximate the UK.

I now shift the analysis to focus on the regional house price index, to investigate the impact the tax holiday had on the house prices for each of the UK NUTS 1 statistical regions.

3.4.2 Regional House Price Index

To investigate further the extent and the variation of the SDLT holiday, I extend the analysis to look into the UK regional level. The main underlying assumption that I will be relying on is that the UK regional house price index would have developed as the synthetic control units had it not been for the SDLT holiday. That is to say, I quantify the SDLT holiday impact as the difference between the realized house price index year on year growth rate values and their synthetic counterpart. Table 3.7 shows the top synthetic control weights assigned to each of the regions in the control group to construct

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^{N}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	1.214**	−0.008*	
	(0.456)	(0.004)	
UK		−0.022**	
		(0.010)	
$d_t \times UK$		0.016	0.015
		(0.019)	(0.019)
PI		0.003***	0.003***
		(0.0002)	(0.0003)
CPI		0.001	−0.003***
		(0.001)	(0.001)
Employment		−0.001***	0.0002
		(0.0003)	(0.002)
Interest Rate		−0.001	−0.002***
		(0.002)	(0.002)
Constant	0.039	−0.251***	
	(0.136)	(0.065)	
Observations	34	425	425
R^2	0.181	0.354	0.208
Adjusted R^2	0.155	0.343	0.114
Residual Std. Error	0.755 (df = 32)	0.033 (df = 417)	
F Statistic	7.069** (df = 1; 32)	32.597*** (df = 7; 417)	19.957*** (df = 5; 379)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 3.6: Effect of the SDLT holiday on the UK aggregate HPI YoY.

Estimating the effect of the SDLT holiday on the UK aggregate HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the UK and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns (2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

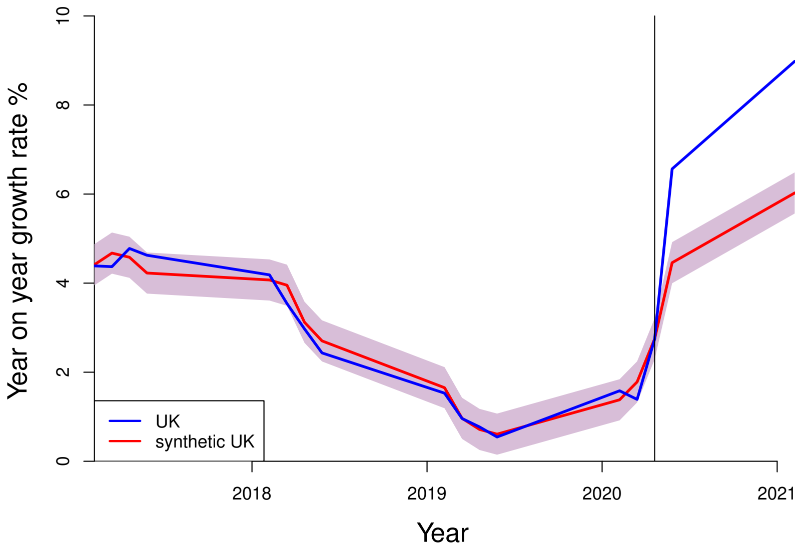

Figure 3.5: Impact of SDLT holiday on UK housing market.

UK against Synthetic UK house price index percentage change from the same time as the year before. The shaded area represents two standard deviations of the pre-treatment gap between the UK and its synthetic counterpart.

the synthetic control units for the three nations and the nine large English regions.

Figure 3.6 displays the year on year growth rate of the house price index over time across the three nations Wales, Scotland and Northern Ireland. The graphs show that the synthetic control unit for each of these nations consistently mimics the evolution of the actual series in the period prior to the SDLT holiday. All these three nations experienced divergence from their synthetic counterparts in 2020 Q3 onwards, where Wales exhibits the greatest increase in its house price index.

Figure 3.7 displays the synthetic control estimates for the remaining English regions. For all of the regions, the graph shows that they have a very close fit to their synthetic counterpart in the pre-treatment period. All the regions except Greater London exhibit a significant increase in their house price index following the SDLT holiday. Greater London had the smallest increase in its house price index in 2020 Q4, however, this increase only stayed for one quarter, before it decreases and underperforms its synthetic counterpart. Table 3.8 displays the synthetic control regression estimator λ of equation 3.3.3 in column (1) and the DiD estimator δ of equation 3.3.5 in columns (2) and (3), tables A.13 to A.34 in the appendix show the estimation of the regression for each region separately. The results of the synthetic control regression go along with the synthetic

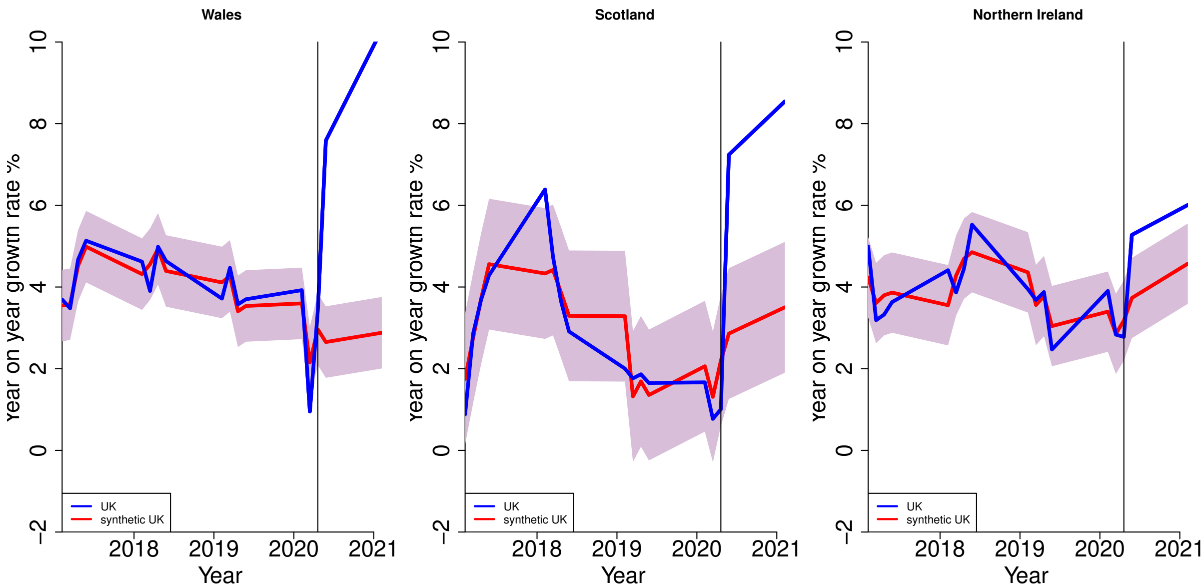

Figure 3.6: Impact of SDLT holiday on the three British nations.

Wales, Scotland and Northern Ireland against their Synthetic counterparts house price index (year on year growth rate). The shaded area is two standard deviations of the pre-intervention gap between the regions and their synthetic counterpart. Due to unavailability of economic covariates for the control group on regional level, I use the approach of Ferman et al. (2020) to match the treated and the control units based on all the pre-intervention lags of the HPI year on year growth rate instead.

control figures 3.6 and 3.7 that all the regions exhibited significant positive impact with the exception of East of England and Greater London. However, for the DiD regression, the DiD estimator which measures the impact of the treatment (SDLT holiday) on the treated unit (UK) is barely significant for any of the regions.

Region	Synthetic control country weights
Wales	Valencia =65.50%, South Australia Adelaide = 20.40%, Iceland = 3.50%, Madrid = 3.50%.
Scotland	Basque Country = 52.70% Western Australia Perth =28.50, Cantabaria = 13.20% and Tasmania = 5.30%.
Northern Ireland	Ceuta =34%, Aragon = 25.10%, Northern Territory Darwin Australia = 20.60%, Central and Western Lithuania =6.90%.
Yorkshire and the Humber	South Australia Adelaide = 50.50%, Valencia = 17.90%, Pomerania = 5.40%, Murcia =4.40%.
North East of England	Northern Territory Darwin Australia = 34.10%, Canary Islands = 24.90%, Basque Country = 22.40%.
North West of England	South Australia Adelaide = 54.70%, Basque Country = 25.70%, Ceuta = 13.50%.
East Midlands	South Australia Adelaide = 27.40%, Madrid =22.80%, Western Australia Perth = 17.30%, Pomerania = 8.20%
West Midlands	South Australia Adelaide = 38.40%, Madrid 20.30%, Western Australia Perth = 18.10%, Ceuta = 12.90%
South West of England	Western Australia Perth = 22%, South Australia Adelaide = 17.30%, Madrid = 15%, Northern Territory Darwin Australia = 13.20%.
South East of England	South Australia Adelaide = 59.20%, Northern Territory Darwin Australia = 12.50%, Madrid = 12.40%, Victoria, Melbourne = 9.50%.
East of England	South Australia Adelaide =83.70%, Madrid =8.50%,Victoria, Melbourne = 7.80%
Greater London	Northern Territory Darwin Australia = 35.40%, South Australia Adelaide = 22.70%, Murcia = 12% ,New South Wales, Sydney = 11.90%.

Table 3.7: Synthetic Control weights for each UK region HPI year on year growth rate.

Regional weights to construct synthetic unit for each region. Weights are estimated based on matching the vector of pre-intervention characteristics of the treated unit $\mathbf{X}_0$ with that of the control group $\mathbf{X}_1$. The weights illustrate the importance of each region j in the donor pool to approximate the synthetic unit.

3.5 Placebo Tests:

To evaluate the reliability of the results obtained in the previous section using synthetic control methodology, in this section I run several placebo tests. The first test is an in-time placebo test, the second test is a country placebo test and lastly leave one control unit out test.

3.5.1 In-time placebo test

I compare the SDLT holiday impact estimated earlier to a placebo impact obtained by reassigning the SDLT holiday to a different period before the SDLT holiday actually took place. Having a large placebo estimate will undermine the confidence that the results in figures 3.5, 3.6 and 3.7 are indeed indicative of the impact of SDLT holiday and are not merely driven by lack of predictive power. The synthetic control algorithm is estimated with a different treatment period and different country weights. The treatment is assigned to a different date, 2019 Q4, which is during the pre-pandemic period when the house prices were already on an upward trend. Figure 3.9 illustrates that there is no significant divergence between the actual UK and synthetic UK during the 2017 Q1 - 2020 Q2 period and there is no substantial effect estimated for 2019 Q4 placebo SDLT holiday. This strengthens the case for the predictive power of the synthetic control method and that the divergence in figure 3.5 successfully estimate the causal impact of the SDLT holiday.

Figures A.5 and A.6 in the appendix display the in-time placebo test applied for all the NUTS 1 statistical regions of the UK. Again the figures show that when applying the SDLT holiday to 2019 Q4, all of the regions were unaffected, except for the South West of England, South East of England and East of England which shows that they underperformed their synthetic counterpart. Together figures 3.6 and 3.7 provide evidence that the increase in the house price index in 2020 Q3 onwards is due to the causal impact of the SDLT holiday.

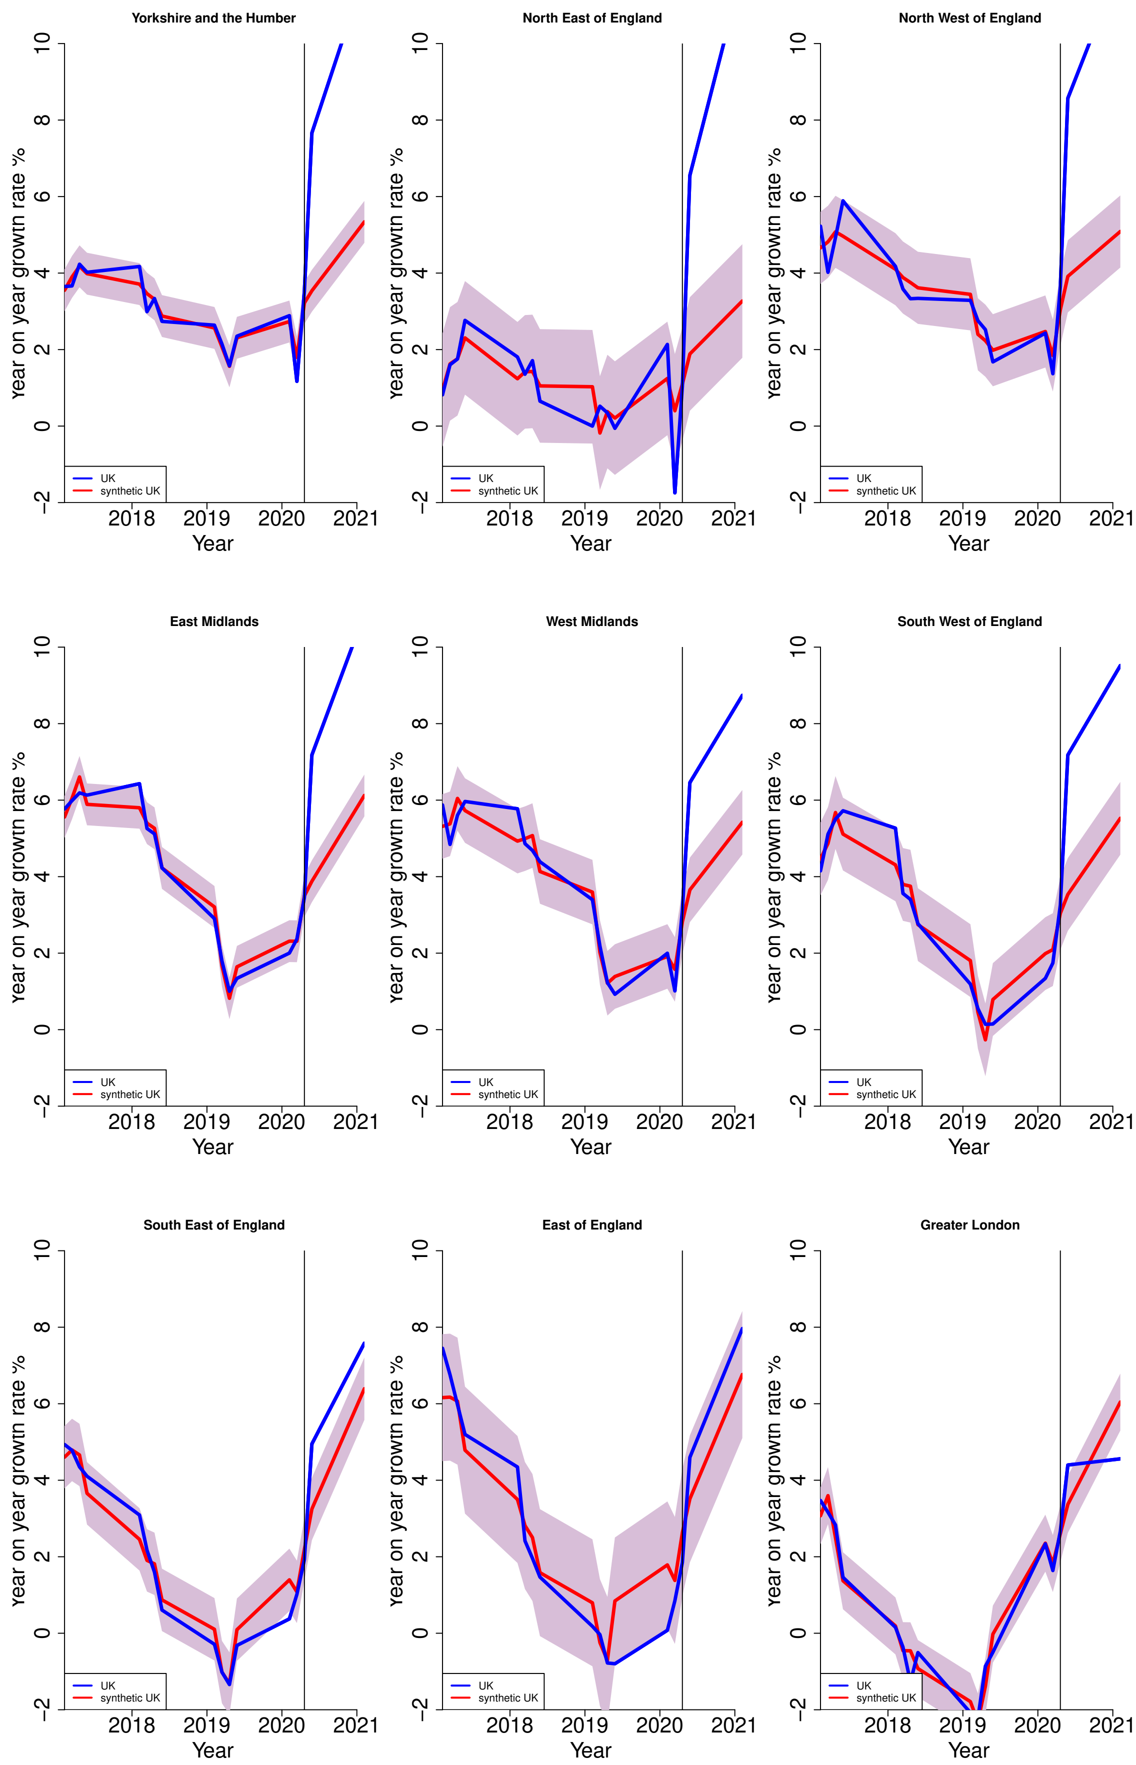

Figure 3.7: Impact of SDLT holiday on all the English regions.

The English regions against their synthetic counterparts house price index (year on year growth rate). The shaded area represents two standard deviations of the pre-intervention gap between the regions and their synthetic counterpart.

Region	SC Estimator λ (1)	DiD Estimator δ (2)	DiD Estimator δ (3)
Wales	3.892***	0.038*	0.012*
	(1.094)	(0.003)	(0.016)
Scotland	2.515**	0.006	0.017
	(0.976)	(0.020)	(0.016)
Northern Ireland	0.762*	−0.006	0.005
	(0.388)	(0.020)	(0.016)
Yorkshire and the Humber	3.312***	0.021	0.032**
	(0.938)	(0.020)	(0.016)
North East of England	4.062***	0.025	0.035**
	(1.291)	(0.020)	(0.016)
North West of England	3.733***	0.020	0.031*
	(1.03)	(0.020)	(0.016)
East Midlands	2.425***	0.010	0.021
	(0.723)	(0.020)	(0.016)
West Midlands	1.968***	0.0001	0.011
	(0.566)	(0.020)	(0.016)
South West of England	2.430***	0.016	0.027*
	(0.702)	(0.020)	(0.016)
East of England	0.606	0.003	0.014
	(0.532)	(0.020)	(0.016)
South East of England	0.857**	0.013	0.024
	(0.358)	(0.020)	(0.016)
Greater London	−0.096	0.021	0.032
	(0.344)	(0.020)	(0.016)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 3.8: Effect of the SDLT holiday on the regional UK HPI YoY.

Estimating the effect of the SDLT holiday on the regional UK HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the UK and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns (2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

(a) 2020 Q4 (b) 2021 Q1

Figure 3.8: Heat map: Impact of SDLT holiday on NUTS 1 Level HPI.

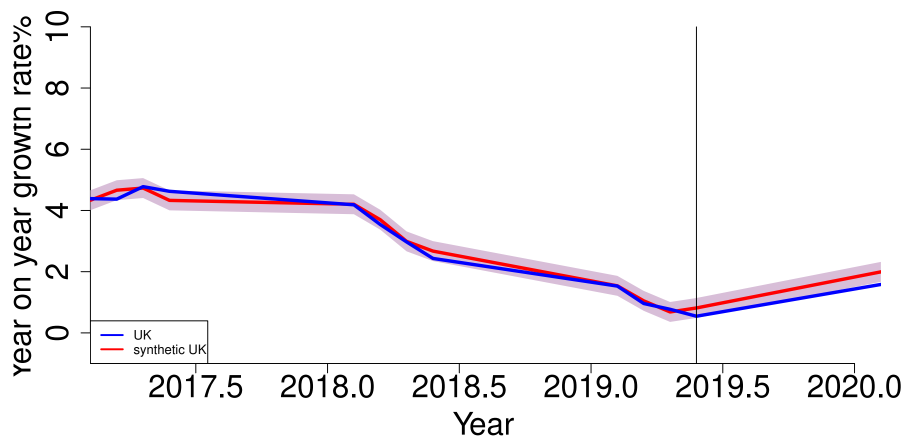

Figure 3.9: In time placebo test.

The synthetic control algorithm is estimated with treatment period 2019Q4. The weights assigned to the countries in the control group are as follows: Korea: 58%, Ireland: 15.7%, Canada: 10.6%, Austria: 9.2% and Norway: 6.3%. The in-sample period is from 2013Q1 to 2019Q4. The shaded area represents two standard deviations of the pre-intervention gap between the regions and their synthetic counterpart.

3.5.2 Country placebo test

To asses the significance of the results demonstrated earlier, I iteratively apply the synthetic control method to each country in the control group that is used to estimate the impact of SDLT holiday. In each iteration, the SDLT holiday is assigned to a country in the control group and the UK is removed from the donor pool. That is, I assess what would happen to their house prices if each country in the control group applied SDLT holiday instead of the UK. Figure 3.10 displays the estimated effect of each country placebo run. Each line is a gap plot derived from the difference of the placebo synthetic control unit of that country and its actual series. The gap between the UK and its synthetic counterpart is illustrated by the black line. The synthetic control method provides an excellent fit for the UK's HPI annual growth rate in the pre-intervention period better than any other country in the control group. The only countries that have a larger gap than the UK in 2021 Q1 are Denmark and Norway, however, they both had a bad fit in the pre-intervention period, making the gap in the post-intervention period unreliable.

According to Abadie et al. (2010), if the synthetic unit failed to fit the HPI annual growth rate in the pre-intervention period, then I would interpret that much of the gap between the actual and the synthetic UK in the post-intervention period is artificially created by the lack of fit rather than the effect of the SDLT holiday, which was the case for Norway.

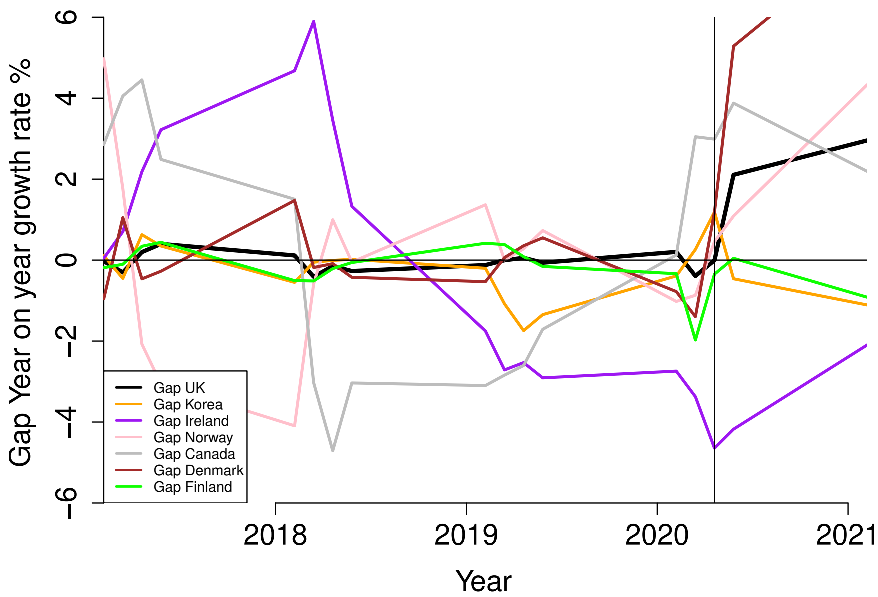

Figure 3.10: Country placebo test.

Gaps in the UK and the significant countries in the donor pool. Synthetic control algorithm is re-estimated for each country that has a non-zero weight in the construction of synthetic UK. The UK is not used in the control group in the construction of any of the country placebo tests.

Similarly figures A.7 and A.8 show the country placebo test applied to each of the regions of the donor pool. The gap between all the British regions and their synthetic counterparts in the post-intervention period, illustrated by the red line, is greater and more positive from the gap, between any region and their synthetic counterpart, when the treatment was assigned to regions in the control group. Therefore, this strengthens the case for the predictive power of the synthetic control method in modelling the impact of the SDLT holiday and indicates that the divergences illustrated in figures 3.6 and 3.7 show the causal impact of the SDLT holiday.

3.5.3 Mean Square Error test

Another way to measure the robustness of the results, is to compute the ratio of the post-SDLT holiday MSE to the pre-SDLT holiday MSE for the United Kingdom and the countries in the control group. The MSE measures the lack of fit between the treated unit and its synthetic counterpart. This ratio can be calculated as:

$$\frac{MSE_{j,post-SDLT-Holiday}}{MSE_{j,pre-SDLT-Holiday}} = \frac{\left(\sum_{t=T_{0+1}}^{T}\left(Y_{jt}-\hat{Y}_{jt}^{N}\right)^{2}\right)/\left(T-T_{0}\right)}{\left(\sum_{t=1}^{T_0}\left(Y_{jt}-\hat{Y}_{jt}^{N}\right)^{2}\right)/T_{0}}, \tag{3.5.1}$$

Where $\left(Y_{jt}-\hat{Y}_{jt}^{N}\right)$ is the treatment effect. This ratio of the MSE is estimated for each country in the control group. Having big value for the ratio indicates that there is a divergence between the trajectories of the treated unit and the synthetic control unit. Since the pre-intervention period had a close fit between the UK and its synthetic counterpart, thus the pre-SDLT holiday MSE is small. The ratio of the post-SDLT holiday MSE to the pre-SDLT holiday MSE indicates the effect of the SDLT holiday in the post holiday period. As illustrated in the previous section, I assign the treatment, SDLT holiday, to all the countries in the control group. I then estimate the ratio of the post-SDLT holiday MSE to the pre-SDLT holiday MSE for each of the control units as illustrated in 3.5.1. This algorithm calculates the ratio for all the countries in the control group including the ones with very small weights. I compare the ratio of the post-SDLT holiday MSE to pre-SDLT holiday MSE of these countries to that of the United Kingdom (UK). Having a large ratio indicates that the gap between the treated unit and its synthetic counterpart in the post-SDLT holiday period is greater than the gap in the pre-SDLT holiday period. Figure 3.11 displays that the UK stands out in comparison to the other donor countries. Therefore, this illustrates the causal impact of the SDLT holiday on the UK housing market.

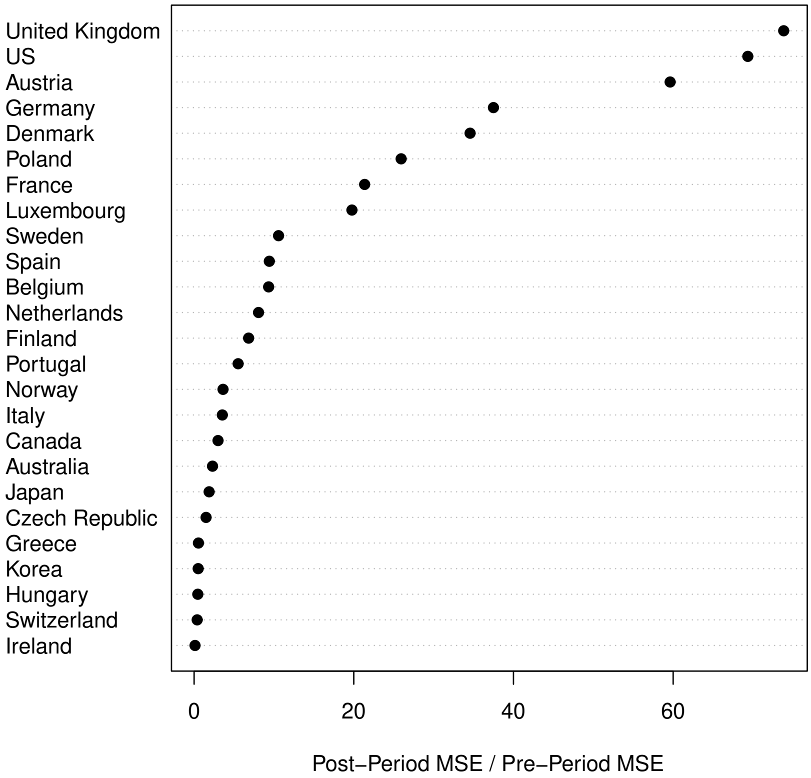

:

Figure 3.11: Ratio of post-SDLT holiday MSE to pre-SDLT holiday MSE.

Synthetic control algorithm is estimated for the United Kingdom and all the 24 countries in the control group. Having biggest value of MSE ratio for the UK indicates the causal effect of the SDLT holiday on the house price index. The Pre-Period is from 2013 Q1 to 2020 Q2, the Post-Period MSE is from 2020 Q3 to 2021 Q1.

To calculate the p-value as illustrated in 3.5.1, if one were to pick a country at random from the sample the chances of obtaining a ratio as high as that of the UK is $\frac{1}{25} \approx 0.084$. According to the p-value, one would reject the null hypothesis that the SDLT holiday had no impact of the UK aggregate house price. When doing the same robustness exercise for each region in figures A.11 to A.22 display the results. I use the package SCtools to perform this placebo test and calculate the p-value for the individual regions. The purpose of this package is that it is suitable when having very large control group as in our case, the control group includes 53 regions as it excludes regions from the control group that have pre-SDLT holiday MSE more than 5 times of that of the UK regions. As displayed for the three British nations, I reject the null hypothesis that the SDLT holiday had no effect on the housing market for Scotland and Wales as their p-values are 0.071 and 0.047 respectively. For the nine English regions, I reject the null hypothesis that the SDLT holiday had no effect on six of the nine regions. The six regions are Yorkshire and the Humber, North East of England, North West of England, East Midlands, West Midlands and South West of England.

3.5.4 Leave one control unit out

I run a series of placebo test to evaluate the sensitivity of my results to changes in the country weights W^*. I iteratively estimate the synthetic control algorithm that solves the minimization problem of equation 3.3.2 to construct synthetic UK. In each iteration, one of the countries in the donor pool is omitted, affecting the goodness of fit of the synthetic control unit. This exercise is re-estimated for 21 iterations. This exercise identifies which country drives my results the most. Table 3.9 shows the weights assigned to each country in the leave-one-unit exercise, the table shows that the results are driven by a combination of Korea and Ireland as they exhibit the biggest weights in all the exercises. Figure 3.12 shows the results of the leave-one-out exercise, the leave-one-out estimates are the grey lines. The results are robust to the exclusion of any particular country from the sample of comparison countries, such that the UK series is above a synthetic control combination.

Synthetic control combination	Countries and W-Weights				
Five units control group	Korea 39%	Finland 24%	Ireland 22%	Canada 9%	Norway 7%
Four units control group	Korea 39%	Finland 29%	Ireland 22%	Norway 11%	
Three units control group	Korea 59%	Ireland 24%	Norway 16%		
Two units control group	Korea 71%	Ireland 29%			
One unit control group	Korea 100%				

Table 3.9: Leave one control out exercise.

The baseline model is re-estimated iteratively to show the countries that contribute the most to the sparse version of the synthetic UK.

Similarly, I apply the same exercise for each region to evaluate the goodness of fit for the synthetic control unit for each of the three nations and the 9 NUTS1 regions. Figures A.9 and A.10 in the appendix show that for all the 9 regions, the leave-one-out estimates

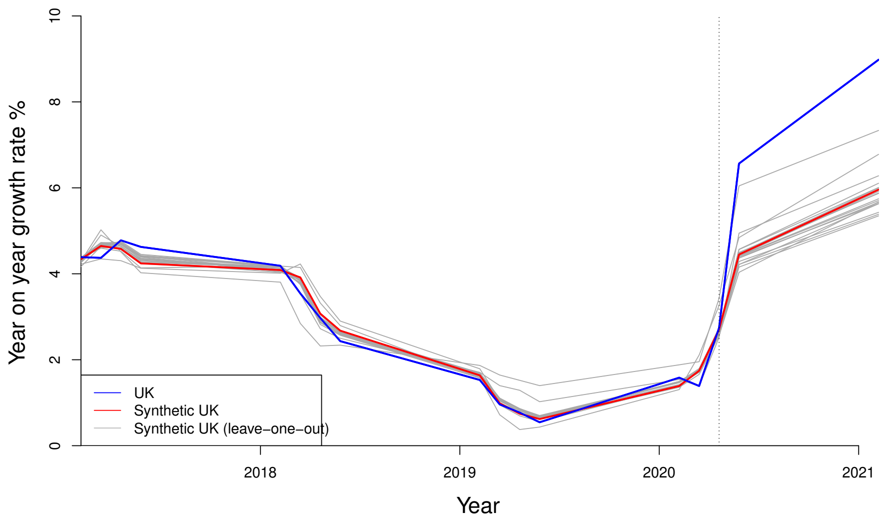

Figure 3.12: Leave one control out exercise.

Leave-one-out robustness exercise for synthetic UK HPI year on year growth rate.

are robust to the exclusion of any particular region from the sample comparison regions.

3.5.5 Change in Covariates

I change the specification of the synthetic control estimation such that the synthetic control unit is estimated using two approaches. The first approach is based on economic covariates along with all the pre-intervention lags of the dependent variable (Billmeier and Nannicini 2013; Bilgel and Galle 2015; Stearns 2015; Villar and Papyrakis 2017; Pellegrini et al. 2021; Papyrakis et al. 2023). The second approach is based on using all the pre-intervention lags only, without the inclusion of economic covariates (Ferman et al. 2020; Breinlich et al. 2020; Botosaru and Ferman 2019). As displayed, across all the specifications and the original estimation of the synthetic control in figure 3.5 the results are unchanged, the UK HPI YoY growth rate significantly over perform its synthetic counterpart following the SDLT holiday till the end of the sample.

3.6 Discussion

It is vital to note that the main question that this paper aim to answer is "What if the SDLT holiday was never introduced?" The paper answers this questions by performing a

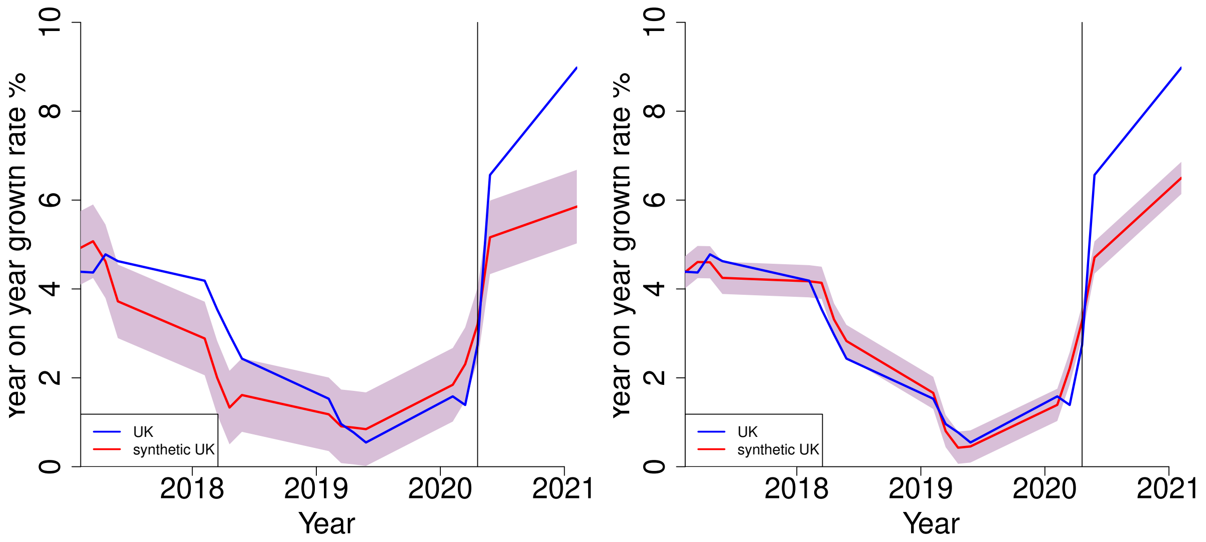

(a) Synthetic UK with economic covariates and all lags of dependent variable

(b) Synthetic UK with all lags of dependent variable as predictors and no economic covariates

Figure 3.13: Covariates Robustness Check

counterfactual analysis to create a unit that mimics the performance of the UK trajectory in the pre-SDLT holiday period. The impact of the SDLT holiday varies across different regions depending on several factors, including local housing affordability, housing market demand and supply and the demographics associated with different locations. To visualise this impact, figure 3.8 shows the difference between the UK regions and their synthetic counterparts on a heat map. The map on the left hand side displays the impact in 2020 Q4, while the map on the right hand side displays the impact in 2021 Q1. It is evident that all the regions have benefited from the SDLT holiday in 2020 Q4 and this benefit increased in 2021 Q1, with the exception of the South East of England and Greater London. The heat map shows that the North West of England and Wales had the greatest increase in HPI in comparison to their synthetic counterpart in 2020 Q4, while the North East of England and Wales had the greatest increase in 2021 Q1. Greater London is the only region coloured in red, indicating the decrease in its house price index in 2021 Q1. The combination of the hybrid/remote work opportunities and the reduced transaction cost due to the SDLT holiday may have encouraged households to explore the possibilities of moving outside of the capital. The number of newly built completed and unsold homes in London increased in 2021 Q1, such that there are 3,100 unsold homes (Savills 2021).

Also, the sales volume has dropped by 16% on an annual basis in 2021 Q1. Similarly the prime market had a decrease in its annual sales by 39% in 2021 Q1.

The performance of the housing market in London can be explained by several factors such as that the average house price in Greater London in 2020 being £633,445 which is above the threshold of the zero SDLT, making buyers unable to benefit from the SDLT holiday. In addition to this, following a year of staying at home, renters and buyers chose to move to larger homes with a garden and garage rather than staying in the capital. A survey by the London Assembly Housing Committee postulates that 15% of the respondents wanted to leave the city as a result of the pandemic.[7] Zoopla (2020a) shows that terms such as garden, garage, parking, detached, rural, secluded, bungalow and annex are among the top 10 most searched terms for property features on Zoopla. These findings contradict the hypothesis presented in Oswald (1996), Van Leuvensteijn and Koning (2004) and Battu et al. (2008) suggesting that home ownership reduces mobility and thus increases unemployment. The introduction of the working from home mandate allowed employees to live further from their place of work whilst working remotely. For example, an employee can have a high paying job in London, while living in the South East or the South West of England. These results confirm the findings of Ramani and Bloom (2021) and Boesel et al. (2021) who show that as a result of the pandemic, the consumer demand has shifted towards more residential space and for properties in communities with lower population density, causing an increase in the rental price of detached houses. In addition to this, Rappaport et al. (2021) postulate that if working from home becomes permanent, employees will consider relocating farther away from their employment, to the outer suburbs of the metropolitan areas.

As discussed above, house prices responded sharply to the SDLT holiday. Land tax prevents transactions that would otherwise be worthwhile from taking place. Agents are highly responsive to property taxes, causing large welfare costs from behavioural responses. Recalling that the SDLT cannot be mortgaged, Best and Kleven (2018) in their theoretical model set out that home buyers with constrained downpayment and high loan-to-value ratios on their mortgages are highly responsive to tax changes because

[7] The same survey show that 416,000 people responded that they would definitely move out of the city in the next 12 months.

they need to pay their transaction taxes upfront. That is to say, with no search frictions, when the SDLT is cut or set on a holiday, there will be an efficient allocation between buyers and sellers. Also, Kopczuk and Munroe (2015) show in their theoretical model the presence of 'unraveling' phenomenon, which can be explained as in the presence of property taxes, transactions with positive surplus may not occur as buyers may prefer to continue searching in order to benefit from locally depressed prices. Therefore, the introduction of the SDLT holiday will increase the buyers surplus, causing an increase in the amount of transactions and housing prices.

Although, the stamp duty generates substantial revenue for the government, the welfare loss associated with its impact on mobility and housing misallocation is significant (Hilber and Lyytikäinen 2017a). The policy implications from the reduction of the SDLT can be positive for households and first time buyers. That is to say, the SDLT acts as a barrier that prevents households from owning their own home, the SDLT can negatively impact housing affordability, especially for the first-time buyers and those looking to move up the property ladder. The tax burden can add significant costs to property transactions, making it harder for individuals and families to enter or move within the housing market. Therefore significant reforms to the SDLT is needed in order to boost the housing market and support households entering or moving up or down the property ladder and prevent housing affordability crisis. The Institute of Fiscal Studies argues that the SDLT is one of the most damaging taxes due to the negative impact it has on the housing market and the labour market (Johnson 2023). Economists argue that abolishing the SDLT completely will result in a loss of £6 billions to the treasury, alternatively reducing SDLT to its previous level, which was a low and reasonably small tax that encouraged mobility without discouraging transactions. However, still this can affect the receipts the treasury can receive. One alternative way that was suggested by Morton (2017) is to reduce the SDLT so that the majority of homes are exempt from paying it. For example, no SDLT would be payable on residential property transactions with values up to £500,000, for higher-value properties, SDLT will still be paid but on a marginal-not-slab system, where rate are applied only over specific thresholds, rather than on the entire property value. As the focus of this paper is to look at the consequences of the SDLT holiday up to £500,000 in England and Northern Ireland and £250,000 in Scotland and Wales, I provide evidence that any cuts or reforms to the SDLT will have strong positive impact on the housing market.

The main assumption that the paper is based on is that the uncertainty around the COVID-19 pandemic and lockdown and economic concerns might have led potential buyers to postpone or cancel their purchasing decisions, reducing overall demand in the housing market. The SDLT holiday was introduced to stimulate the housing market, the paper looks at if the SDLT holiday did offset this decrease in the housing market. I achieve this through removing any country from the donor pool that had similar tax reliefs. However, the estimated effects of the SDLT holiday may be influenced by the broader pandemic-related factors such as the the UK intervention and support provided in the form furlough schemes, mortgage payment holidays, and business support loans, could have indirectly influenced the effectiveness of the SDLT holiday. Another assumption of the Synthetic Control Method is that the treated unit is not affected from any other treatment. One aspect this analysis did not look at is the impact of the Help to Buy Scheme on the UK housing market during the period of the SDLT holiday. The Help to Buy Scheme is a scheme where the government provides an equity loan to home buyers buying a new built home. The scheme was introduced in 2013 and ended in March 2021, which falls during the sample estimation period. Although, it is hard to disentangle the effect of the SDLT holiday from that of the Help to Buy Scheme on the dependent variable (HPI YoY growth rate) in our analysis. The Help to Buy scheme is only applied to newly built homes, in figure 3.14 I re-estimate the synthetic control method, however with the HPI YoY growth rate for existing (old) dwellings as the dependent variable and the HPI YoY growth rate for newly built dwellings as dependent variable separately which are assumed to be the target of the Help to Buy Scheme. It can be visible that both types of dwellings benefited during the period of the SDLT holiday. The existing dwellings had a perfect fit with its synthetic counterpart in the pre-SDLT holiday period and the p-value based on equation 3.5.1 is 0.08, unlike the new dwellings which was not perfectly captured by its synthetic counterpart in the pre-SDLT period and the p-value is 0.25. That is, even in the absence of the Help to Buy scheme, the UK housing market

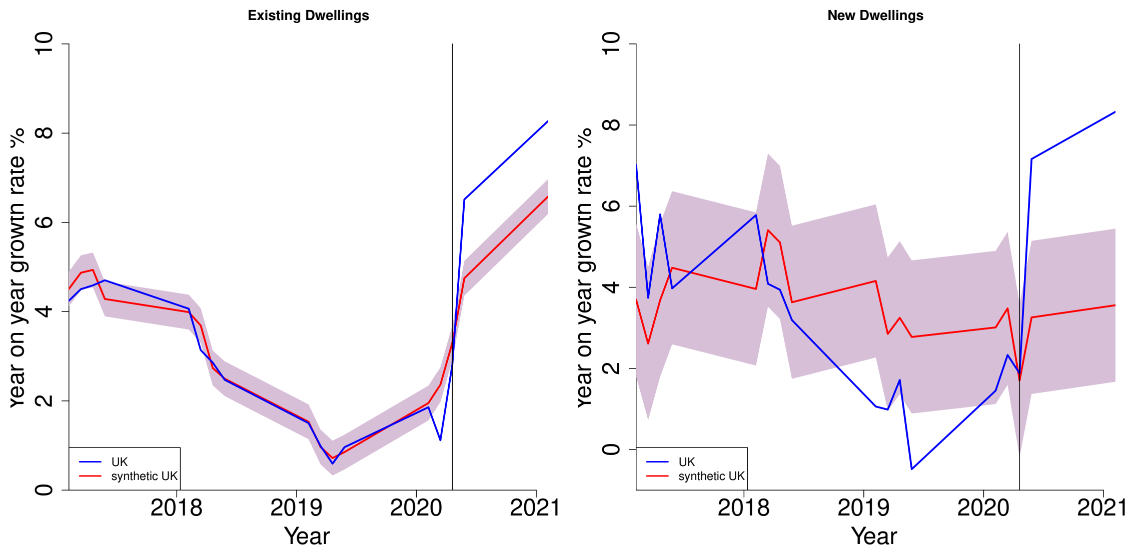

Figure 3.14: Synthetic Control Method estimation on dwellings type.

Synthetic control UK for existing dwellings is made of Korea, Ireland, Norway and Austria, p-value = $\frac{1}{13} = 0.08$. Synthetic control UK for new dwellings is made of Belgium, Italy, Germany and Norway, p-value=$\frac{4}{16} = 0.25$.

would still benefit from the effect of the SDLT holiday.

3.7 Conclusion

This paper investigates the impact of the SDLT holiday using the synthetic control methodology. Most of the literature studying the impact of the SDLT holiday looks at the economic wide impact using surveys data, which ignores the underlying economic geography. This paper fills this gap in the literature by looking at the impact on the aggregate UK and regional house prices. The results of this paper suggest that lowering the tax rate can stimulate the activity in the housing market as agents took advantage from widening the non-taxable bracket.

My findings show that the UK housing market is highly responsive to changes in the property transaction taxes. The house prices responded significantly to the SDLT holiday as the economic theory predicts. The price responses illustrated by the gap between the UK and its synthetic counterpart show that buyers took advantage of the temporary tax cut resulting in large-scale re-timing of the property transactions. The 12 months SDLT holiday was able to stimulate the housing market, causing an increase in the house price index for the UK.

More precisely, my findings suggest that the SDLT holiday caused the aggregate UK house price index to increase by 2.11% and 2.96% in the two quarters following the SDLT holiday. In terms of the regional analysis, all the NUTS 1 statistical regions of the UK benefited from the SDLT holiday with Wales and North West of England and North East of England having the greatest increase in their house prices. However, Greater London had a decrease in house prices in 2021 Q1. These results confirm the findings of growing literature on the effect of COVID-19 on work displacement. As a result of the pandemic homeowners moved outside of the capital and big cities in favour of the countryside and coastal cities.

The results of this paper confirms the findings of policy literature that stamp duty land tax should be reformed in order to benefit the economic activity and provide support for households. The results also shed light on the importance of fiscal stimulus especially property tax cuts in simulating the housing market and the real economic activity during times of economic slowdown.

Chapter 4

Predicting the Volatility of the S&P 500 Index: The Role of News Sentiment

Abstract

This chapter decomposes the textual data in the U.S. Newspapers into different news topics using unsupervised machine learning approach. I then measure the tonality of the news topics using dictionary based approach and investigate their abilities in predicting the volatility of stock returns. The results of this paper demonstrate that incorporating topical news sentiment as exogenous regressors in GARCH(1,1) and GJR-GARCH (1,1) models explains the stock returns volatility. More precisely, the news sentiment indices for topics related to currency, monetary policy, manufacturing, trade and trade policy can explain the variation in the volatility of stock returns. When combining the individual forecasts using regression based method, the combined forecasts provide accuracy in predicting the volatility of stock returns.

4.1 Introduction

Understanding the link between news media and the financial markets is a growing field in financial economics. According to Shiller (2000) media reflects and shapes the investors' and managers' beliefs, which affect their trading behaviour. Negative news can influence the desire of investments for risk averse investors. Such that during uncertain times, the tone of the news reported has effects on investors' expectations. In finance, texts from different sources of news such as financial news, social media, and company filings can be used as a source of prediction for asset prices. The aim of this paper is to identify the sources of economic news that can be used to enhance the prediction of the stock returns volatility.

Tetlock (2007) is one of the earliest to employ a dictionary based method, using Harvard GI dictionary, to evaluate the impact of the sentiment of the news on the stock market using VAR methodology, specifically looking at media pessimism. Their results show that pessimistic news from the Wall Street Journal "Abreast of the Market" column predict the daily returns and volume of the U.S. Dow Jones Index. Similarly, Garcia (2013) constructs a measure of news sentiment based on two columns of financial news from the New York Times. Their results show that their measure of news sentiment has ability to predict the average return of the Dow Jones Index over a one day horizon. A big strand of the literature on the relationship between news sentiment and the stock returns uses already classified aggregate sentiment measures. These papers use news analytics software packages such as Thomson Reuters News Analytics (TRNA) and Ravenpack. For example, Uhl (2014) uses the positive and negative Thomson Reuters news sentiment index to forecast the returns of the Dow Jones index. Smales (2015), using Thomson Reuters News Analytics firms specific news to construct an aggregate measure of news sentiment, shows that news sentiment can predict stock returns for different industries. Smales (2016) show that the Ravenpack news sentiment index can predict the stock returns and the volatility (proxied by VIX). One limitation of the sentiment indicators constructed using these software packages is that the software vendors focus on firm-specific news and ignore news outside the stock market index. Constructing measures of news sentiment using primarily firms and stock market news ignore the influence of news related to other factors such as macroeconomic news, political news and geo-political news which can have an affect on the stock market. Manela and Moreira (2017) construct news based measure of uncertainty based on the Wall Street Journal. They show that the majority of the variations in their index are driven by news categories related to Government, Wars and Stock Market. Their findings are that their measure predicts returns at frequencies from six months up to 24 months. Similarly, Adämmer and Schüssler (2020) find that news related to geo-political events predicts the equity premium for the S&P 500 index measured by the Out-of-Sample R-squared.

According to Peng and Xiong (2006), investors have limited cognitive ability thus they should be selective with the information they process. Da et al. (2011) demonstrate that the wealth of information creates a poverty of investors attention. Taking these theories in consideration in this paper, I use the unsupervised machine learning approach to decompose textual data in the U.S. newspapers into different news topics using the Latent Dirichlet Allocation (LDA) of Blei et al. (2003). The LDA can uncover hidden themes in large corpus of written documents without the need to provide any prior information on the documents themselves (Adämmer and Schüssler 2020). I then measure the tonality of the news topics over time using the dictionary based approach of Harvard General Inquirer dictionary. I use the final product of combining the topic modelling with sentiment modelling to have a time series for each news topic and evaluate the ability of news tonality in forecasting the volatility of the stock returns. The aim of this paper is to evaluate which category of news sentiment can provide an improvement in forecasting the volatility of the S&P 500 returns. To perform this analysis, I develop volatility forecasting models, Generalized Auto Regressive Conditional Heteroscedasticity (GARCH) and GJR-GARCH models and introduce topical news sentiments as exogenous regressors in these models. I evaluate the volatility model by considering the significance of the in-sample results and the out-of-sample forecasts. My results show that the volatility forecasts are improved when introducing topical news sentiment indices as exogenous variables in the GJR-GARCH models. More precisely, modelling the asymmetry effect in the GJR-GARCH model with using the topics monetary policy, trade, trade policy, manufacturing, currency and geo Politics can explain the volatility of stock returns. When performing a forecast pooling exercise by combining the individual models forecasts using regression based method, I find that the combined news sentiment topics offer better out-of-sample forecast.

Other papers looked at testing the ability of aggregate news sentiments in forecasting the volatility of stock returns. Hsu et al. (2021) decompose aggregate news sentiment into a positive measure and a negative measure. They show that the negative aggregate news sentiment index can forecast the volatility of the Taiwanese stock index. C. Liang et al. (2020) compares the sentiment from the newspaper to sentiment from social media and internet media in forecasting the Chinese volatility of stock returns. They show that social media and online media can improve volatility forecasting. The contribution of this paper to the literature on news sentiment and volatility prediction is that this is the first paper that looks on decomposing the news into different topics and evaluate the sentiment of each of these topics in forecasting the volatility of the S&P 500 returns. As well as combing the individual models forecasts into one measure of economic forecasts that provides better out-of-sample predictions. The closest papers to this analysis are that of Adämmer and Schüssler (2020), Larsen and Thorsrud (2019), and Ardia et al. (2019). Adämmer and Schüssler (2020) use Correlated Topic Modelling (CTM) to decompose the U.S. newspaper into different topics and evaluate the ability of media coverage on different topics in predicting the equity premium. Their findings are that geopolitical news topic has the greatest forecasting gains among all the other topics. This paper is different in the aspects that I complement the topic modelling by sentiment analysis to model the tonality of the news topics over time and evaluate their ability in predicting the volatility of stock returns. Larsen and Thorsrud (2019) paper is different to this one, as in their paper they decompose the Norwegian newspapers into topics and link the topics to companies to assess their ability to predict the stock prices. Their news sentiment is more akin to company/industry specific news. Whereas, the news sentiment in this paper is more of a narrow economic measure coverage. Lastly, to decompose the news articles into different news topics, Ardia et al. (2019) use the LexisNexis Smart Indexing tool. Using this tool results in the articles could be related to multiple topics, which increases the chance of having correlated topics. However, by using probabilistic models such as Latent Dirichlet Allocation for topics identification, as in this paper, I mitigate that problem especially that the approach I use to choose the optimal number of topics reduces the correlation within topics.

The rest of this paper is organised as follows. Section 4.2 describes the data, LDA topic modelling, and the creation of the sentiment indices. Section 4.3 introduces the GARCH and GJR-GARCH models. Section 4.4 show the results of incorporating news sentiment in the volatility models. Section 4.5 concludes.

4.2 Data

In this paper, I evaluate the relationship between the news and the stock returns. The raw data of the newspaper corpus is obtained from the LexisNexis database, which includes a variety of journalistic articles, business reports and legal documents. To preserve space, details about the news corpus is available in the appendix.

In the so-called information age, investors' attention is a scarce cognitive resources. Instead of reading entire newspaper articles, investors tend to process more market and sector wide information (Peng and Xiong 2006). In identifying the relevant news to be analysed, I use the Latent Dirichlet Allocation (LDA) approach of Blei et al. (2003) to decompose the news articles of the major U.S. newspapers into different topics to capture their development over time. I then evaluate the sentiment of each news topic, to evaluate how the tonality of each news topic changed over time using Harvard GI Psychological dictionary. Figure 4.1 below shows a flow chart of the undertaken steps to construct the topical news sentiment indices.

The use of unsupervised vs supervised machine learning approaches in the context of textual data analysis has been growing in the economics and data science literature. Unsupervised machine learning aims to finding patterns or structures in the data without any prior labeled examples or specific guidance. This approach is useful when the data is unstructured or when the aim is to explore and gain insights from the data. Unsupervised learning algorithms include clustering methods such as k-means clustering, hierarchical clustering and dimensionality reduction techniques like principal component analysis (PCA) and topic modelling techniques such as Latent Dirichlet Allocation (LDA). The LDA is unsupervised technique as it assumes that each document is a mixture of various topics, and each topic is represented by a distribution of words. LDA iteratively assigns words to topics and topics to documents, trying to maximize the probability of generat-

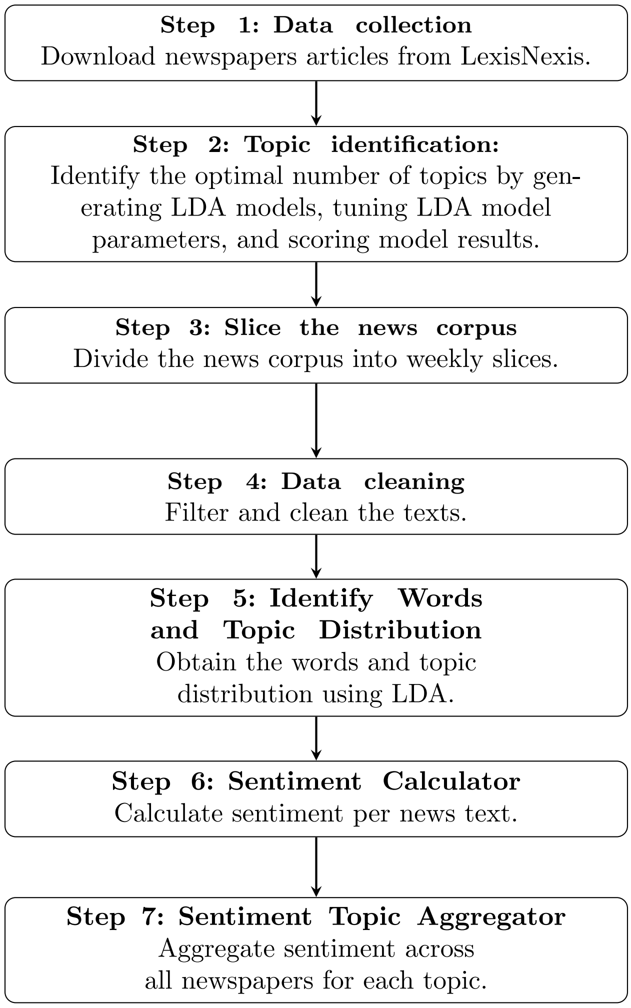

Figure 4.1: Topic modelling flow chart

ing the observed data. It independently identifies patterns and relationships within the text corpus, clustering similar documents together and assigning words to different topics based on their co-occurrence patterns.

On the other hand, supervised machine learning approach involves training a model using labeled data, where each input example is associated with a corresponding output or target value. The algorithm learns from these labeled examples and generalizes to make predictions on new, unseen data. In the context of textual analysis, this approach is applicable to be used to perform sentiment analysis using dictionary based methods. In the dictionary based methods, the dictionaries are collections of word categories or classes that are manually defined and annotated by human experts. Each word in the dictionary is assigned to one or more predefined categories based on its semantic or contextual meaning. The most commonly used dictionaries are the Harvard General Inquirer

(GI) dictionary and the Loughran and McDonald (2011) dictionary. The approach uses dictionaries as references to determine the categories or classes to which the words in the input text belong.

Combining both approaches together provides an illustration of the use of machine learning in economics and finance.

4.2.1 Newspaper data

The following steps were undertaken to construct the topical news sentiment indices: To perform newspaper topic modelling and sentiment analysis, I retrieved a set of economic news articles from LexisNexis for major US newspapers.[1] The articles come from newspapers such as The New York Times, USA Today, Daily News (New York) Pittsburgh Post-Gazette, Tampa Bay Times, St.Louis Post-Dispatch (Missouri), The Atlanta Journal-Constitution, The Tampa Tribune (Florida), The Philadelphia Inquirer, The Daily Oklahoman (Oklahoma City, OK), Star Tribune (Minneapolis MN), The Christian Science Monitor, The Philadelphia Daily News, Daily Journal of Commerce (Portland, OR), and Alaska Journal of Commerce. To ensure that the news corpus covers economic news, I restrict all the newspaper documents to include "Econ" as the stem of the word economics. The sample starts from 1st of January 1990 till 31st December 2015, the news corpus includes 115,371 articles. The newspaper corpus satisfy the following criteria:

1. The geography is browsed by the United States to ensure that the news corpus focus on the news related to the US.

2. The articles are restricted to articles written in English.

3. The subject terms are Economy & Economic Indicators and Trade & Development.

4. The news is filtered to Newspapers only and I eliminate any articles that are less than 250 words as very short articles are likely to be noisy.

[1] The license used to obtain the newspapers articles does not include access to The Wall Street Journal (WSJ) and Washington Post which are highly circulated newspapers in the United States.

4.2.2 Data cleaning

As a starting point to clean and organise my newspaper corpus. For each article, I remove all repeated articles, the punctuation and stop words, numbers, all the words that occur less than 5 times.[2] I apply a stemming procedure such that all the words are converted to their roots, for example, the stem of the word economics is "econom". This results in having 83,000 articles.

The topic modelling analysis that I perform involves estimating the topics in real time on a weekly basis. My news corpus has been divided into weekly slices. Such that for every week, I collapse all the articles published in that week into one document and compute the topics and word distributions in this one newly formed document. Performing topic modelling in that way not only captures the evolution of topics over time, but also ensures that the set of keywords for each topic does not include any future information beyond that week that can distort the forecasting exercise performed in 4.3

When reading a newspaper article, it can be made of multiple topics, and the words in that article reflect the particular set of topics it addresses. In order to model the contribution of different topics to that article, the topics should be treated as a probability distribution over words (Griffiths and Steyvers 2004). To decompose my news corpus into different news topics, I use the unsupervised machine learning algorithm, Latent Dirichlet Allocation (LDA) of Blei et al. (2003) and Blei (2012).

4.2.3 Topic Modelling: Latent Dirichlet Allocation

The main idea of the LDA analysis is to learn the word and the topic distribution of the underlying news corpus. It is an unsupervised machine learning algorithm meaning that there is no need for prior labeling of the articles or their content. The algorithm is known as Latent Dirichlet Allocation as the data generating process of the algorithm depends on two sets of Dirichlet distribution. A document can be viewed as a probabilistic mixture of different topics.

Figure 4.2 below displays a plate notation for the LDA algorithm. I use the notations that were highlighted in Larsen and Thorsrud (2022). The outer plate represents the

[2] Stop words are words that do not add much meanings such as I, me, they,... etc.

entire corpus with M newspaper articles. The inner plate represents the total number of words for all the documents N. The distribution of topics for any given article is given by θ_m, $m \in [1, 2, ..., M]$. The distribution of words for each topic is given by φ_k, $k \in [1, 2, ..., K]$, where K is the number of topics estimated. As illustrated by their links, both θ_m and φ_k follow a Dirichlet distribution with hyperparameters α and β respectively. For each word position n in the news document we draw a topic assignment $Z_{m,n}$ and words $W_{m,n}$. Consequently, each news article consists of repeated choice of topics $Z_{m,n}$ (topic assignment) and words $W_{m,n}$ that are drawn from Multinomial distribution using θ_m and φ_k. Unlike all the other letters, $W_{m,n}$ is coloured in grey to indicate that it is the only observable variable in the model.

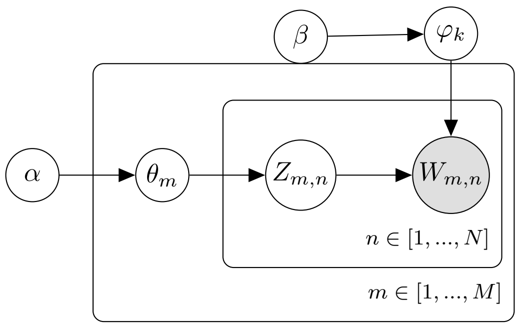

Figure 4.2: LDA model illustration using plate notation.

The hyperparameter α is the Dirichlet prior parameter of the per document topic distribution. The hyperparameter β is the Dirichlet prior parameter of the per topic word distribution. θ_m is the article-topic distribution. φ_k is the topic-word distribution.

As described in figure 4.2, I follow the multi-stage process highlighted in Blei (2012).[3] In this process the per document topic distribution and the per topic word distributions depend on the hyperparameters α and β respectively. High values of the α hyperparameters will result in less dispersion of topics in an article, while low values for it indicate more disperse distribution of topics within an article. That is to say, a low value for α means a higher chance of documents being assigned to one single topic with a high probability, resulting in the articles documents consisting of fewer topics. Similarly, high values for β indicate that the probability of distribution of words to topics being even, while low values indicate that fewer words have a much higher probability of defining each topic. The multi stage process can be explained through two sets of distributions:

1. Distribution over topics: The theme of the articles is picked by giving the articles a

[3] I use the R package "topicmodels" Grün and Hornik (2011) to perform topic modelling.

distribution over the topics, given by the θ_m, the allocation of the topic distributions for the document is determined by

$$\theta_m \sim \text{Dirichlet}(\alpha),\ m \in [1, ..., M].$$

2. Distribution over words: The words distribution for each topic is given by:

$$\varphi_k \sim \text{Dirichlet}(\beta),\ k \in [1, ...K].$$

3. For each word in the document:

   (a) Choose a topic from the distribution of topics in step 1, $Z_{m,n} \sim \text{Multinominal}(\theta_m)$.

   (b) Given that chosen topic, choose a word randomly from the corresponding distribution over the vocabulary, $W_{m,n} \sim \text{Multinominal}(\varphi_{Z_{m,n}})$.

According to steps 1 and 2 θ_m and φ_j follow Dirichlet distribution with hyperparameters α and β respectively. Steps 1, 2 and 3 can be summarised as for each article, each word is chosen according to first selecting a topic (step 1) and then selecting a word that is associated with this topic (step 2). Step 3a illustrates that the selected topic is chosen from the per-article distribution over topics. Step 3b defines that each word in each article is drawn from one of the topics.[4] The calibration of the hyperparameters α follows Griffiths and Steyvers (2004), such that it is allowed to vary based on the number of topics K, while the hyperparameter β is allowed to vary based on the number of words in all the documents W. Although Griffiths and Steyvers (2004) use a calibrated value for β as 0.1, I calibrated β in a data driven approach such that the hyperparameter varies according to the number of words in the documents. In Griffiths and Steyvers (2004), the corpus included scientific research. Using the same value of 0.1 in economics news may not be appropriate.[5]

$$\alpha = \frac{50}{K} \quad , \quad \beta = \frac{200}{N}.$$

Choosing K is a model selection problem. A small number of topics could result in overlapping between topics i.e. some words may appear in more than one topic. However,

[4] See Blei et al. (2003) for details of the data generating process of LDA topic modelling.
[5] Larsen and Thorsrud (2022) also calibrate α and estimate β in the same way.

a large number of topics could result in them being redundant and hard to interpret (Q. Liang et al. 2021). I follow the Gibbs sampling algorithm of Griffiths and Steyvers (2004) and the density based method of Cao et al. (2009) in choosing the optimal number of topics. Griffiths and Steyvers (2004) use a Bayesian statistical method in determining the optimal number of topics using Gibbs simulations.[6] The optimal number of topics is the one that maximised the likelihood $P(w|K)$ where w is the words in the corpus and K is the number of topics.[7] Figure 4.3 shows the estimates of $P(w|K)$ across different number of topics K ranging from 5 to 25 topics. The $P(w|K)$ increases as a function of the number of topics, peaking at 17 topics then decreasing before it exceeds that value for the number of topics of 20 to 25.

On the other hand, the optimal number of topics according to Cao et al. (2009) relies on the topic density through minimizing the average cosine similarity across topics. The cosine similarity (S_C) is a common metric that is used to measure the degree of topic similarity between documents. The cosine similarity measures the similarity between two vectors of an inner product space. It is measured by the cosine of the angle θ, where θ is the angle between the two vectors x and y. The cosine similarity between the two vectors x and y can be written as:

$$S_C(x, y) = \cos\theta = \frac{\mathrm{x}^{\mathbf{T}}\mathrm{y}}{\|\mathrm{x}\|\,\|\mathrm{y}\|} = \frac{\sum_{i=1}^{n} x_i y_i}{\sqrt{\sum_{i=1}^{n} x_i^2}\sqrt{\sum_{i=1}^{n} y_i^2}},$$

where $\|x\|$ is the Euclidean norm of vector x and $\|y\|$ is the Euclidean norm of vector y. The resulting similarity can range between -1 to 1. Having a cosine similarity of -1 meaning that the vectors are exactly opposite, while a similarity of 1 indicates that they are exactly the same. Having values in between 0 and 1 indicates an intermediate degree of similarity.[8] A cosine similarity of 0 indicates orthogonality between the vectors, meaning that the two vectors are at 90 degrees of each other. A cosine similarity close to 1 indicates a smaller value for the angle θ and therefore high similarity between the vectors.

In the context of topic modelling, the vectors x_i and y_i are vector representation of

[6] To determine the optimal number of topics, I use the R package *lda tuning* (Nikita and Chaney 2020).
[7] $P(w|K)$ can be estimated as the harmonic average of $P(w|z, K)$ when z is sampled from the posterior $P(z|w, K)$.
[8] Cao et al. (2009) verified that in topic modelling the cosine similarity falls between 0 and 1.

each topic, obtained based on the topic–word probability distribution matrix. As I am dealing with economics news, they all have some degree of similarity, thus, I will expect that the angle between two topic vectors θ cannot be greater than 90°. This makes the value of similarity between topics to be between 0 and 1. The cosine similarity (S_C) between the topic pairs T_i and T_j based on the topic-word probability distribution can be written as:

$$S_C(T_i, T_j) = \cos\theta = \frac{T_i^{\mathbf{T}} T_j}{\|T_i\| \, \|T_j\|} = \frac{\sum_{v=0}^{V} T_{iv} T_{jv}}{\sqrt{\sum_{v=0}^{V} (T_{iv})^2} \sqrt{\sum_{v=0}^{V} (T_{jv})^2}}. \tag{4.2.1}$$

A small value of similarity indicates that the topics are more independent and therefore little topic overlap.[9] The average distance of the topic can be calculated as:

$$\text{average distance (structure)} = \frac{\sum_{i=0}^{K} \sum_{j=i+1}^{K} S_C(T_i, T_j)}{K \times (K-1)/2}. \tag{4.2.2}$$

A smaller value indicates that the model is more stable i.e. fewer overlap. As illustrated in figure 4.4 the average distance has a big dip when the number of topics K is 17 before it starts to increase again. Therefore, I will proceed with using 17 topics as the optimal number of topics.

Once the news corpus is distributed across 17 topics and the word distribution across different topics have obtained. The output is unlabelled topics, just numbered topics, for example topic 1, topic 2,. . ., etc. Referring to the topics by numbers does not give economic intuition. The topics are then labeled by visual inspection of the top keywords distributions. The label of the topic reflects the theme of that topic. For instance, if a newspaper article has high proportion of words such as "interest", "inflation", "rate" then it is accounted for as a part of the monetary policy topic. Table 4.1 shows the top five words that appear for each topic over time. The words may change from time to another as some words may appear and disappear based on the state of the events happening in the economy. For example in the monetary policy topic, as the topic models follow real time estimation, in 1998 the word "greenspan" for Alan Greenspan appears as one of the

[9] Cosine similarity S_C differs slightly from Pearson correlation measure $\rho_{x,y}$. The Pearson correlation can be written as: $\rho_{x,y} = \frac{\sum_{i=1}^{n}(x_i - \bar{x}_i)(y_i - \bar{y}_i)}{\sqrt{\sum_{i=1}^{n}(x_i - \bar{x}_i)^2 \sum_{i=1}^{n}(y_i - \bar{y}_i)^2}}$ Pearson correlation and the Cosine similarity can be equal if and only if the vectors x and y have zero mean, which is not the case here as the vectors are obtained based on the topic-word probability distribution.

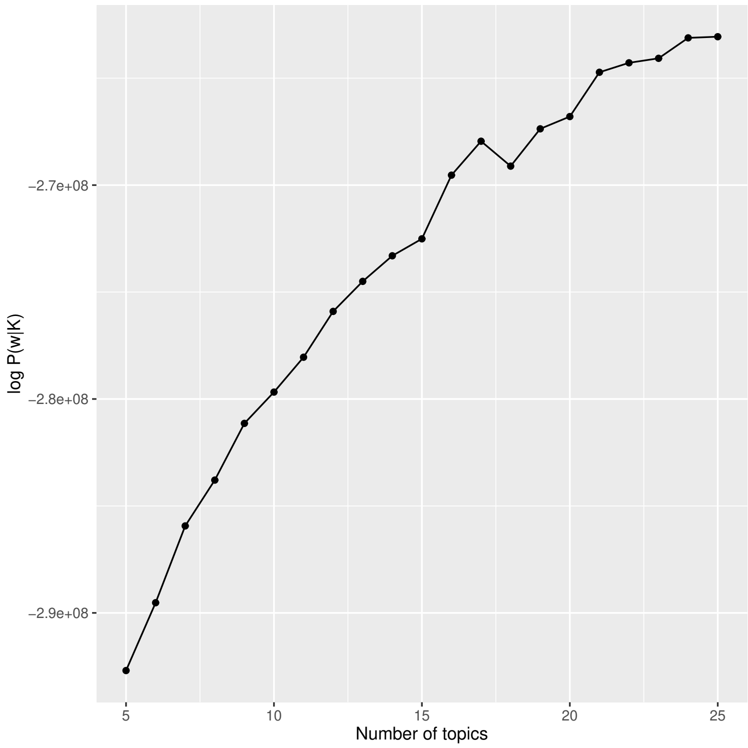

Figure 4.3: Griffiths and Steyvers (2004) optimal number of topic selection.

Displaying the log-likelihood of the corpus for different number of K topics.

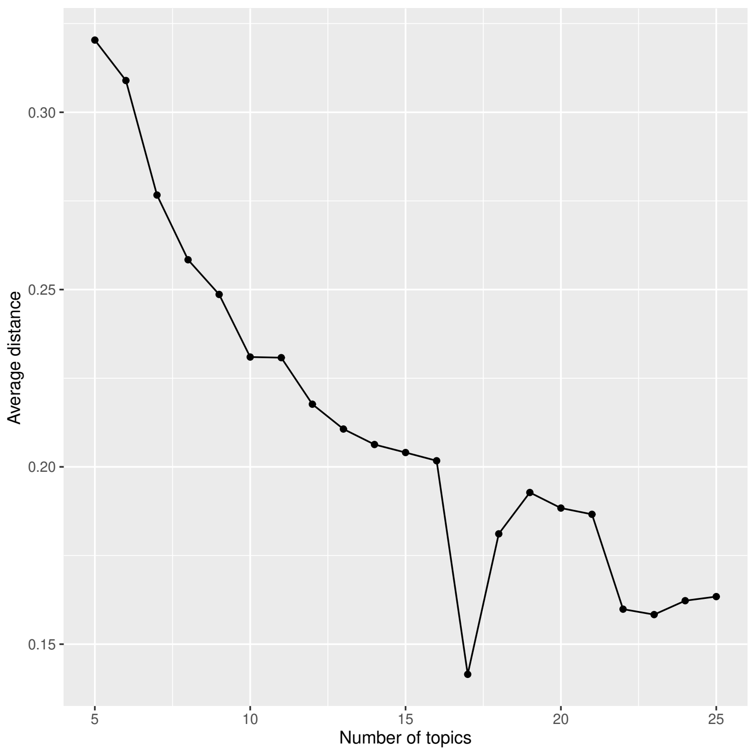

Figure 4.4: Cao et al. (2009) optimal number of topic selection.

Lower average distance indicates more stable model i.e. better model performance.

keywords. However, during the global financial crisis, among the top words that always appeared were "credit", "debt", "bernank", "mortgage" appeared. The word "bernank" stands for Ben Bernanke the Fed Chairman during that period. Performing the LDA topic modelling in this way ensures that that the keywords for each topic follows a real time estimation. For example, the word "Trump" prior to 2015 used to appear in the company topic as it was associated with Trump's business ventures, while afterwards it appears in the topic political economy as this was the time when Donald Trump was running for the U.S elections. Similarly there are some words that can appear at some point of time and disappears afterwards.

Table 4.1 shows the topic-word distribution for each topical news sentiment. The top words column shows the top five words that appear all the time for each topics. Other words can be mentioned as keywords depending on the time and the events taking place. For example, in the geo-political topic during the time of the Gulf war and the Iraq war, the word war appeared more frequently. However in 2013 around the time of the nuclear deal negotiations between Iran and the U.S., the words Iran and nuclear appeared more frequently.

Label	Top words
Monetary policy	rate, fed, bank, money, inflat
Stock market	market, stock, share, investor, dow
Commodities and energy	gas, energi, oil, drill, coal
Fiscal policy	budget, tax, fiscal, stimulus, spend
Economic performance	report, economi, unemploy, growth, quarter
Local economy	counti, region, local, state, tampa
Politics	republican, democrat, presid, senat, congress
Geo-politics	nuclear, war, iran, sanction, bush
Trade	import, trade, export, deficit, surplus
Trade policy	nafta, zone, mexico, global europ
Foreign economy	europ, imf, america, mexico, britain
Manufacturing	manufactur, factori, order, product, order
Company and business	inc, execut, compani, board, market
Travel and tourism	travel, tourism, hotel, florida, visitor
Housing	hous, home, properti, mortgage, foreclosur
Construction and development	develop, build, project, construct, plan
Currency	dollar, rate, currenc, yen, japan

Table 4.1: Descriptive analysis for each news topic and their associated top words.

The first column shows the label assigned to each topic based on their top key words. The second column shows the top words for each topic.

Different words have different weights in their topics across different contexts. For instance, the word credit becomes important for the monetary policy topic during the recession period, while disappears or have very small weight during the post-recession period. To illustrate this, I show in figure 4.5 word cloud for the monetary policy topic during the recession period, which starts from December 2007 to June 2009 and the post recession period which starts from July 2009 till December 2015. The words with larger font have higher probability of appearing in that topic during that period based on the posterior distribution of the words and the documents. As displayed in panel (a) during the recession period, the words credit, crisi, mortgage, loan and secur have bigger fonts due to their importance during this period. However, for the post recession period, these words either disappeared or appear with smaller font to display the lack of the importance for them.

4.2.4 News Sentiment

Following the decomposing of the news into different topics, to understand the tonality of the news, I transform the news into a sentiment measure. Each news article document will be assigned a sentiment score as either positive, negative or neutral. I use a dictionary based machine learning approach to perform lexical sentiment analysis for each of the 17 topics obtained by the LDA method over time. Lexicon-based sentiment analysis has an advantage over the commonly used measures of economic sentiment in terms of cost, timeliness and scope. There are several concerns of the commonly used measures of economic sentiment, as consumers and business surveys they suffer from several problems such as being costly to obtain and ignoring key information beyond the firms and the stock market news.

The main idea of the Lexicon-based sentiment analysis is the qualification of linguistic patterns (e.g. words or sentences) as positive, negative, or neutral using predefined lists called lexicons,(Ardia et al. 2019). I follow the literature of news sentiment in economics and finance by using the General Inquirer (GI) Harvard IV-4 dictionary when applying the standard bag-of-words approach to compute the net sentiment of each document. The dictionary is composed of 3628 words which are decomposed of 1637 positive words

(a) Monetary Policy topic during recession period.

(b) Monetary Policy topic during the post-recession period.

Figure 4.5: Topic Modelling Over different periods.

Estimation of the topic modelling during the recession and post recession periods. The size of each word is based on posterior distribution of words and documents to topics.

and 1991 negative words.[10] According to Tetlock (2007), Larsen and Thorsrud (2022), the Harvard GI dictionary is more applicable for economic and non-economic news than the financial dictionary of Loughran and McDonald (2011) which performs well when applied to financial and company reports such as the 10-K reports.

Similar to Ardia et al. (2019), we consider for valence shifting words. Valence shifting words are words such as absolutely, hugely, very, barely, not that can affect the nearby words. An example of the effect of valence shifting words is "not genuine", under the default dictionary approach, the word "genuine" is a positive with a score of 1, but due to the negator valence shifting word that precedes it is negative and has a score of -1. After accounting for all the positive and negative words of each news article and valence shifting words, the sentiment is calculated on a sentence level in every paragraph in each article. The net sentiment for an entire document is therefore constructed as the proportion of words in the article that are positive minus the proportion of words that are negative:

$$s_{n,t} = \frac{N^{+}_{n,t} - N^{-}_{n,t}}{N^{+}_{n,t} + N^{-}_{n,t} + N^{0}_{n,t}}, \ \in [-1, 1] \,, \tag{4.2.3}$$

Where $N^{+}_{n,t}$ is the number of positive words for article n on day t , $N^{-}_{n,t}$ is the number of negative words, $N^{0}_{n,t}$ is the neutral words. The weekly sentiment per topic K is constructed by aggregating the sentiment of all the articles published during that week.[11] The higher the value of net sentiment indicates more optimistic news for topic K. The news sentiment per each news article takes a value between -1 and 1. Figure A.23 in the appendix shows an example of an article from the economic news topic that was published in the Associated Press in March 2014. This article has the lowest sentiment as it discusses the situation of the labour market in the U.S. following the great recession.

In figure 4.6 below I show the evolution the sentiment on a weekly basis for the entire news corpus from January to December 2015. In the appendix, I display figures A.24, A.25 and A.26 which illustrate the sentiment evolution for each news topic per week from January 1990 to December 2015. The graph shows that the sentiment for each news

[10] Details about the Harvard GI dictionary are available at the General Inquirer webpage.

[11] The recursive approach that is applied to topic modelling is that for every week, all the news articles are collapsed into one document then the topic modelling algorithm finds the top words associated with each topic. Hence, this does not differentiate between the effects of some keywords that are not considered in the Harvard GI dictionary. Therefore, one would not be able to see the effect of delaying any of the keywords.

topic captures the U.S. business cycle and other economic and political developments. For example, economic performance and the stock markets topics have their greatest dips during the time of U.S. recession, similarly applies to the housing market topic. In addition to this, the Political Economy topic have fluctuations during the times of tight U.S. elections, for example it had a dip during the 2000 Bush elections and 2008 Obama elections. Similarly, figures A.28, A.29 and A.30 show the application of the Loughran and McDonald (2011) dictionary on the news corpus, this dictionary consists of 2709 words, with 368 positive words and 2340 negative words. The aggregate news sentiment of the GI dictionary and the Loughran and McDonald (2011) have very similar evolution overtime and the Pearson correlation coefficient is 0.6.

Figure A.32 shows heat map of the correlation matrix between individual news sentiment indices. As illustrated, the highest correlation coefficient is between the company news topic and the commodities news topic such that it the Pearson correlation coefficient is 0.157. The correlation between the topics tend not to be strong and this can be explained by the choice of the optimal number of topics according to the Cao et al. (2009) approach that reduces the correlation between individual topics.

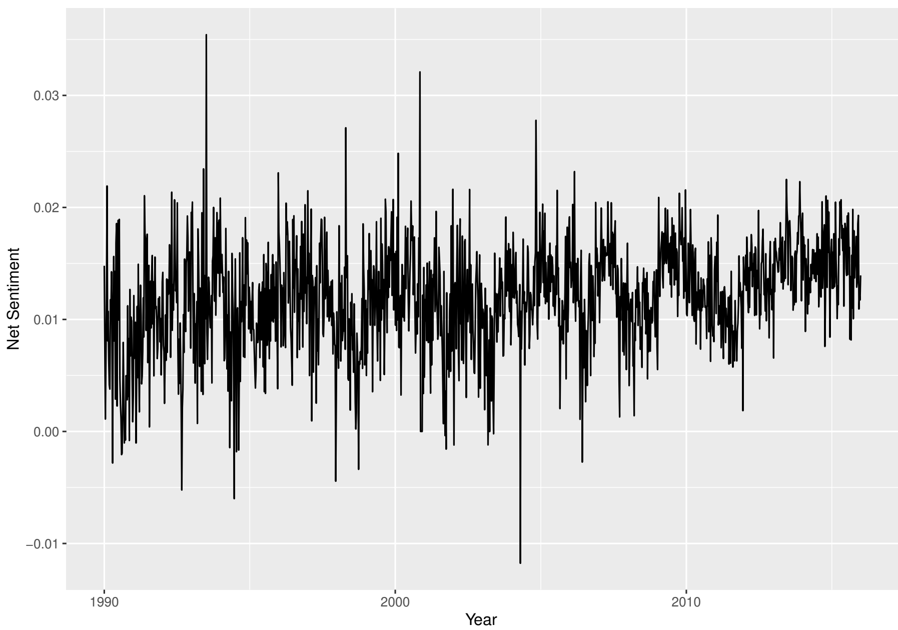

Figure 4.6: News corpus net sentiment.

Net sentiment for all the news corpus that combines 17 news topics on a weekly basis from 03/01/1990 to 28/12/2015. The net sentiment is calculated as $s_{n,t} = \frac{N^+_{n,t} - N^-_{n,t}}{N^+_{n,t} + N^-_{n,t} + N^0_{n,t}}$.

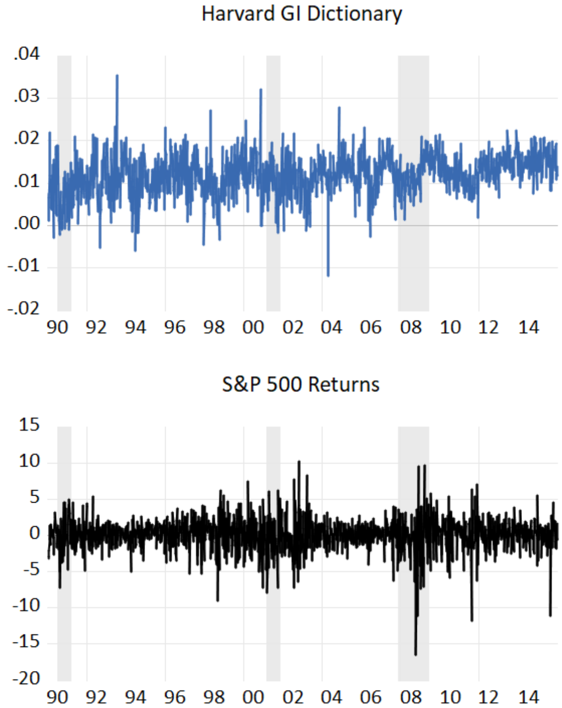

Figure 4.7: Harvard GI Dictionary and S&P 500 returns.

Harvard GI Dictionary versus S&P 500 returns.

The negative (positive) news sentiment implies an increase in the pessimism (optimism) about the market expectations that can affect the investor's trading behaviour. In the following section, I will investigate which topical news sentiment index can enhance the predictions of the volatility of stock returns. Figure 4.7 shows that the Harvard GI dictionary can approximate the S&P 500 index returns. Both figures tends to have decrease during time of recessions as highlighted by the US recessionary bars. The degree of correlation between both indices is 0.30.

The news sentiment is calculated using the "Sentometrics" package on R. In this package under both dictionaries, the positive words have a sentiment score of 1, while a negative words has a sentiment score of -1. This makes negative words and positive words have same impact. However, the inclusion of valence shifting words, which have different scores based on their intensity, the sentiment analysis becomes more nuanced and acknowledges that certain words can have a stronger or weaker impact on the overall sentiment.[12] It allows for the consideration of the intensity or strength of the sentiment

[12] The valence shifting words have scores either 1.8, 0.2, or -1 based on their intensity. For instance, the word totally has a score of 1.8 and the word barely has a score of 0.2 but the word can't has a score of -1 as it indicates negative valence shifting word.

expressed by each word. An alternative dictionary that is used is the Baccianella et al. (2010) SentiWordNet dictionary that is made of 8898 positive words and 11,029 negative words. This dictionary is a weighted lexicons, where words are weighted according to their degree of positiveness or negativeness. For example, the word accident has a sentiment score of -0.437, while the word acquire has a sentiment score of 0.250. Figure A.31 shows the net sentiment of this dictionary overtime. The SentiWordNet dictionary has a correlation of 0.54 with the Harvard GI dictionary.

One alternative for the textual analysis is the Term Frequency Inverse Document Frequency (Tfidf) approach. This technique measures the importance of a term in a document within a corpus by considering both the term frequency (TF) and the inverse document frequency (IDF). However using this approach is not feasible given the use of Sentometrics R package which only performs the weighting using IDF, identify terms that are relatively rare or unique to certain documents and therefore carry more meaningful information.

4.3 Model:

The time series of the stock prices for the S&P 500 composite price index is obtained from Bloomberg for the period 1990 to 2015 on a weekly basis for a total 1357 observations.The weekly price returns r_t are computed as the continuously compounded return on the asset over the period $t-1$ and t.

$$r_t = \log\left(\frac{P_t}{P_{t-1}}\right) \times 100 \tag{4.3.1}$$

P_t and P_{t-1} are the end of week prices for the current and the previous week respectively. Figure 4.8 displays the weekly prices and returns from 1990 to 2015, as displayed, the stock prices and returns decrease during the time of the recessions.

Engle (1982) proposed the Autoregressive Conditional Heteroscedasticity ARCH model in order to analyse the volatility of financial time series. In financial time series the variance of the errors is not constant, thus this invalidates the homoscedasticity assumption of the linear regression model. The plots in figure 4.8 are useful visual diagnostics for de-

(a) End of week S&P 500 index prices.

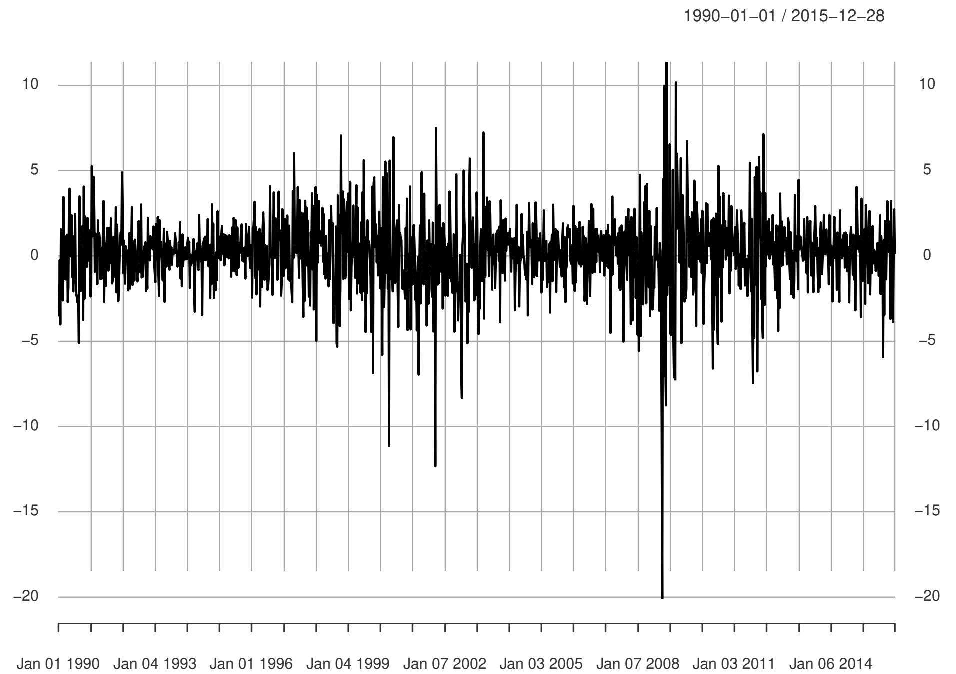

(b) End of week S&P 500 index price returns.

Figure 4.8: S&P 500 prices and returns.

Panel (b) shows the weekly prices are based on the end of week daily prices, returns are calculated based on equation 4.3.1.

tecting ARCH. Therefore it is important to consider using a model that does not assume a constant variance. If the conditional mean equation of the stock return r_t is:

$$r_t = \mu + u_t. \tag{4.3.2}$$

The residual term at time t, $r_t - \mu$ can be defined as

$$u_t = \sqrt{h_t}\eta_t \quad \eta_t \sim i.i.d.N(0,1). \tag{4.3.3}$$

A simple ARCH (p) model can then be written as:

$$h_t = \alpha_0 + \sum_{i=1}^{p} \alpha_i u_{t-i}^2. \tag{4.3.4}$$

Where η_t is iid normal innovation with a mean of zero and a variance of one, u_t is normally distributed with a mean of zero and a variance σ^2. As h_t is the conditional variance, it must be positive to ensure this, all the coefficients for an ARCH(p) model should be positive. That is, where $\alpha_0 > 0$ and $\alpha_i \geq 0 \, \forall \, i = 1, 2, ..., p$.

Bollerslev (1986) extended the ARCH model by introducing the Generalized Autoregressive Conditional Heteroscedasticity (GARCH) model such that the conditional variance is dependent on q of its previous lags in addition to p lags of squared residuals from the mean equation 4.3.2.

$$h_t = \alpha_0 + \sum_{i=1}^{p} \alpha_i u_{t-i}^2 + \sum_{j=1}^{q} \beta_j h_{t-j}. \tag{4.3.5}$$

The non-negativity constraints ensure that $\alpha_0 \geq 0$, $\alpha_i \geq 0 \;\; \forall i = 1, 2, ..., p$. $\beta_j \geq 0 \;\; \forall j = 1, 2, ..., q$ and $\alpha_1 + \beta_1 \leq 1$. Similar to the ARCH(p) model the kurtosis of the GARCH model tends to be far less than the sample kurtosis observed for most financial return series. Bollerslev (1986) finds evidence of conditional leptokurtosis in monthly S&P 500 index returns, i.e. larger than that of the normal distribution.

Although the GARCH(p,q) model is successful in capturing volatility clustering, it does not account for the leverage effect of shocks, as in financial time series models a negative shock is likely to cause volatility to rise by more than a positive shock of a similar magnitude, however by squaring the lagged residuals the magnitude of the signs is lost. I will use the GJR-GARCH of Glosten et al. (1993) and Zakoian (1994) as an extension to the basic GARCH model. This extension includes a dummy variable I_t that permits a response of volatility to positive and negative shocks, the dummy variable I_t takes a value of 1 if there is a negative shock and zero otherwise. Asymmetric leverage effect is present if $\gamma_k > 0 \, \forall \, k = 1, 2, ...r$. Therefore, in this model good news and bad news have different impacts on conditional volatility, the good news has an impact of α_i, while bad news has a greater impact i.e. the impact of $(\alpha_i + \gamma_k)$. That is to say, the GJR-GARCH model can be written as:

$$h_t = \alpha_0 + \sum_{i=1}^{p} \alpha_i u_{t-i}^2 + \sum_{j=1}^{q} \beta_j h_{t-j} + \sum_{k=1}^{r} \gamma_k u_{t-k}^2 I_{u_{t-k}<0}, \tag{4.3.6}$$

$$I_{t-k} = \begin{cases} 1, & \text{if } u_{t-k} < 0 \\ 0, & \text{otherwise} \end{cases} \tag{4.3.7}$$

The aim of this paper is to investigate the ability of topical news sentiment indices that are constructed and explained in section 4.2 to predict the volatility of stock returns. In other words, I ask: does the tone of the news about different news topics impact the volatility of stock returns? To answer this question, I introduce one lag of the news sentiment variable as an exogenous variable to the GARCH and GJR-GARCH of equations 4.3.5 and 4.3.6. That is to say the GARCH-X(p,q) and the GJR-GARCH-X(p,q,r) can be written as:

$$h_t = \alpha_0 + \sum_{i=1}^{p} \alpha_i u_{t-i}^2 + \sum_{j=1}^{q} \beta_j h_{t-j} + \psi X_{t-1}, \tag{4.3.8}$$

$$h_t = \alpha_0 + \sum_{i=1}^{p} \alpha_i u_{t-i}^2 + \sum_{j=1}^{q} \beta_j h_{t-j} + \sum_{k=1}^{r} \gamma u_{t-k}^2 I_{u_{t-1}<0} + \psi X_{t-1}, \tag{4.3.9}$$

where X is an exogenous variable which shows the net sentiment for each topic. My analysis is mainly driven by the findings of the literature that trading in the stock markets can be affected by different types of news such as macroeconomic, geo-political, company-specific news, local level news,...etc. News events across different news categories can offer an economic explanation for the stock returns volatility. Peng and Xiong (2006) and Ben-Rephael et al. (2017) find that the reaction of the institutional investors and the retail investors differ in regard to major news events, which can affect the magnitude of the volatility. Investigating which source of news events can influence the stock returns has not been examined in the literature of volatility predictions before. In order to evaluate the ability of the news sentiment in predicting the volatility of the stock returns, I extend the GARCH model to include the news sentiment of different news categories as exogenous regressors. As shown in Antweiler and Frank (2004) and Tetlock (2007) that texts from the internet and newspaper news can predict the stock prices. Including the news sentiment as an exogenous variable which can provide an explanation for the volatility of the stock returns. The literature of news sentiment and GARCH modelling found a negative relationship between the tones of the news and the volatility of returns. Smales (2015) shows that there is a negative relationship between news sentiment and volatility. Illustrated in the sign of the coefficient of news sentiment in the GARCH model, his findings illustrates that positive news causes a decrease in the volatility in the gold futures market. Similarly, Hsu et al. (2021) confirm this negative relationship by showing that negative news tends to increase the volatility of the returns from the Taiwanese stock market. In a similar context, Puzanova and Eratalay (2021), find that an increase in the good news sentiment of the housing market causes the volatility of returns of Scandinavian banks to decrease.

The choice of lags in the above mentioned GARCH models is based on the Akaike Information Criterion (AIC) and the Schwarz Information Criterion (SIC).

$$\mathrm{AIC} = 2P - 2\ln(L), \tag{4.3.10}$$

$$\mathrm{SIC} = \ln(N)P - 2\ln(L), \tag{4.3.11}$$

where P is the number of estimated parameters in the model, L is the maximum value of the likelihood function for the estimated model and N is the number of observations.

The empirical analysis is making use of the GARCH and GJR-GARCH model in particular as of the GARCH family for the following reasons. I start by looking at the GARCH model which is criticised for its simplicity in not capturing the asymmetric response of the volatility to negative and positive shocks. I include the news sentiment as exogenous variable to understand if in a simple GARCH model can the volatility be explained by the exogenous variable. I then utilise the GJR-GARCH model which is able to overcome the simplicity of the GARCH model in my analysis such that:

- The GJR-GARCH model incorporates a parameter that captures the notion of different volatility dynamics in different market regimes. In another words, during a period of economic downturn, the coefficient γ of the dummy variable $I_{u_{t-k}}$ will have a non-zero value, capturing the effects of the negative shock. It is plausible to assume that there may be distinct periods of high and low volatility regimes depending on the prevailing market sentiment. The dummy variable $I_{u_{t-k}}$ in the GJR-GARCH model allows for capturing such regime shifts, enabling a more nuanced analysis of volatility responses to news sentiment. Which is a phenomenon that can not be applied with other models of the GARCH family such as a basic GARCH model or IGARCH model which focuses primarily on modeling the persistence of volatility over time rather than capturing the volatility asymmetry.

- Including news sentiment as an exogenous variable in the GJR-GARCH model allows for examining the impact of news sentiment on the stock market volatility across different times of volatility regimes, potentially providing insights into how market participants react to news events.

4.4 Results:

In this section I evaluate if a volatility model incorporating the topical news sentiment indices can capture the investor reaction and thus explain the stock returns volatility in terms of in-sample and out-sample periods. I now change the measure of news sentiments such that the weekly news sentiment indices are measured based on the end of the week observations to match the way the weekly stock returns are measured. The in-sample period is from 08/01/1990 to 30/12/2013, this includes 1251 data points, while the out-of-sample period is from 06/01/2014 to 28/12/2015, including 105 data points.

As a starting point to the analysis, table 4.2 below shows the summary statistic of the returns of the S&P 500 index for the entire sample period from 08/01/1990 to 30/12/2015. As displayed in the table, the weekly average return is 13.03 percent, the weekly standard deviation is 2.303 and the weekly variance is 5.303. The coefficient of Kurtosis, which measures the thickness of the tails of the distribution is 9.739 higher than that of the normal distribution of 3, rejecting the Gaussianity assumption i.e. implying that the stock returns do not follow a normal distribution. Similarly, the Jarque Bera test for normality of the returns distributions yield a statistic of 2687.565 which rejects the null hypothesis of normally distributed returns at any critical value. Engle and Patton (2007) discuss that among the stylized facts of financial returns, the unconditional distribution of asset returns has heavy tails. Figure 4.9 below displays a Q-Q plot that compares the distribution of the S&P 500 returns to that of the normal distribution. As the data is not symmetrically lying on the curve, this confirms that the S&P 500 weekly price returns has more extreme values than the normal distribution. A visualisation of the autocorrelation functions of the returns and the squared returns is in figure 4.10. As displayed in the panel of the returns, there is weak dependence in the mean of the series. For the squared returns however, the plot displays dependence in the volatility of returns.

	r_t
Mean	0.130
Median	0.241
Maximum	11.356
Minimum	−20.836
Std. Dev.	2.303
Skewness	−0.733
Kurtosis	9.739
Jarque-Bera	2687.565
Observations	1356

Table 4.2: S&P 500 index weekly returns summary statistics.

The entire sample period is from 08/01/1990 to 28/12/2015, a total of 1356 observations.

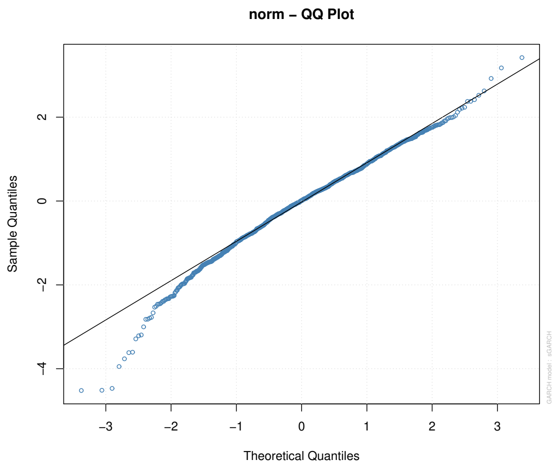

Figure 4.9: Q-Q plot for the S&P 500 returns.

Quantiles of the S&P 500 returns against that of the normal distribution.

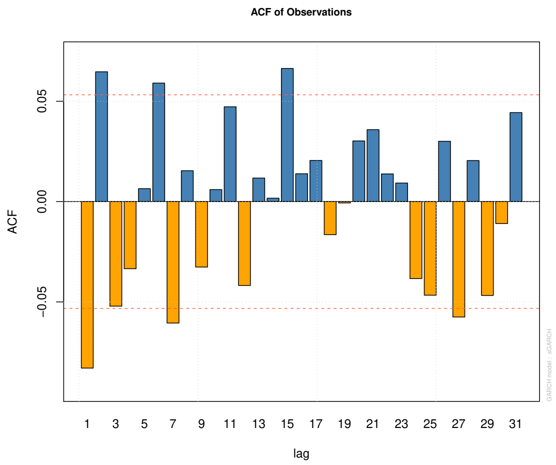

(a) Autocorrelation Function for returns.

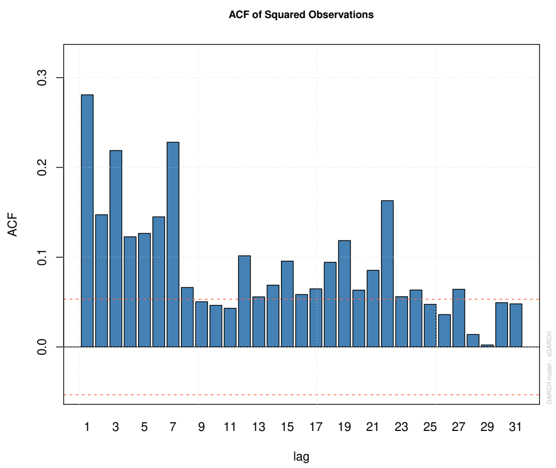

(b) Autocorrelation Function for squared returns.

Figure 4.10: Autocorrelation functions for S&P 500 returns and squared returns.

4.4.1 GARCH Model:

The first step in performing the volatility modelling is to identify the appropriate model that fits the conditional volatility best. Table 4.3 shows that the optimal choice that yields the lowest Akaike Information Criterion and Schwarz Information Criterion values 1 lag for the squared residuals and 1 lag for the conditional variance. Hence, I will follow the information criterion and apply 1 lag for the GARCH and GJR-GARCH models.

	AIC		SIC	
p	q=1	q=2	q=1	q=2
1	4.247	4.247	4.262	4.266
2	4.2469	4.248	4.266	4.271
3	4.248	4.249	4.271	4.276
4	4.240	4.240	4.267	4.276

Table 4.3: Optimal lag selection.

AIC is the Akaike Information Criterion, SIC is the Schwarz Information Criterion. The choice of 1 lag for squared residuals and conditional variance is the optimal choice as it reduces the information criterion values for both the AIC and SIC.

The estimations of the GARCH(1,1) and the GARCH(1,1)-X models are displayed in table 4.7 and the robust Bollerslev-Wooldridge standard errors are in the parentheses. For the GARCH(1,1) the mean and variance equations are illustrated in the second column. As illustrated in the table, the volatility persistence $(\alpha_1 + \beta_1)$ is 0.978. I then test the null hypothesis of Integrated GARCH that $\alpha_1 + \beta_1 = 1$ in other words, that the model has a unit root in variance.[13]

Test Statistic	Value	df	Probability
t-statistic	-7.801	1247	0.000
F-statistic	60.848	(1, 1247)	0.000
Chi-square	60.848	1	0.000

Table 4.4: Testing of Integrated GARCH (IGARCH).

The Wald test of null hypothesis is that $\alpha_1 + \beta_1 = 1$ i.e. unit root in variance.

As the primary focus of the paper is to consider the conditional relationship between the volatility of stock returns and the topical news sentiments. Panel (a) of figure 4.11 shows the square root of the conditional volatility $\sqrt{h_t}$ of GARCH(1,1) model for the in-sample period. Columns 4 to 20 of table 4.7 shows the estimation of the conditional

[13]Table 4.4 shows the estimation of the Wald test. As displayed in the table, I reject the null hypothesis in favour of the alternative hypothesis that the conditional variance is not integrated.

volatility with topical news sentiments as exogenous variables. The coefficients of the ARCH term u_{t-1}^2 and the coefficients of the GARCH term h_{t-1} are highly significant indicating persistence in the volatility levels for all the specifications. That is, high volatility is followed by high volatility and low volatility is followed by low volatility. To evaluate the ability of topical news sentiment in explaining the conditional volatility, the coefficients ψ of the topics Monetary Policy, Trade, Trade Policy, Local Economy and Manufacturing are negative and highly significant, indicating that the conditional volatility decreases as tone of the news becomes more positive. However, for the topics Commodities, Currencies and Politics, the coefficients are positive and significant. This can be explained as positive news for the commodities market are associated with an increase in the commodity prices which can have an effect on the volatility of stock returns. Similarly, positive news about appreciation in the currency market can increase the volatility of stock returns in the stock market.

4.4.2 GJR-GARCH Model:

To model the leverage effect, I estimate the following GJR-GARCH(1,1) model:

$$h_t = \alpha_0 + \alpha_1 u_{t-1}^2 + \beta_1 h_{t-1} + \gamma u_{t-1}^2 I_{u_{t-1}<0}. \tag{4.4.1}$$

Panel (b) of figure 4.11 shows the square root of the conditional volatility $\sqrt{h_t}$ of GJR-GARCH(1,1) model for the in-sample period. The graph shows that during the times of economic recessions, the volatility of the GJR-GARCH(1,1) is higher than that of the GARCH(1,1) model, confirming the leverage effect phenomenon that negative news shocks have a greater impact on volatility than positive ones.

As illustrated in table 4.8 , the persistence of the volatility drops when estimating the GJR-GARCH(1,1) model. The volatility persistence $(\alpha_1 + \beta_1)$ is 0.837. The coefficient of the asymmetric leverage effect $u_{t-1}^2 I_{u_{t-1}<0}$ is positive and highly significant indicating asymmetry in the conditional volatility arising from the innovations in the stock returns. I now introduce the topical news sentiment as exogenous variables for the GJR-GARCH to investigate the forward looking predictability of the news sentiment, such that the following GJR-GARCH(1,1)-X is estimated:

$$h_t = \alpha_0 + \alpha_1 u_{t-1}^2 + \beta_1 h_{t-1} + \gamma u_{t-1}^2 I_{u_{t-1}<0} + \psi X_{t-1}. \tag{4.4.2}$$

As explained in section 4.3, X_{t-1} is the one lag of the net sentiment for each topic. The output of the estimation of this equation is in columns 4 to 20 in table 4.8 below. The coefficients of the exogenous variables X_{t-1} are negative and statistically significant for the topics Monetary Policy, Trade, Trade Policy, Geo Politics, Currency and Manufacturing. This confirms the results of the GARCH-X(1,1) model that these topics can explain the conditional volatility for the S&P 500 weekly returns. These findings are similar to that of Adämmer and Schüssler (2020) they find that out of 100 topics, the volume of news topics related to politics and war only can forecast the stock returns.[14]

Finding a negative relationship between the news sentiment and the volatility of stock returns was documented in the literature.[15] To show the contribution of each variable to the conditional volatility, figure 4.12 displays the estimated conditional volatilites $\sqrt{h_t}$ for the models with statistically significant coefficients for the news sentiment exogenous variables plotted along with the conditional volatility of the GJR-GARCH(1,1) model. As displayed, there exists variation in the conditional variance for each the GJR-GARCH(1,1)-X model. For instance, during the time of the financial crisis, highlighted by the third recessionary bar, the conditional volatility of the monetary policy was the highest, reaching 10.7. This variation in the conditional volatility encourage the construction of a series that combines the forecasts for each of these individual models. It is important to stress that the positivity of the conditional variance for all the models can be visualised for the GJR-GARCH(1,1)-X models.

I then divide my estimation sample into two sub-sample periods and estimate GARCH and GJR-GARCH models for each period to evaluate the ability of the topical news sentiment indices in explaining the conditional volatility during these periods. The first sub-sample is the recession period and the second sub-sample is the post-recession (expansion) period. I follow the NBER recession dates when dividing my sample into these

[14] Manela and Moreira (2017) find that news related to war and government can predict the stock returns.

[15] Smales (2015), Hsu et al. (2021) find a negative relationship between net news sentiment and conditional volatility. Similarly, Puzanova and Eratalay (2021) find a negative relationship between good news and stock return conditional volatility.

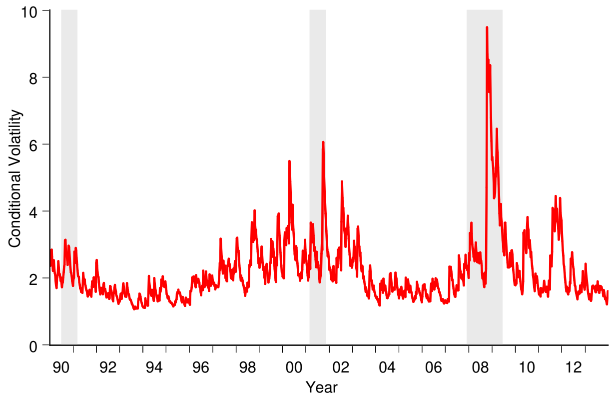

(a) GARCH(1,1) Conditional volatility.

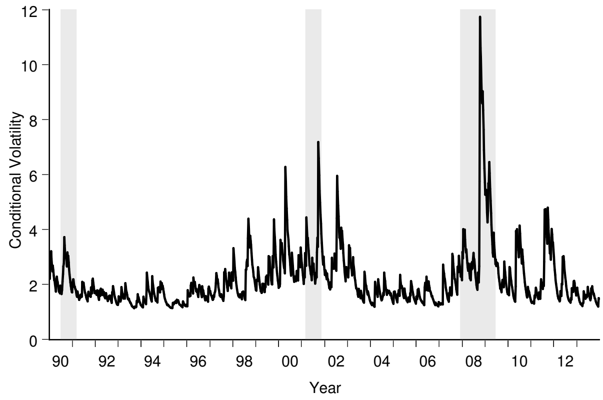

(b) GJR-GARCH(1,1) Conditional volatility

Figure 4.11: Estimated conditional volatility.

$\sqrt{h_t}$ for GARCH(1,1), GJR-GARCH(1,1) The shaded bars represent the NBER recession periods.

periods. The NBER defines a recessionary period as a period that starts at the peak of a business cycle and ends at the trough. This means that the recession started in December 2007 and ended in June 2009. The post-recession period, covering the period of July 2009 to the end of the sample. Smales (2015) and Hsu et al. (2021) find that negative news has a strong effect on the volatility of stock returns during the recessionary period.

As displayed in the appendix, tables A.35 and A.36 show the output of the GARCH(1,1)-X in the recession and post-recession periods respectively. While, tables A.37 and A.38 show the output of the GJR-GARCH(1,1)-X in the recession and the post-recession periods. The topics Monetary Policy, Currency show a significant asymmetric effect on the volatility of the stock returns in the GARCH(1,1)-X and the GJR-GARCH(1,1)-X during the recession and post-recession period. The magnitude of this effect depends on the state of the economy. However, there are some indices that show significant effect on the volatility depending on the period such as the Commodities, Construction and Tourism topical indices which have a significant effect in the post-recession periods.

4.4.3 Out-of-sample Predictability

To evaluate the ability of the topical news sentiment indices in enhancing the volatility prediction, I perform an exercise to predict the one step and the multi step ahead stock returns conditional volatility. Such that the out-of-sample period covers 104 observations

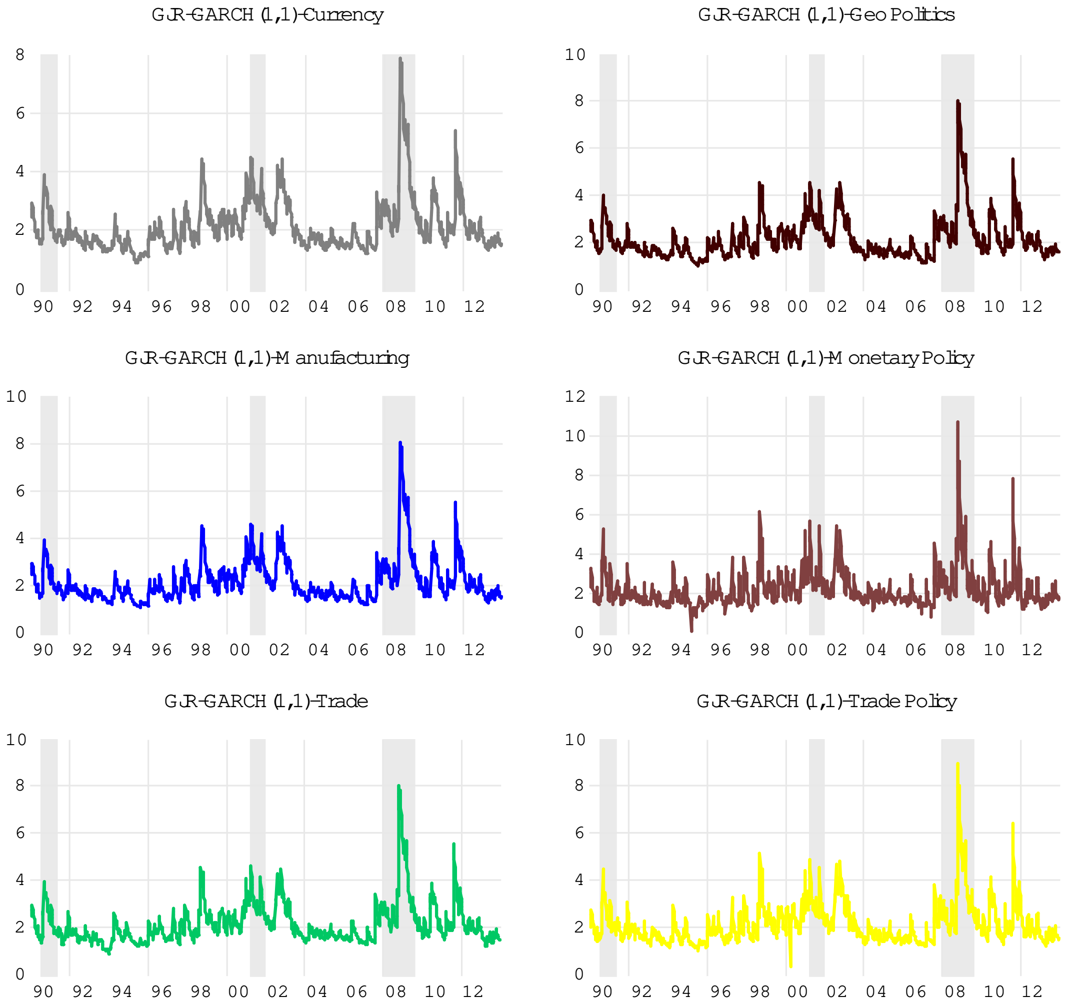

Figure 4.12: Estimated Conditional Volatility for the GJR-GARCH-X(1,1) model.

$\sqrt{h_t}$ for GARCH(1,1), GJR-GARCH(1,1) The shaded bars represent the NBER recession periods.

and runs from 06/01/2014 to 28/12/2015 and obtain the fitted volatility series for the S&P 500 index for each specification.

I perform a one step ahead out of sample forecasts (h=1), then I perform multi step ahead forecasts (h>1). The one step ahead forecasts are computed in a recursive manner, this involves forecasting one week ahead, then refitting the model after every time step based on the new information that feeds into the forecasts. For the out-of-sample period, I generate the forecasts for the one step ahead conditional variance h_{t+1} for the GARCH(1,1) model as:

$$\mathbb{E}_t\left[h_{t+1}\right] = \alpha_0 + \alpha_1 u_t^2 + \beta_1 h_t. \tag{4.4.3}$$

Similarly, in the GJR-GARCH(1,1) the one step ahead forecasts can be written as:

$$\mathbb{E}_t\left[h_{t+1}\right] = \alpha_0 + \alpha_1 u_t^2 + \beta_1 h_t + \gamma u_t^2 I_{u_t<0}. \tag{4.4.4}$$

The one step ahead forecasts for the conditional volatility obtained from the GARCH(1,1) and the GJR-GARCH(1,1) models are displayed in figure 4.13. In figure A.33 I display the plot of the one step ahead forecasts of the conditional volatility for the GARCH(1,1)-X and GJR-GARCH(1,1)-X for the models that had statistically significant exogenous variables in the in-sample period as explained in the previous section. Panel (a) of figure A.33 displays the one step ahead forecasts of the GARCH(1,1)-X model, while panel (b) displays the one step ahead forecasts of the GJR-GARCH(1,1)-X model. The forecasting exercise performed in this paper is performed using EViews static forecasts which performs one step ahead forecasts in a recursive estimation approach. The forecasting exercise is only performed to one step ahead due to the limitations available with EViews program which is unable to perform forecasts with two step ahead forecasts or longer.[16]

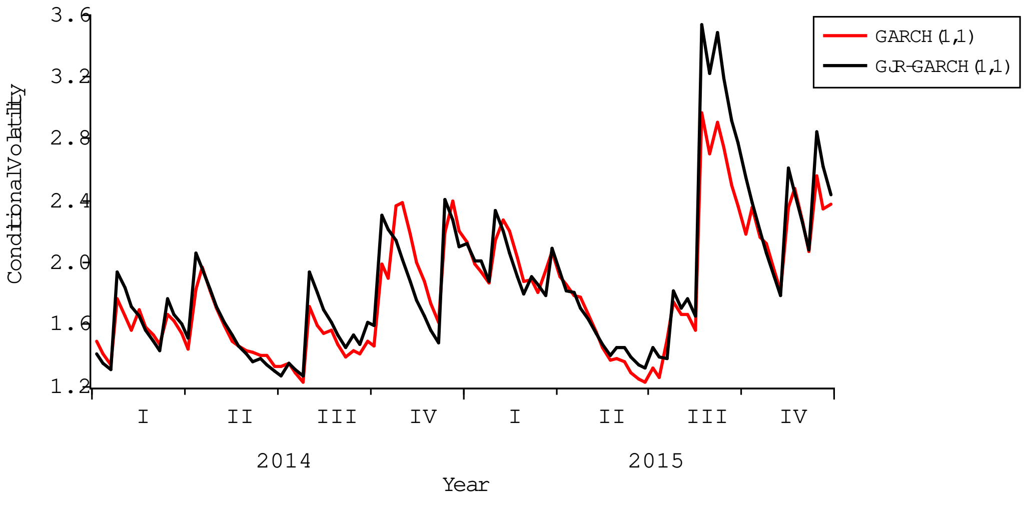

Figure 4.13: One step ahead forecasts of the conditional volatility.

Conditional volatility is plotted as $\sqrt{h_t}$ for the GARCH(1,1) and the GJR-GARCH(1,1). The out-of-sample period is from 06/01/2014 to 28/12/2015.

Bates and Granger (1969) were the first to show that a combination of individual fore-

[16]EViews has the dynamic forecast procedure which forecasts with $h > 1$. However, with this procedure, performing forecasts evaluation and comparing different methods across different horizons could be problematic. That is, the root mean squared error for multi-step ahead ($h > 1$) forecasts is computed with a smaller sample size than that of the one step ahead ($h = 1$) as we approach the end of the data, potentially penalizing long-range forecasts more severely at the beginning of the forecasting sample.

casts can out-perform the individual forecasts themselves. I aggregate the out-of-sample forecasts obtained from the GJR-GARCH(1,1)-X models using a forecast combination approach. Such that I have a set of combined forecasts for the one step ahead forecasts and another combined forecast for the multi step ahead forecasts.[17] The combination of forecasts of $\hat{h}_{c,t+h}$ are weighted averages of the individual N forecasts:

$$h_{c,t+h} = \sum_{i=1}^{N} \omega_i \hat{h}_{i,t+h}, \tag{4.4.5}$$

where ω_i is the weight assigned to each forecast and $h_{i,t+h}$ is the forecast for each individual topic. In this case, I combine forecasts of five topics, $N = 5$. There are several ways to obtain weights for ω_i. They range from simple averaging schemes such as mean, median and trimmed mean to sophisticated methods such as a regression based approach.

Using regression to combine forecasts is based on matching the target series that I need to forecast with the conditional volatilities from individual models. The target series is the realised volatility proxy, I use the squared residuals u_t^2 as a proxy for volatility.[18] That is to say, I run the following regression using the in-sample period observations to estimate weights:

$$\hat{u}_t^2 = \alpha + \sum_{i=1}^{N} w_i \hat{h}_{i,t} + \epsilon_t.$$

According to Rapach et al. (2010), the weights that are formed at time t are functions of the historical forecasting performance of the individual models over the training set, which is the in-sample period. While in the test sample, which is the out-of-sample period, I evaluate the model performance. The weights assigned to each of the individual forecasts ω_i can be estimated using OLS to minimise the sum of squared errors. The combined series can be written as

$$h_{c,t} = \hat{\alpha} + \sum_{i=1}^{N} \hat{\omega}_i \hat{h}_{i,t},$$

Where $\hat{\alpha}$ and $\hat{\omega}_i$ are the OLS estimates for α and ω_i respectively. The one-step ahead

[17] I use the R package ForecastComb of Weiss et al. (2018) to combine the forecasts.

[18] The squared residuals are calculated as: $u_t^2 = (r_t - \mu)^2$. Awartani and Corradi (2005) also used the squared residuals as a proxy for volatility.

combined forecasts, h=1:

$$h_{c,t+1} = \hat{\alpha} + \sum_{i=1}^{N} \hat{\omega}_i \hat{h}_{i,t+1}.$$

Nowotarski et al. (2014) suggest using the Least Absolute Deviation (LAD) regression in the presence of a correlated individual component forecasts. In table 4.5 below I evaluate the accuracy of the combined forecasts across different techniques based on the Mean Absolute Error (MAE), Mean Percentage Error (MPE) and Mean Absolute Percentage Error (MAPE). The choice of using LAD yields the lowest errors, thus I will use the LAD technique to combine the individual five forecasts.

	MAE	MPE	MAPE
Regression based methods:			
OLS	3.787	-33307.20	33330.72
CLS	3.595	-31996.66	32020.69
LAD	2.748	-13849.54	13899.11

Table 4.5: Evaluation of forecast combination techniques.

OLS is the estimations using Ordinary Least Squares regression. CLS is the estimations using Constrained Least Squares regression. LAD is the estimations using Least Absolute Deviation regression. Evaluation is made during in the in-sample period.

Forecast Evaluation

In this section I evaluate if the combined forecasts series described in the previous section can provide gains in forecasting the volatility of the stock returns. That is, I evaluate the ability of the combined forecast $\hat{h}_{c,t+h}$ for the one step ahead (h =1) in predicting the realised volatility. Although Patton and Sheppard (2009) and Sheppard (2019) show that the choice of squared returns, r_t^2, is a reasonable proxy of realised volatility when working with high frequency applications. I choose the squared residuals, $\hat{u}_t^2$ estimated from a model for the conditional mean, as a proxy for realised volatility as I am working with a model with longer horizon. I use the Diebold and Mariano (2002) test of equal predictive accuracy to perform comparisons across two different models at a time. I evaluate whether the forecasts from the benchmark model GJR-GARCH(1,1) without any topical news sentiment indices as exogenous variables can be out-performed by the combined forecasts which combines the five topics Monetary Policy, Trade, Trade Policy, Manufacturing and Currency. The DM tests the null hypothesis of equal predictive accuracy against composite alternatives that indicate which forecasts perform better.

When using a proxy for realised volatility it may be affected by some noise which can affect the conclusion driven from the measures of the forecast accuracy used.[19] Several literature looked at the use of multiple loss functions for forecasts comparison when using a measure of volatility proxy see (Hansen and Lunde 2005). Therefore, a response to using the squared residuals as volatility proxy is to apply several loss functions that were highlighted in the literature to evaluate the volatility forecasts. In that case, I apply the nine loss functions that were used in Patton and Sheppard (2009) and Patton (2011). These loss functions are MSE, MSE-LOG, MSE-SD, MSE-prop, MAE, MAE-LOG, MAE-SD, MAE-prop and QML loss functions for forecasts comparison. Although the MSE and MAE are the main loss functions applied in most of financial econometrics literature. Utilising the nine loss functions can help in coming to a conclusion on the accuracy of the volatility forecasts.

$$H_0 : \mathrm{E}\left[\delta_t\right] = 0$$

$$H_1 : \mathrm{E}\left[\delta_t\right] \neq 0$$

Where the MSE loss function differential δ_t is defined as

$$\delta_t = \left(\hat{u}_{t+h}^2 - h_{A,t+h}\right)^2 - \left(\hat{u}_{t+h}^2 - h_{c,t+h}\right)^2 \tag{4.4.6}$$

The MSE-LOG loss function differential δ_t is defined as

$$\delta_t = \left(\log \hat{u}_{t+h}^2 - \log h_{A,t+h}\right)^2 - \left(\log \hat{u}_{t+h}^2 - \log h_{c,t+h}\right)^2 \tag{4.4.7}$$

The MSE-SD loss function differential δ_t is defined as

$$\delta_t = \left(\hat{u}_{t+h} - \sqrt{h_{A,t+h}}\right)^2 - \left(\log \hat{u}_{t+h} - \sqrt{h_{c,t+h}}\right)^2 \tag{4.4.8}$$

[19] Patton (2011) highlights that if the log stock price follows a Brownian motion then realised volatility based on intra-daily 5-min returns (78 intra-daily observations) could be unbiased and more efficient than the squared return or residual squared. However, using high-frequency data is not feasible in the news sentiment exercise as the news articles available are not timely dated to show whether they were published during trading hours or after.

The MSE-prop loss function differential δ_t is defined as

$$\delta_t = \left(\frac{\hat{u}^2_{t+h}}{h_{A,t+h}} - 1\right)^2 - \left(\frac{\hat{u}^2_{t+h}}{h_{c,t+h}} - 1\right)^2 \tag{4.4.9}$$

The MAE loss function differential δ_t is defined as

$$\delta_t = \left|\hat{u}^2_{t+h} - h_{A,t+h}\right| - \left|\hat{u}^2_{t+h} - h_{c,t+h}\right| \tag{4.4.10}$$

The MAE-LOG loss function differential δ:

$$\delta_t = \left|\log \hat{u}^2_{t+h} - \log h_{A,t+h}\right| - \left|\log \hat{u}^2_{t+h} - \log h_{c,t+h}\right| \tag{4.4.11}$$

The MAE-SD loss function differential δ_t is defined as

$$\delta_t = \left|\hat{u}_{t+h} - \sqrt{h_{A,t+h}}\right| - \left|\log \hat{u}_{t+h} - \sqrt{h_{c,t+h}}\right| \tag{4.4.12}$$

The MAE-prop loss function differential δ:

$$\delta_t = \left|\frac{\hat{u}_{t+h}}{h_{A,t+h}} - 1\right| - \left|\frac{\hat{u}_{t+h}}{h_{c,t+h}} - 1\right| \tag{4.4.13}$$

For the QML-loss function, the loss function differential δ_t is defined as:

$$\delta_t = \left(\log h_{A,t+h} + \frac{u^2_{t+h}}{h_{A,t+h}}\right) - \left(\log h_{c,t+h} + \frac{u^2_{t+h}}{h_{c,t+h}}\right) \tag{4.4.14}$$

where $h_{A,t+h}$ is the forecasts for model A which is the benchmark GJR-GARCH(1,1) model. While $h_{c,t+h}$ is the combined forecast illustrated in the previous section. Rejecting the null hypothesis of the equal forecasting ability according to the DM statistic and the p-value implies that one of the models out-performs the other.[20] To understand which model has a better forecasting ability, the sign of $\bar{\delta}$, which is the sample mean of the loss differential plays an important role. Where $\bar{\delta}$ can be defined as

$$\bar{\delta} = \frac{\sum_{r=1}^{R} \delta_r}{R}$$

[20] Computations of the DM test are made through the "forecast" package on R. This package computes the Harvey et al. (1997) test statistic, which is a bias correction to the DM statistic.

R is the number of out-of-sample observations used to compute the Diebold and Mariano statistic. Positive values for $\bar{\delta}$ indicate that the model with the combined forecasts of the GJR-GARCH(1,1)-X model outperforms the forecasts from model A which is the benchmark GJR-GARCH(1,1) model. Table 4.6 shows the results of the DM test for each loss function. According to the DM statistic and the p-values and the values for $\bar{\delta}$ of all the loss functions with exception of MSE and MAE, the combined forecasts of the one step ahead (h=1) forecasts outperform the forecasts from the GJR-GARCH(1,1) model.

Loss Function	
MSE:	
$\bar{\delta}$	0.242
DM statistic	0.678
P-value	0.809
MSE-LOG	
$\bar{\delta}$	5.035
DM statistic	1.993***
P-value	0.000
MSE SD:	
$\bar{\delta}$	2.322
DM statistic	3.727***
P-value	0.0003
MSE Prop:	
$\bar{\delta}$	75.367
DM statistic	1.690***
P-value	0.094
MAE:	
$\bar{\delta}$	0.766
DM statistic	−1.962*
P-value	0.052
MAE-LOG:	
$\bar{\delta}$	0.319
DM statistic	2.690***
P-value	0.008
MAE-SD:	
$\bar{\delta}$	0.230
DM statistic	3.148***
P-value	0.002
MAE-Prop:	
$\bar{\delta}$	2.282
DM statistic	1.694*
P-value	0.094
QML	
$\bar{\delta}$	0.436
DM statistic	1.813*
P-value	0.073

Table 4.6: Diebold and Mariano forecast evaluation test.

Test of equal prediction accuracy between the forecasts of GJR-GARCH(1,1) and the combined forecasts using the topical news sentiments. The loss functions are MSE, MSE-LOG, MSE-SD, MSE-Prop, MAE, MAE-LOG, MAE-SD, MAE-Prop and QML (QLIKE).

4.5 Conclusion

This paper develops an approach to decompose textual data in the U.S newspapers into different news topics in a real time manner using the machine learning approach Latent Dirichlet Allocation. The tonality of the news articles is then measured using the Harvard GI dictionary to construct a time series for each topic.

I investigate the ability of news sentiment across different topics to explain the stock returns volatility for the S&P 500 index using GARCH and GJR-GARCH models. I consider the effects of classified news such as macroeconomic news, company specific news, financial markets news and political and geo-political news in capturing the stock returns volatility. The findings of this paper illustrate that the volatility of stock returns can be explained by the sentiment of topics such as monetary policy, manufacturing, trade, trade policy and currency market. I show that there is a negative relationship between the news sentiments and the volatility of stock returns, confirming the findings of previous literature.

When I perform a forecast combination exercise using regression based method, I find that combining the out-of-sample forecasts of these five topics using regression based methods outperforms the forecasts of a simple GJR-GARCH(1,1) model according to the DM test.

To conclude, this paper shows that the news sentiment across different economic topics can serve as a gauge for public information, that can drive the attention of the risk averse investors, as they cover various areas of the economy, which Keynes referred to as "Animal Spirits".

Model	GARCH(1,1)-X																		
	GARCH(1,1)	All	Monetary Policy	Commodities	Housing	Local Economy	Economic Performance	Currency	Company	Fiscal Policy	Political Economy	Geo Political	Trade Policy	Trade	Construction	Manufacturing	Tourism	Stock Market	Foreign Economy
Mean equation																			
μ	0.242***	0.242***	0.240***	0.245***	0.238***	0.243***	0.240***	0.249***	0.164**	0.242***	0.243***	0.241***	0.247***	0.235***	0.242***	0.278***	0.245***	0.241***	0.242***
	(0.050)	(0.050)	(0.050)	(0.050)	(0.049)	(0.046)	(0.050)	(0.051)	(0.066)	(0.052)	(0.052)	(0.048)	(0.050)	(0.050)	(0.052)	(0.052)	(0.052)	(0.050)	(0.051)
Variance equation																			
α_0	0.178***	0.198**	0.214***	0.163***	0.150***	0.275***	0.171**	0.184***	3.859	0.179**	0.177***	0.162**	0.145***	0.182***	0.190***	2.184***	0.192**	0.192***	0.170***
	(0.045)	(0.091)	(0.055)	(0.045)	(0.042)	(0.071)	(0.057)	(0.050)	(13.166)	(0.059)	(0.058)	(0.049)	(0.048)	(0.051)	(0.043)	(0.546)	(0.058)	(0.064)	(0.064)
α_1	0.175***	0.175***	0.179***	0.175***	0.165***	0.175***	0.173***	0.170***	0.129	0.175***	0.175***	0.175***	0.119***	0.151***	0.177***	0.323***	0.169***	0.174***	0.178***
	(0.016)	(0.016)	(0.0176)	(0.0175)	(0.017)	(0.018)	(0.017)	(0.043)	(1.493)	(0.043)	(0.043)	(0.017)	(0.022)	(0.039)	(0.436)	(0.073)	(0.042)	(0.043)	(0.044)
β_1	0.801***	0.799***	0.790***	0.800***	0.810***	0.792***	0.803***	0.804***	0.468	0.801***	0.801***	0.801***	0.862***	0.825***	0.798***	0.302*	0.806***	0.800***	0.797***
	0.020)	(0.021)	(0.022)	(0.021)	(0.020)	(0.021)	(0.020)	(0.036)	(2.555)	(0.038)	(0.038)	(0.020)	(0.023)	(0.032)	(0.038)	(0.159)	(0.037)	(0.039)	(0.039)
ψ		−0.002	−0.018*	0.023*	0.013	−0.0010**	0.003	0.088***	−0.119***	−0.0001	0.0002	0.007	−0.021***	−0.021***	−0.004	−0.084***	−0.021*	−0.005	0.008
		(0.007)	(0.010)	(0.014)	(0.009)	(0.005)	(0.007)	(0.021)	(0.009)	(0.008)	(0.009)	(0.009)	(0.007)	(0.007)	(0.005)	(0.013)	(0.012)	(0.011)	(0.011)

Table 4.7: Estimation of the GARCH models.

I use the Berndt-Hall-Hall-Hausman (BHHH) optimisation algorithm to estimate the parameters of the LLF. Column (2) shows an estimation of GARCH(1,1) model. Column (3) to column (20) show the estimation of GARCH-X(1,1) with different news sentiment topics. The Bollerslev-Wooldridge standard errors (sandwich estimator) are reported in the parentheses. The significant levels are $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$ respectively. In sample period is on a weekly basis from 8$^{\text{th}}$ of January 1990 to 30$^{\text{th}}$ of December 2013. Coefficients in the colour green indicate the statistical significance of the topical news sentiment variables in the GARCH-X equations.

Model		GJR-GARCH(1,1)-X																	
	GJR-GARCH(1,1)	All	Monetary Policy	Commodities	Housing	Local Economy	Economic Performance	Currency	Company	Fiscal Policy	Political Economy	Geo Political	Trade	Trade Policy	Construction	Manufacturing	Tourism	Stock Market	Foreign Economy
Mean equation																			
μ	0.148**	0.147**	0.148**	0.151**	0.148**	0.151**	0.148**	0.158**	0.155**	0.148**	0.148**	0.115**	0.152***	0.190***	0.182***	0.153**	0.153**	0.148**	0.148**
	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)	(0.051)	(0.058)	(0.050)	(0.045)	(0.051)	(0.050)	(0.051)	(0.051)	(0.051)
Variance equation																			
α_0	0.214***	0.259**	0.233***	0.204***	0.207***	0.287***	0.212***	0.207***	3.531**	0.218***	0.212***	3.366***	0.214***	1.391***	0.192***	0.244***	0.026	0.220***	0.210***
	(0.052)	(0.094)	(0.054)	(0.047)	(0.030)	(0.066)	(0.057)	(0.046)	(1.663)	(0.055)	(0.053)	(0.649)	(0.037)	(0.253)	(0.054)	(0.053)	(0.030)	(0.056)	(0.058)
α_1	0.026	0.026***	0.027	0.029	0.026	0.026	0.026	0.022	0.134	0.027	0.026	0.021	0.022	0.166**	0.019	0.025	0.026	0.026	0.027
	(0.030)	(0.031)	(0.031)	(0.030)	(0.030)	(0.031)	(0.030)	(0.027)	(0.111)	(0.030)	(0.030)	(0.063)	(0.016)	(0.058)	(0.024)	(0.031)	(0.030)	(0.030)	(0.030)
γ	0.245***	0.249***	0.248***	0.243***	0.242***	0.247***	0.244***	0.228***	0.912	0.247***	0.245***	0.268***	0.220***	0.296**	0.209***	0.244***	0.243***	0.245***	0.245***
	(0.057)	(0.057)	(0.057)	(0.056)	(0.056)	(0.057)	(0.057)	(0.055)	(0.459)	(0.057)	(0.057)	(0.109)	(0.055)	(0.093)	(0.040)	(0.056)	(0.056)	(0.056)	(0.056)
β_1	0.811***	0.807***	0.805***	0.809***	0.812***	0.804***	0.811***	0.822***	0.454***	0.809***	0.810***	0.432***	0.830***	0.450***	0.834***	0.808***	0.810***	0.810***	0.809***
	(0.024)	(0.027)	(0.027)	(0.026)	(0.026)	(0.028)	(0.026)	(0.025)	(0.092)	(0.027)	(0.026)	(0.122)	(0.021)	(0.090)	(0.028)	(0.027)	(0.026)	(0.027)	(0.027)
ψ		−0.004	−0.008***	0.013	0.003	−0.008	0.001	0.078***	−0.109**	−0.002	0.002	−0.132***	−0.025***	−0.067***	−0.001	−0.027***	−0.017	−0.002	0.003
		(0.007)	(0.004)	(0.016)	(0.008)	(0.006)	(0.009)	(0.019)	(0.045)	(0.006)	(0.007)	(0.013)	(0.004)	(0.006)	(0.008)	(0.003)	(0.011)	(0.009)	(0.009)

Table 4.8: Estimation of the GJR-GARCH models.

I use the Berndt-Hall-Hall-Hausman (BHHH) optimisation algorithm to estimate the parameters of the LLF. Column (2) shows an estimation of GJR-GARCH(1,1) model. Column (3) to column (20) show the estimation of GJR-GARCH(1,1)-X with different news sentiment topics. The Bollerslev-Wooldridge standard errors (sandwich estimator) are reported in the parentheses. The significant levels are $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$ respectively. In sample period is on a weekly basis from 8th of January 1990 to 30th of December 2013. Coefficients in the colour green indicate the statistical significance of the topical news sentiment variables in the GJR-GARCH-X equations

Chapter 5

Summary and Conclusion

This thesis looked at the application of data science in the context of policy evaluation and stock market forecasting. The thesis applied the state-of-the-art causal inference method Synthetic Control to measure the impact of the Brexit vote and the Stamp Duty Land Tax (SDLT) holiday on the UK economy. The application of data science for policy evaluation was complemented with application of data science and machine learning techniques to construct a measure of different news topics that can be used to forecast the volatility of aggregate stock returns.

Chapters 2 and 3 look at applying the Synthetic Control method, such that I created a counterfactual series named "Synthetic UK" that mimicked the performance of the UK economy in the pre-treatment periods. The impact of the treatment is calculated as the difference between the UK and its synthetic counterpart. Based on the results of chapter 2, the Brexit vote resulted in a deterioration in the UK economic activity explained by the reduction of 3.6% in the manufacturing sector employment. The results of chapter 3 show that the SDLT holiday caused the aggregate UK housing market to increase by 2.11% and 2.96% on annual rates in the two quarters following the tax holiday. In terms of the regional level, all the NUTS1 regions with the exception of Greater London benefited from the SDLT holiday. Chapter 4 evaluates the ability of the tones of newspapers topics in forecasting the volatility of stock returns. The findings are that news related to monetary policy, manufacturing, currency exchange, trade and trade policy can be combined to provide a pooled measure that has a forecasting power in predicting the

S&P 500 returns volatility.

This thesis also provides new grounds for future research. Chapter 2 sheds light on the impact of the Brexit vote on the UK labour market in the manufacturing sector. Future work can take net immigration into account as well. Chapter 3 investigates the impact of the SDLT holiday on the UK housing market. Future work could potentially look at the impact of the SDLT cuts introduced by the British Government in the Mini Budget of 2022. Chapter 4, the unsupervised machine learning approach is complemented with a dictionary based approach to provide news sentiment indices that can predict the volatility of stock returns. Future work can look at news sentiment related to US sectors and industries and evaluate their ability in predicting the volatility of returns of the S&P 14 sectors indices.

Appendix A

A.1 Chapter 2 Appendix

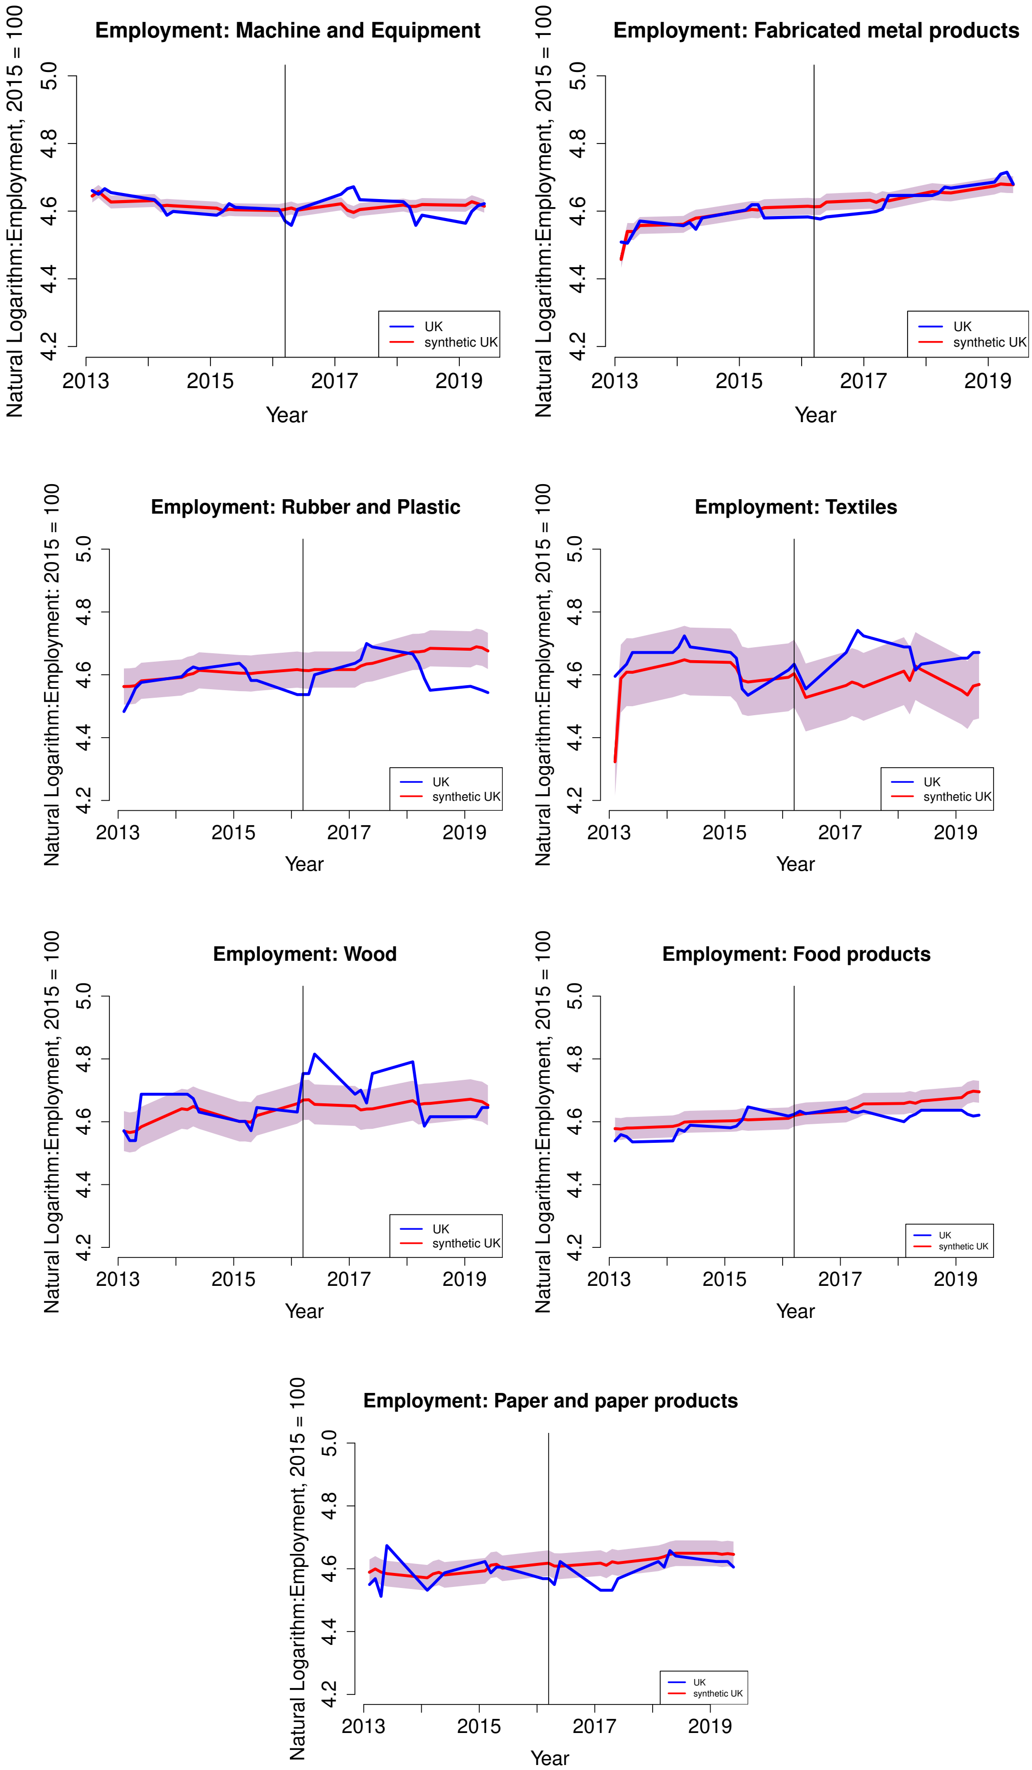

Figure A.1: The unaffected manufacturing divisions.

I follow Born et al. (2019) in constructing the shaded area as 2 standard deviations of the difference between the UK and the synthetic UK in the pre-Brexit period.

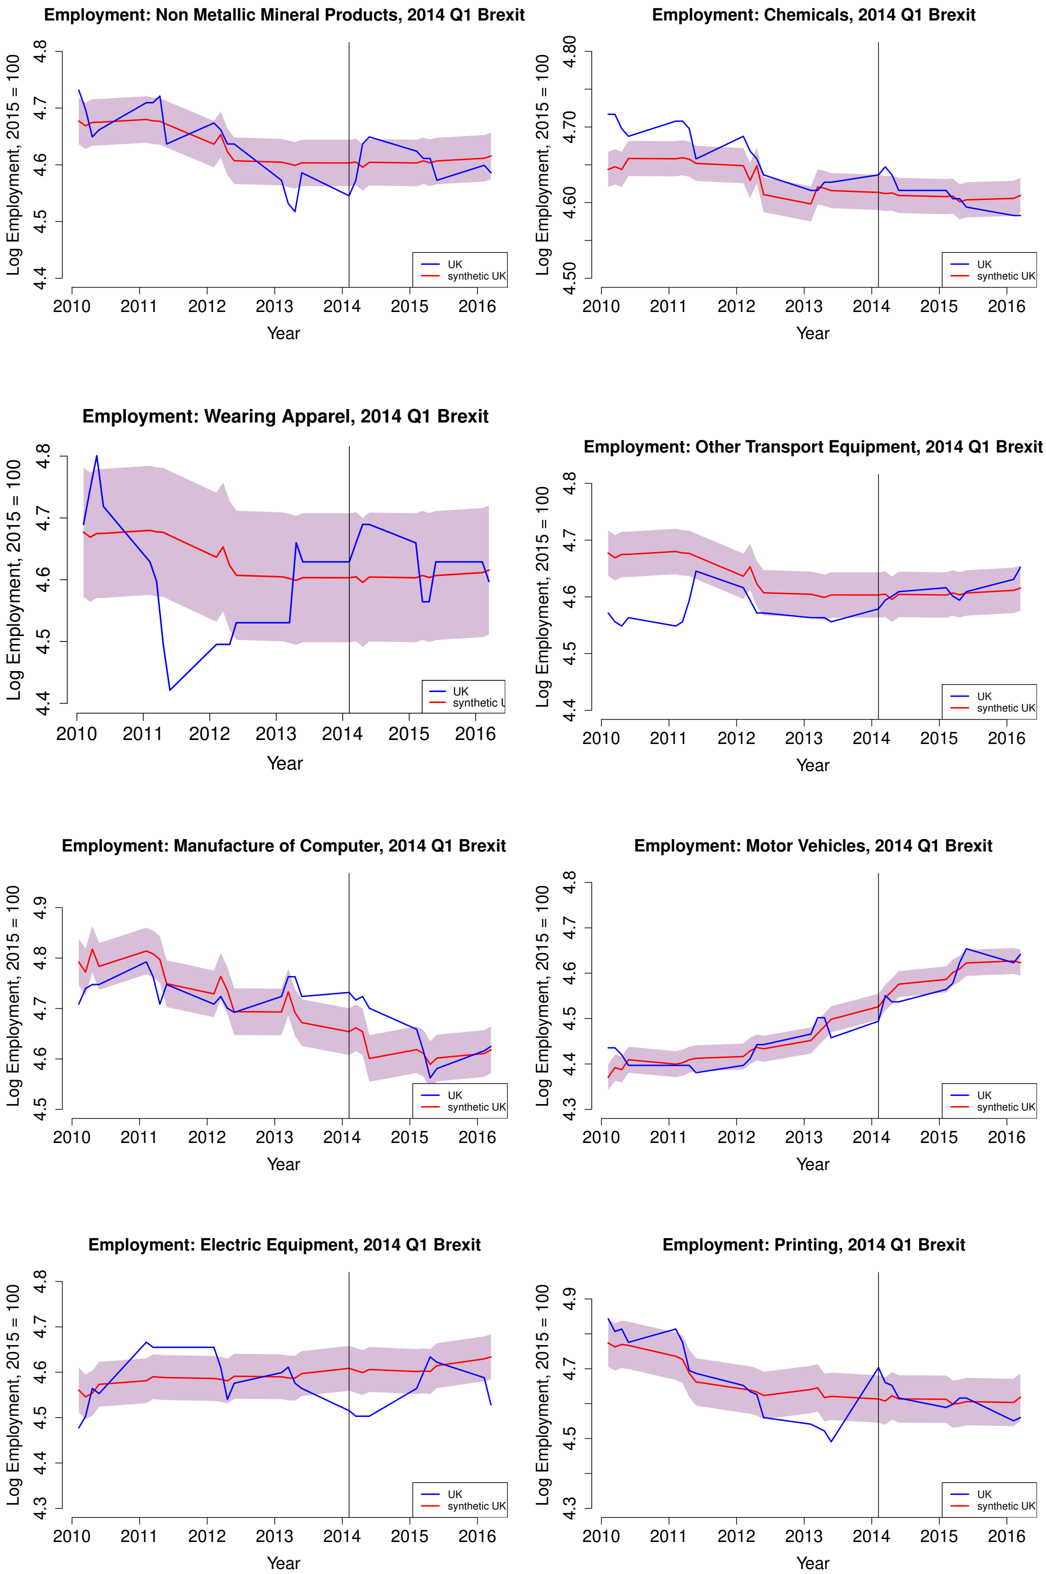

Figure A.2: Time placebo tests for the individual manufacturing divisions.

I follow Born et al. (2019) in constructing the shaded area as 2 standard deviations of the difference between the UK and the synthetic UK in the pre-Brexit period.

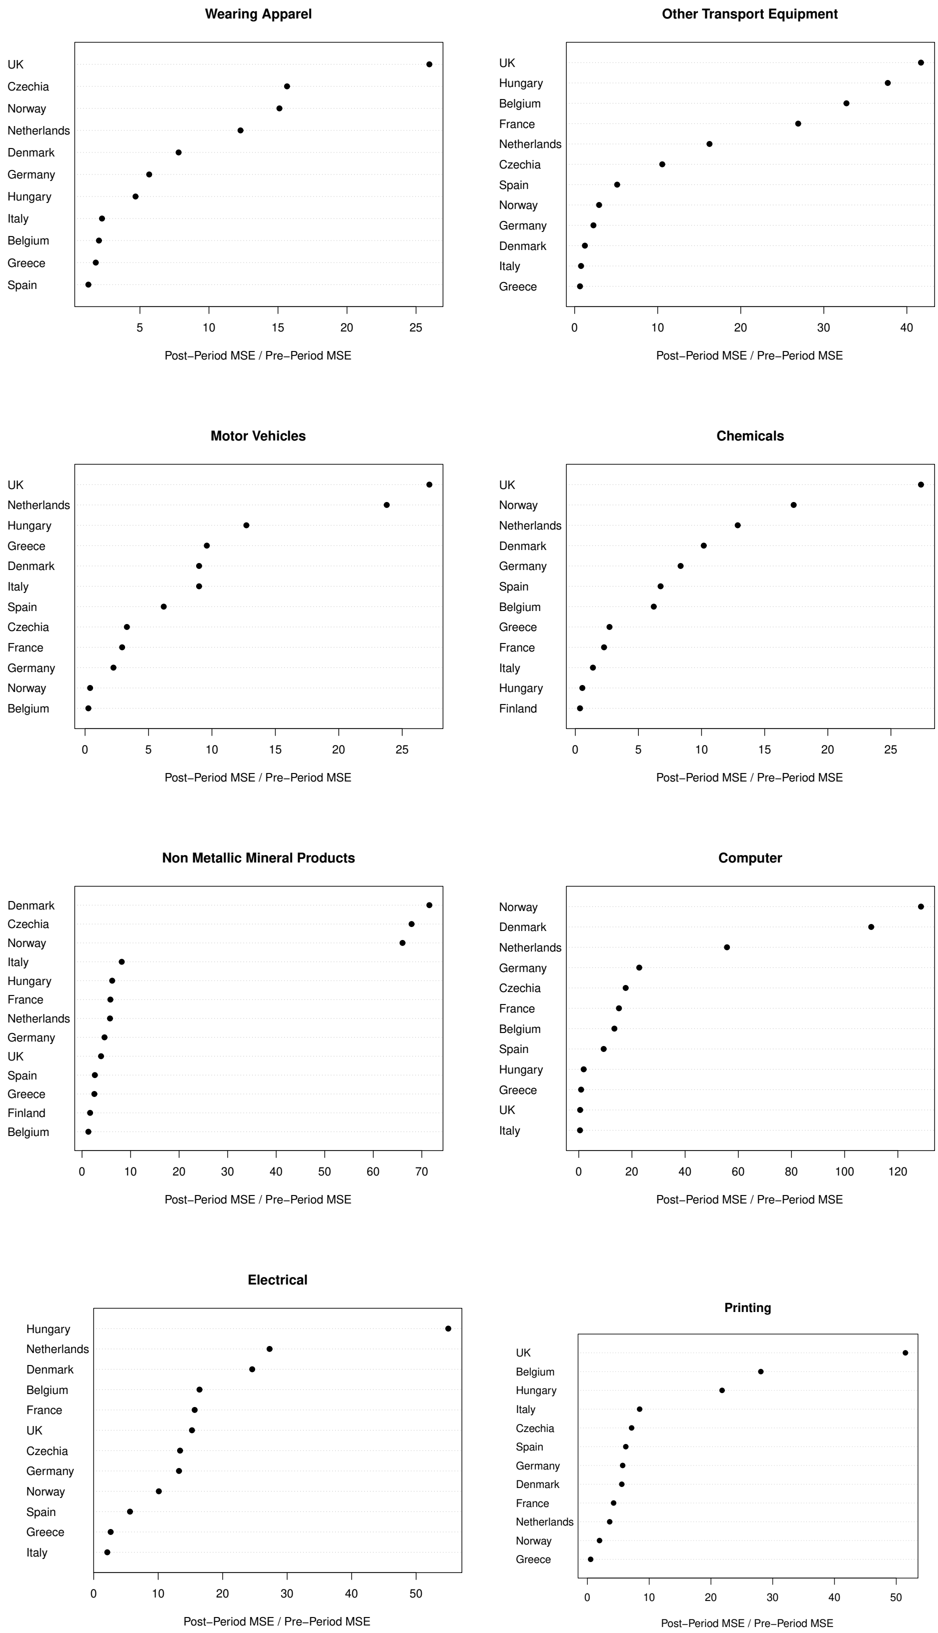

Figure A.3: MSE tests for the individual manufacturing divisions.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	-0.276^{***}	0.024^{**}	
	(0.039)	(0.010)	
UK		0.010	
		(0.023)	
$d_t{\times}UK$		-0.341^{***}	-0.341^{***}
		(0.033)	(0.031)
Constant	0.004	4.608^{***}	
	(0.028)	(0.007)	
Observations	28	336	336
R^2	0.657	0.376	0.295
Adjusted R^2	0.644	0.370	0.202
Residual Std. Error	0.103 (df = 26)	0.084 (df = 332)	
F Statistic	49.894^{***} (df = 1; 26)	66.648^{***} (df = 3; 332)	124.005^{***} (df = 1; 296)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.1: Synthetic Control Estimator Regression and DiD for wearing apparel division.

Estimation of equations 2.2.3 and 2.2.4 for the wearing apparel division. Column (1) shows the estimation of the Synthetic Control Regression Estimator. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^N$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	-0.048^{***}	0.034^{**}	
	(0.009)	(0.010)	
UK		-0.006	
		(0.023)	
$d_t \times UK$		-0.043	-0.043
		(0.034)	(0.031)
Constant	0.005	4.601^{***}	
	(0.006)	(0.007)	
Observations	28	336	336
R^2	0.657	0.376	0.295
Adjusted R^2	0.644	0.370	0.202
Residual Std. Error	0.103 (df = 26)	0.084 (df = 332)	
F Statistic	49.894^{***} (df = 1; 26)	66.648^{***} (df = 3; 332)	124.005^{***} (df = 1; 296)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.2: Synthetic Control Estimator Regression and DiD for other transport equipment division.

Estimation of equations 2.2.3 and 2.2.4 for the other transport equipment division. Column (1) shows the estimation of the Synthetic Control Regression Estimator. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	-0.093^{***}	0.0004	
	(0.013)	(0.003)	
UK		-0.010	
		(0.008)	
$d_t \times UK$		0.021^{*}	0.093^{***}
		(0.011)	(0.004)
Constant	0.005	1.525^{***}	
	(0.006)	(0.002)	
Observations	28	336	336
R^2	0.657	0.376	0.295
Adjusted R^2	0.644	0.370	0.202
Residual Std. Error	0.103 (df = 26)	0.084 (df = 332)	
F Statistic	49.894^{***} (df = 1; 26)	66.648^{***} (df = 3; 332)	124.005^{***} (df = 1; 296)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.3: Synthetic Control Estimator Regression and DiD for motor vehicles division.

Estimation of equations 2.2.3 and 2.2.4 for the motor vehicles division. Column (1) shows the estimation of the Synthetic Control Regression Estimator. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*panel linear*
	(1)	(2)	(3)
d_t	0.058***	0.043***	
	(0.014)	(0.005)	
UK		0.024	
		(0.014)	
$d_t \times UK$		0.002	0.002
		(0.019)	(0.017)
Constant	0.001	4.594***	
	(0.010)	(0.004)	
Observations	28	336	336
R^2	0.657	0.376	0.295
Adjusted R^2	0.644	0.370	0.202
Residual Std. Error	0.103 (df = 26)	0.084 (df = 332)	
F Statistic	49.894*** (df = 1; 26)	66.648*** (df = 3; 332)	124.005*** (df = 1; 296)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.4: Synthetic Control Estimator Regression and DiD for chemical products division.

Estimation of equations 2.2.3 and 2.2.4 for the chemical products division. Column (1) shows the estimation of the Synthetic Control Regression Estimator. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*panel linear*
	(1)	(2)	(3)
d_t	-0.050^{***}	0.029^{***}	
	(0.015)	(0.010)	
UK		0.062^{**}	
		(0.025)	
$d_t \times UK$		-0.078^{**}	-0.078^{***}
		(0.034)	(0.024)
Constant	0.038^{***}	4.622^{***}	
	(0.010)	(0.007)	
Observations	28	336	336
R^2	0.313	0.036	0.035
Adjusted R^2	0.286	0.027	-0.093
Residual Std. Error	0.039 (df = 26)	0.086 (df = 332)	
F Statistic	11.829^{***} (df = 1; 26)	4.077^{***} (df = 3; 332)	10.603^{***} (df = 1; 296)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.6: Synthetic Control Estimator Regression and DiD for Computer division.

Estimation of equations 2.2.3 and 2.2.4 for the Computer division. Column (1) shows the estimation of the Synthetic Control Regression Estimator. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*panel linear*
	(1)	(2)	(3)
d_t	-0.139^{***}	-0.010^{***}	
	(0.025)	(0.002)	
UK		-0.009^{*}	
		(0.005)	
$d_t \times UK$		0.002	0.002
		(0.007)	(0.006)
Constant	-0.024	1.533^{***}	
	(0.017)	(0.001)	
Observations	28	336	336
R^2	0.553	0.094	0.0004
Adjusted R^2	0.535	0.085	-0.131
Residual Std. Error	0.065 (df = 26)	0.017 (df = 332)	
F Statistic	32.118^{***} (df = 1; 26)	11.429^{***} (df = 3; 332)	0.131 (df = 1; 296)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.7: Synthetic Control Estimator Regression and DiD for printing division.

Estimation of equations 2.2.3 and 2.2.4 for the printing division. Column (1) shows the estimation of the Synthetic Control Regression Estimator. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^{N}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	0.094***	0.030***	
	(0.016)	(0.006)	
UK		−0.005	
		(0.015)	
$d_t \times UK$		−0.090***	−0.090***
		(0.021)	(0.019)
Constant	−0.018	4.610***	
	(0.011)	(0.004)	
Observations	28	364	364
R^2	0.580	0.143	0.063
Adjusted R^2	0.564	0.136	−0.053
Residual Std. Error	0.041 (df = 26)	0.053 (df = 360)	
F Statistic	35.972*** (df = 1; 26)	19.999*** (df = 3; 360)	21.675*** (df = 1; 323)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.8: Synthetic Control Estimator Regression and DiD for non metallic mineral products division.

Estimation of equations 2.2.3 and 2.2.4 for the non metallic mineral products division. Column (1) shows the estimation of the Synthetic Control Regression Estimator. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^N$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	0.055***	0.030***	
	(0.016)	(0.006)	
UK		−0.039**	
		(0.019)	
$d_t{\times}UK$		0.096***	0.096***
		(0.027)	(0.025)
Constant	−0.006	4.598***	
	(0.011)	(0.004)	
Observations	28	644	644
R^2	0.307	0.076	0.024
Adjusted R^2	0.280	0.072	−0.058
Residual Std. Error	0.043 (df = 26)	0.069 (df = 640)	
F Statistic	11.494*** (df = 1; 26)	17.518*** (df = 3; 640)	14.553*** (df = 1; 593)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.9: Synthetic Control Estimator Regression and DiD for furniture division.

Estimation of equations 2.2.3 and 2.2.4 for the furniture division. Column (1) shows the estimation of the Synthetic Control Regression Estimator. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^{N}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	-0.183^{***}	0.014^{*}	
	(0.028)	(0.007)	
UK		-0.054^{***}	
		(0.018)	
$d_t \times UK$		-0.071^{***}	-0.071^{***}
		(0.026)	(0.022)
Constant	-0.038^{*}	4.619^{***}	
	(0.020)	(0.005)	
Observations	28	336	336
R^2	0.627	0.147	0.035
Adjusted R^2	0.613	0.139	−0.093
Residual Std. Error	0.073 (df = 26)	0.065 (df = 332)	
F Statistic	43.786^{***} (df = 1; 26)	19.075^{***} (df = 3; 332)	10.582^{***} (df = 1; 296)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.11: Synthetic Control Estimator Regression and DiD for electrical products division.

Estimation of equations 2.2.3 and 2.2.4 for the electrical division. Column (1) shows the estimation of the Synthetic Control Regression Estimator. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 2.2.4, where the dependent variable is the log employment index is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables.

A.2 Chapter 3 Appendix

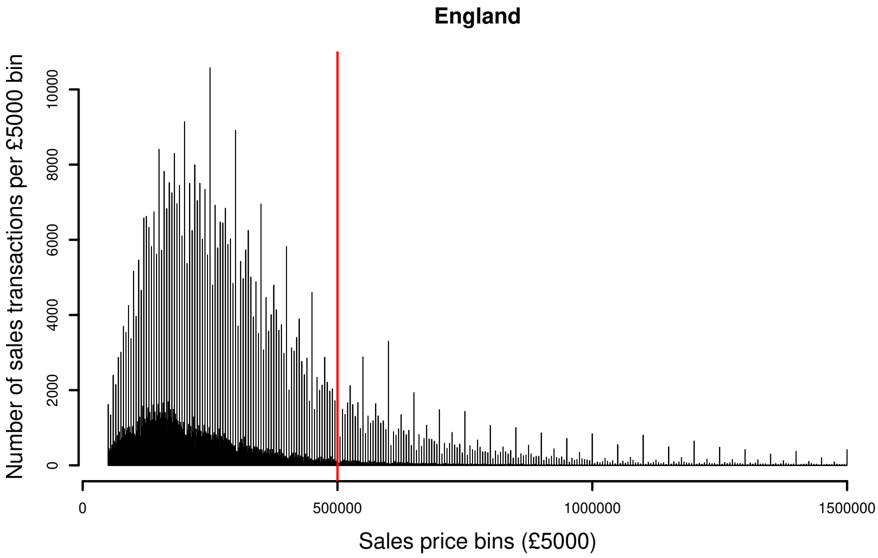

(a) England

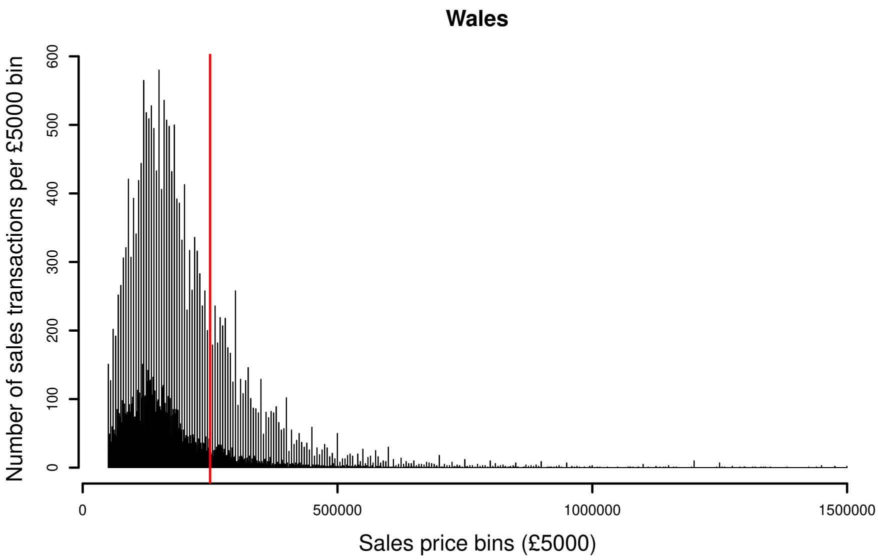

(b) Wales

Figure A.4: Density of property transactions for 2019.

The vertical red line indicates the cut off amount for SDLT exemption, £500,000 for England and £250,000 for Wales. Source: HM Land Registry.

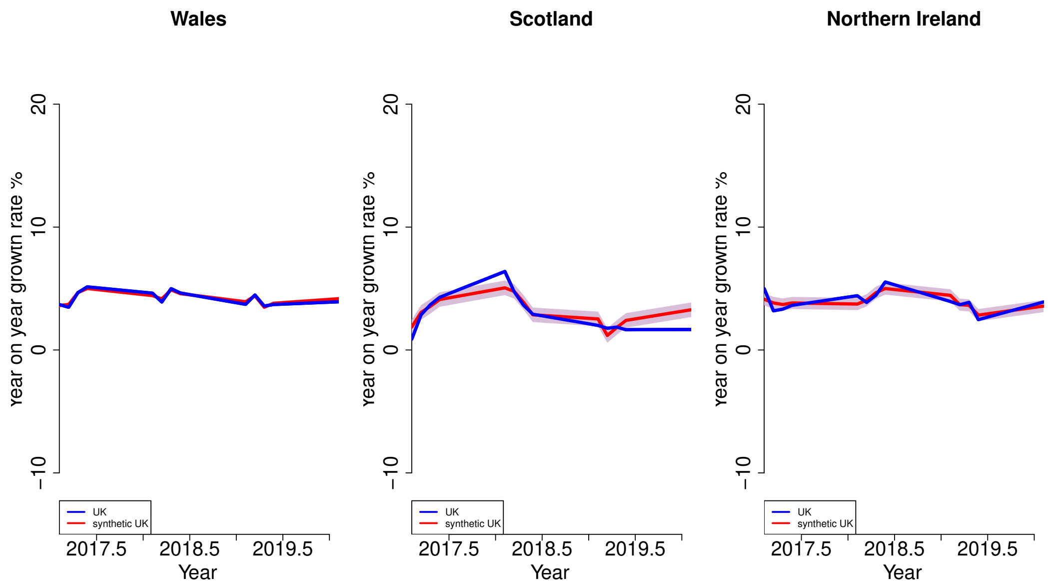

Figure A.5: In-time placebo test: British nations.

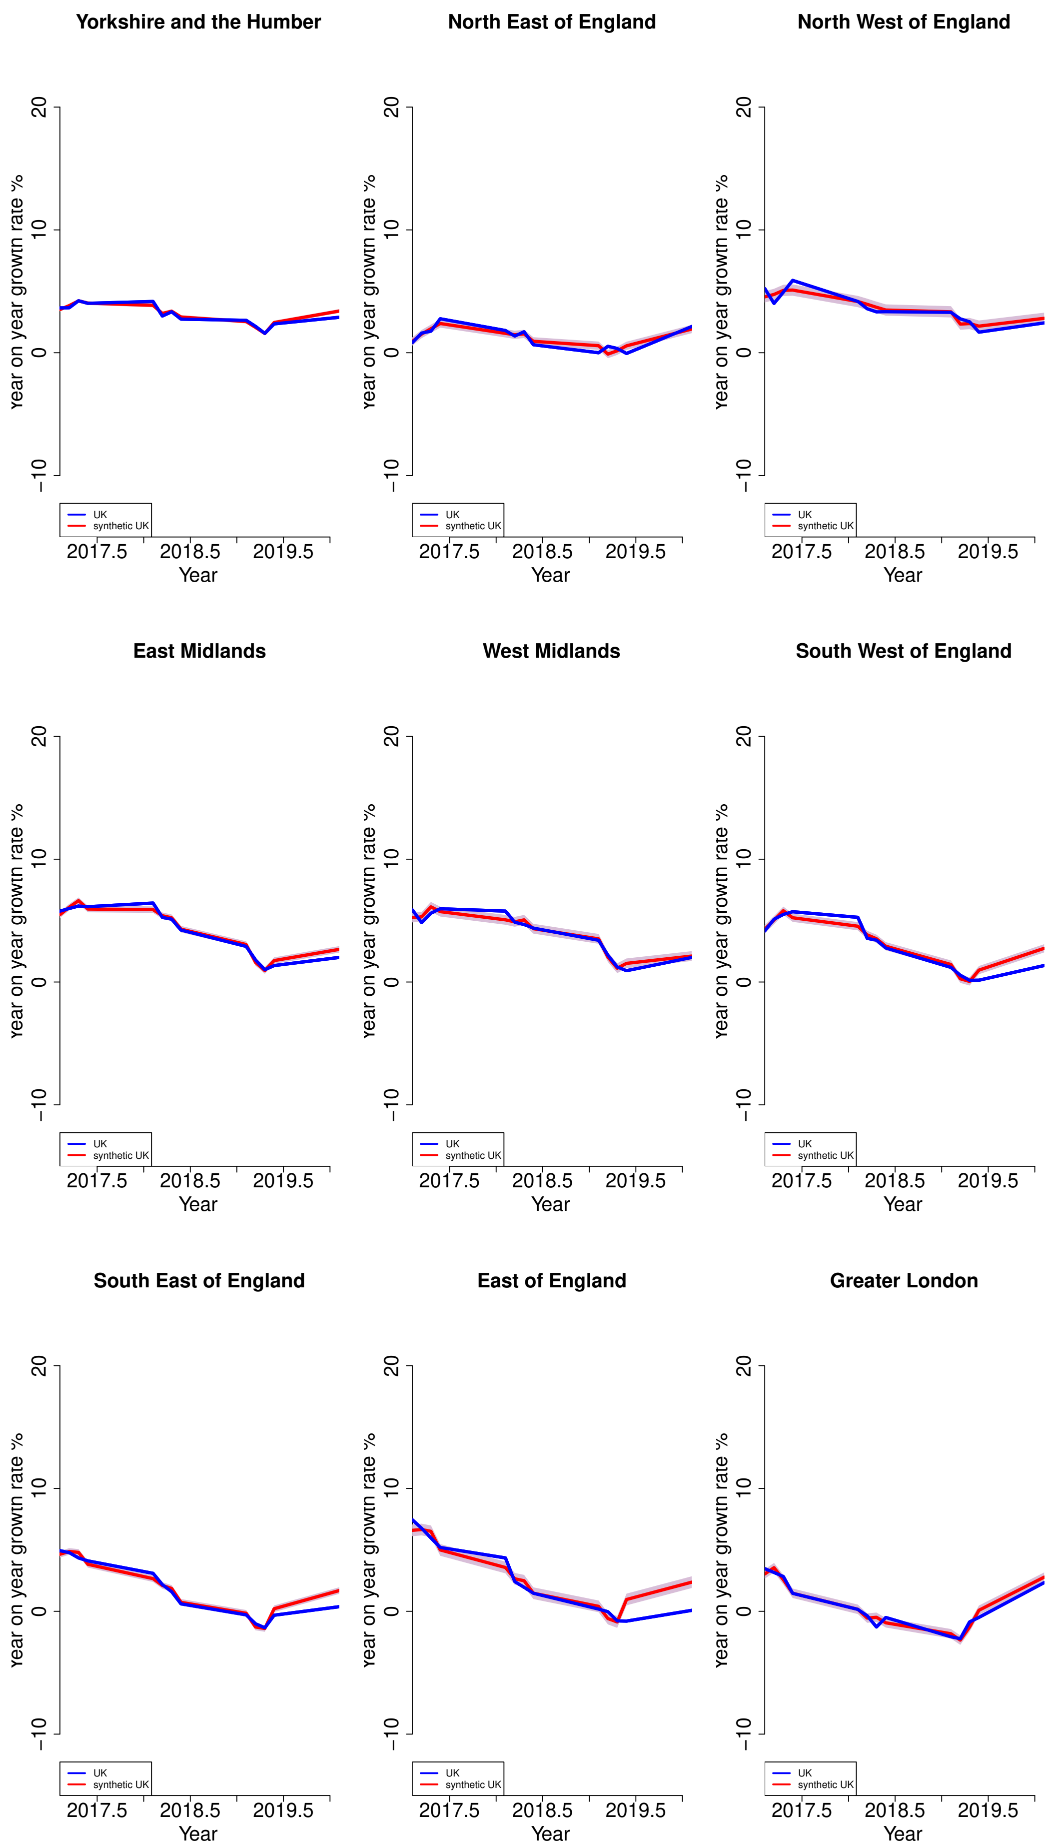

Figure A.6: In-time-placebo test: English regions.

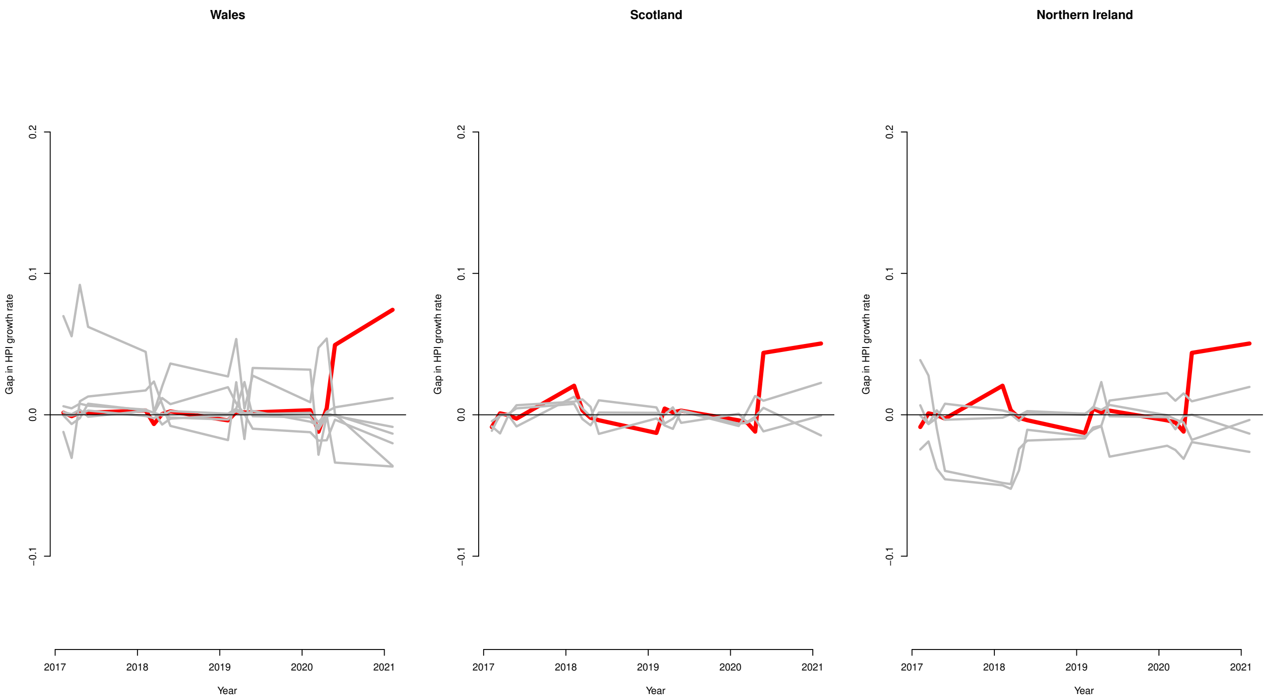

Figure A.7: Country placebo test: British regions.

Gaps in the British regions and the significant regions in the donor pool

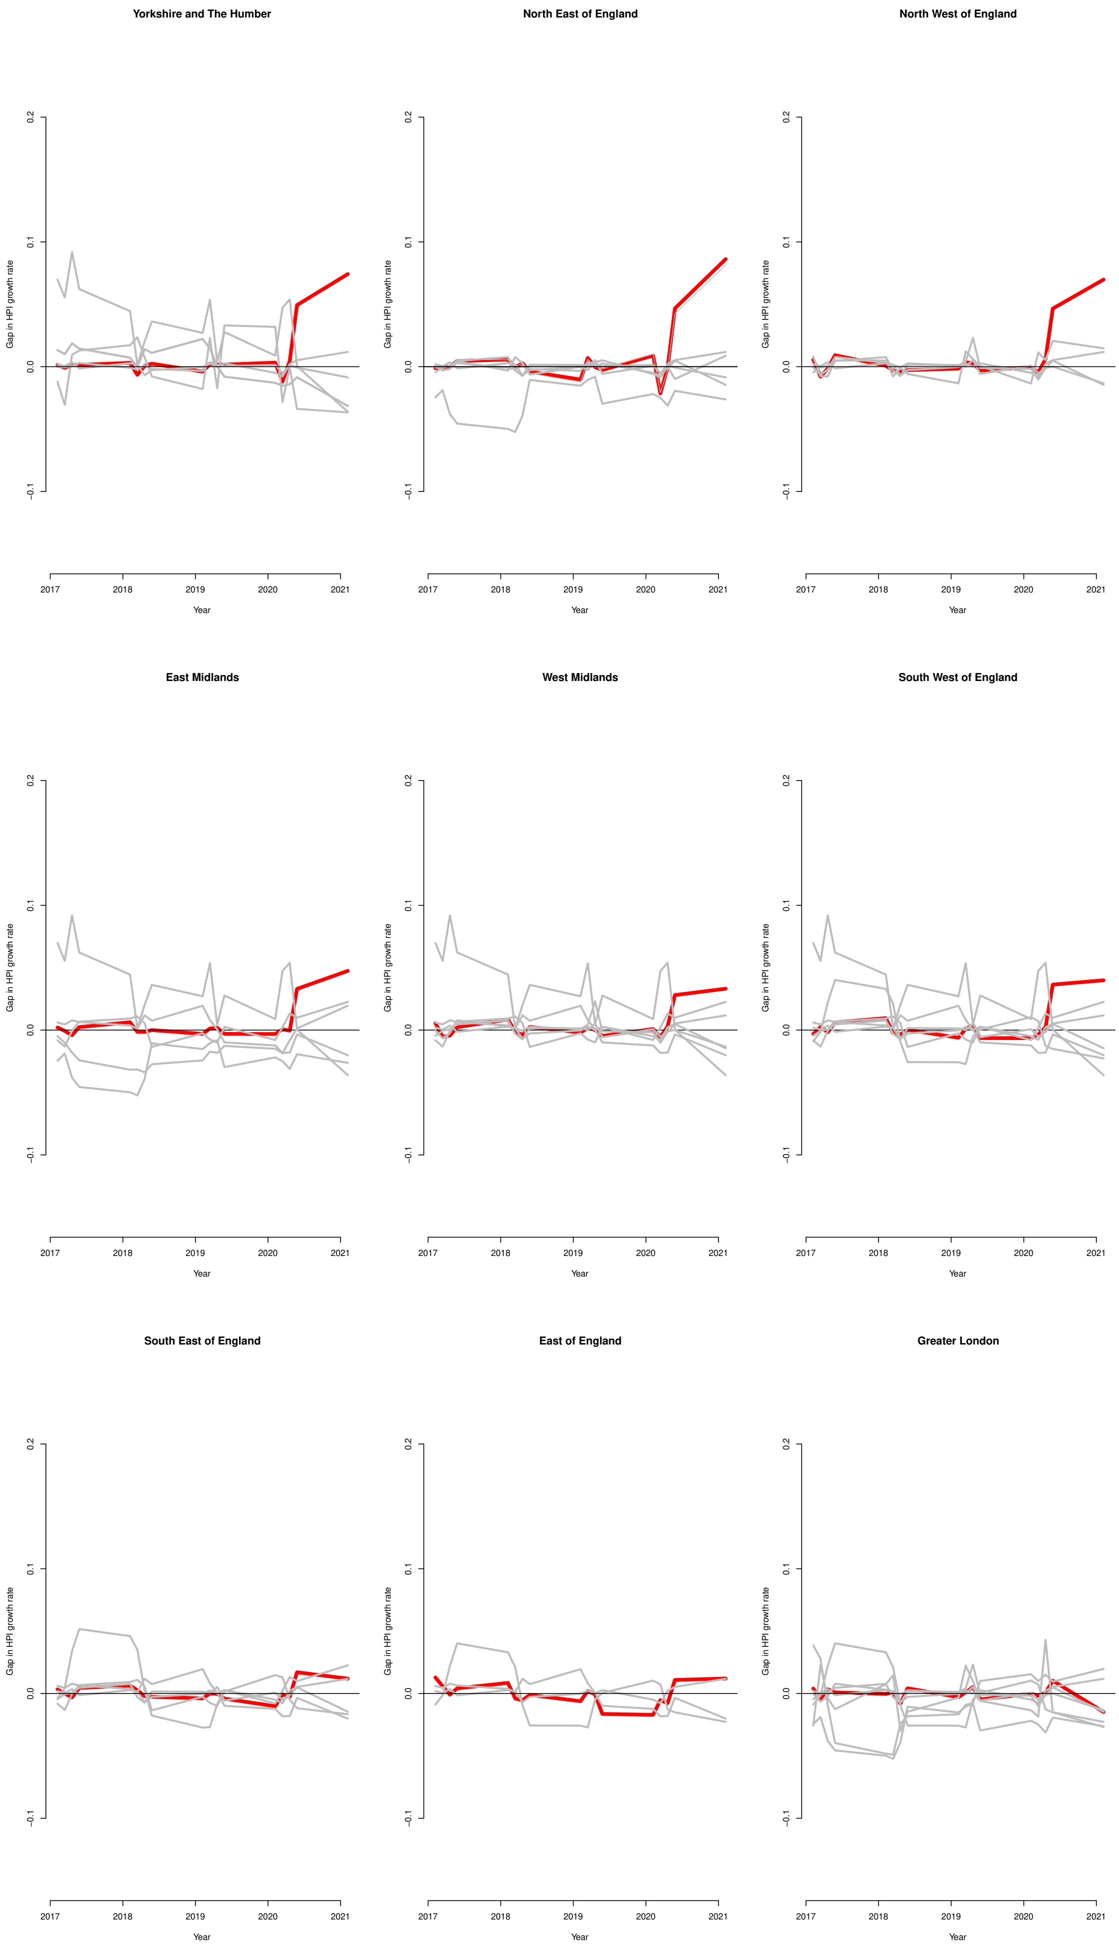

Figure A.8: Country placebo test: English regions.

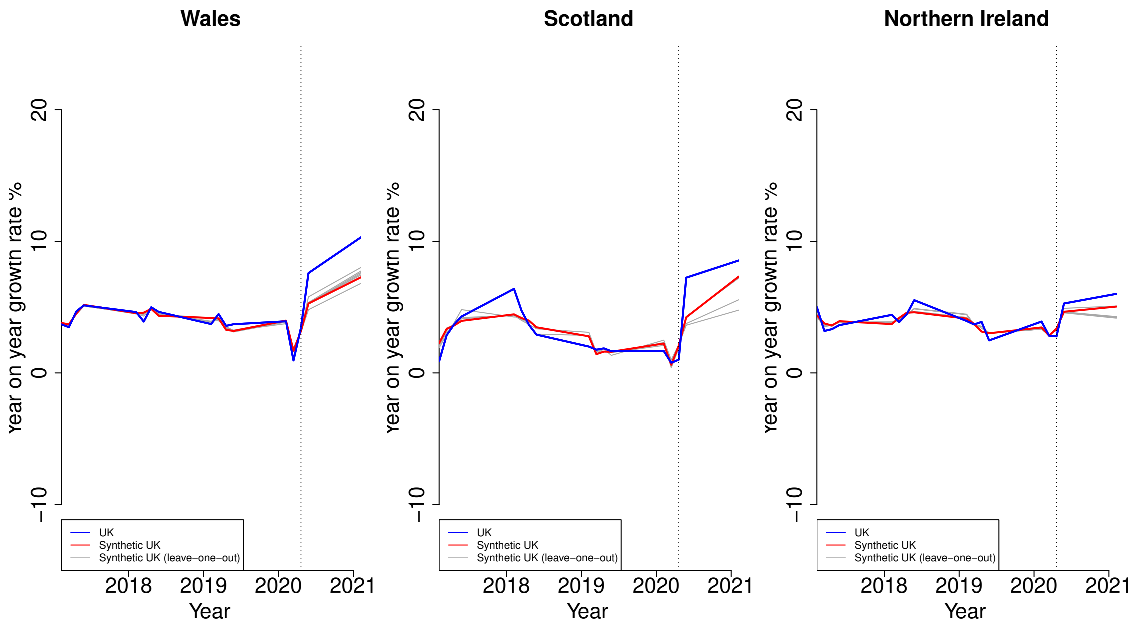

Figure A.9: Leave one control unit out: British regions

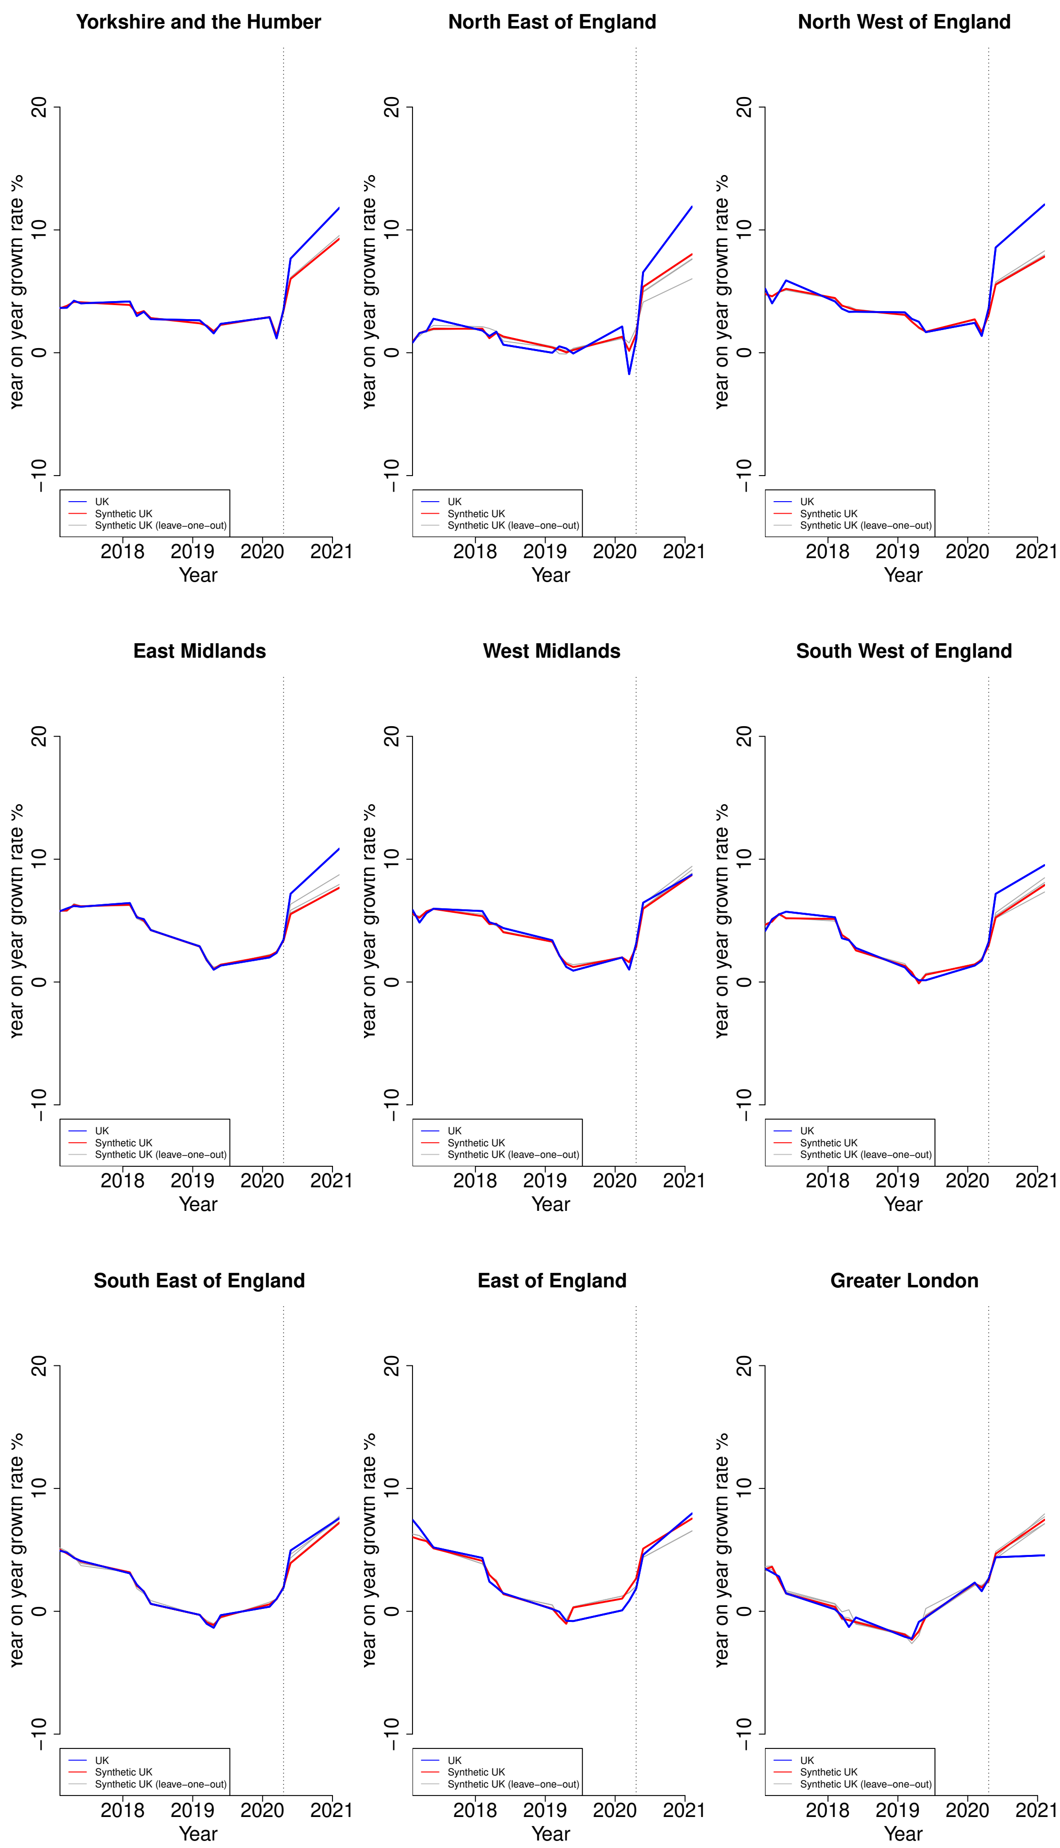

Figure A.10: Leave one control unit out: English regions

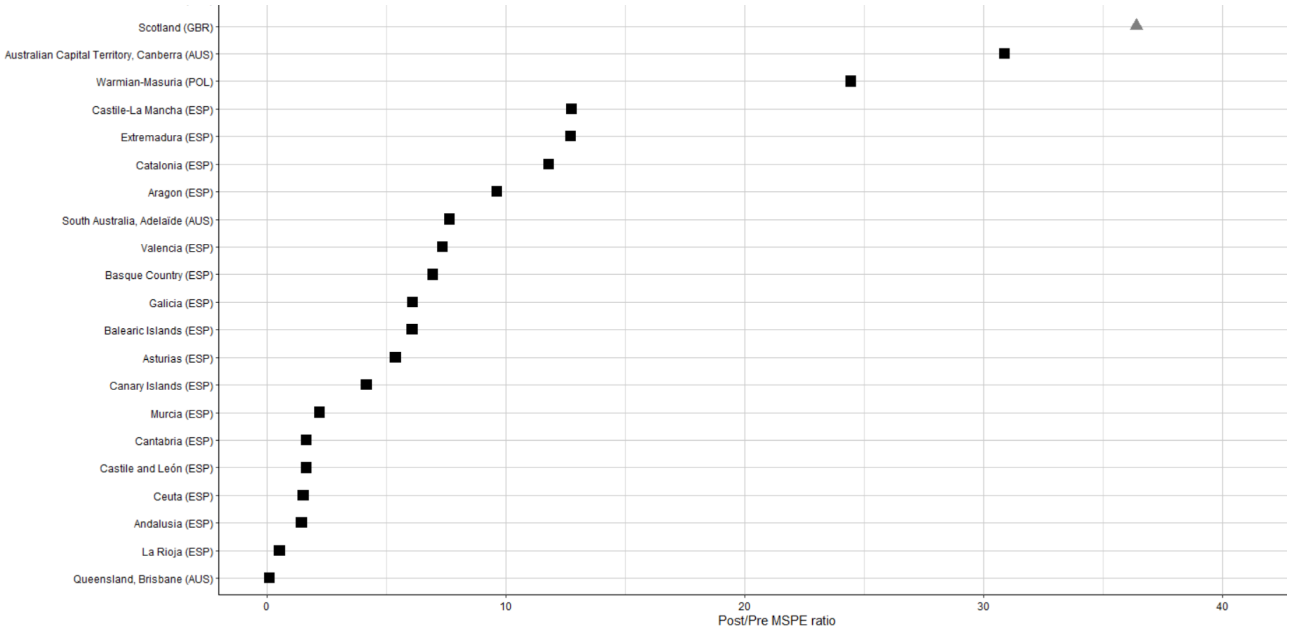

Figure A.11: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for Scotland (GBR)

The p-value = $\frac{1}{14} \approx 0.0714$. The plot is generated Using the package SCtools.

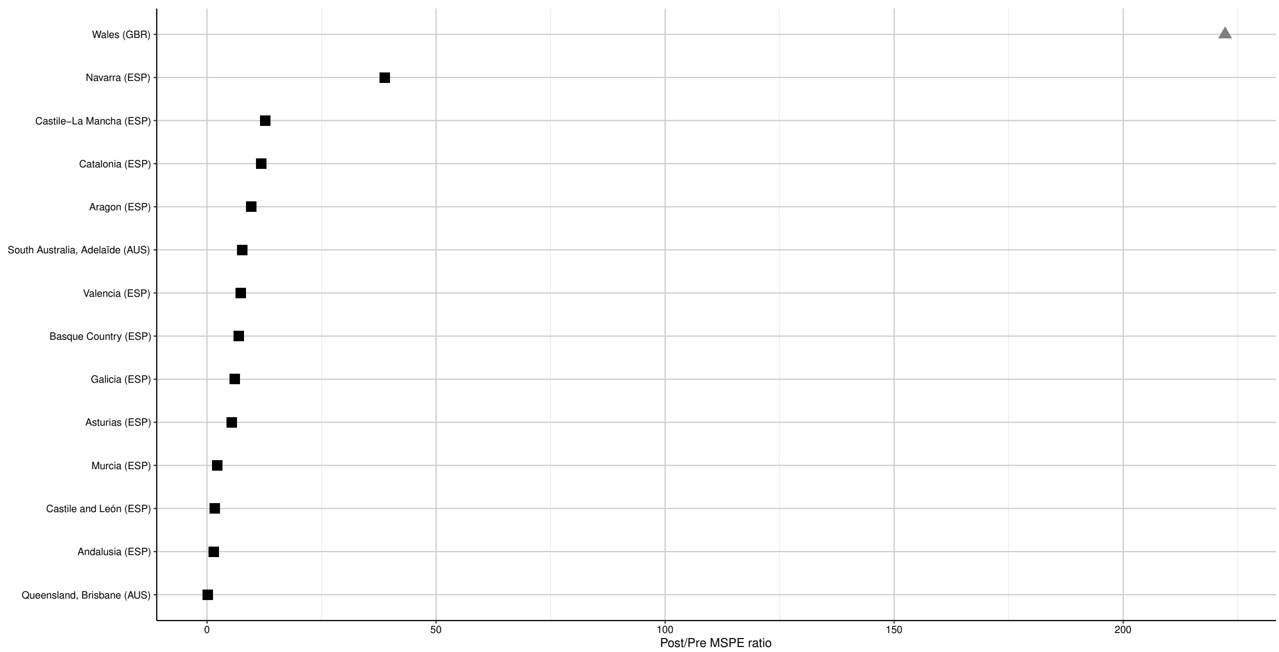

Figure A.12: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for Wales (GBR)

The p-value = $\frac{1}{21} \approx 0.047$. The plot is generated Using the package SCtools.

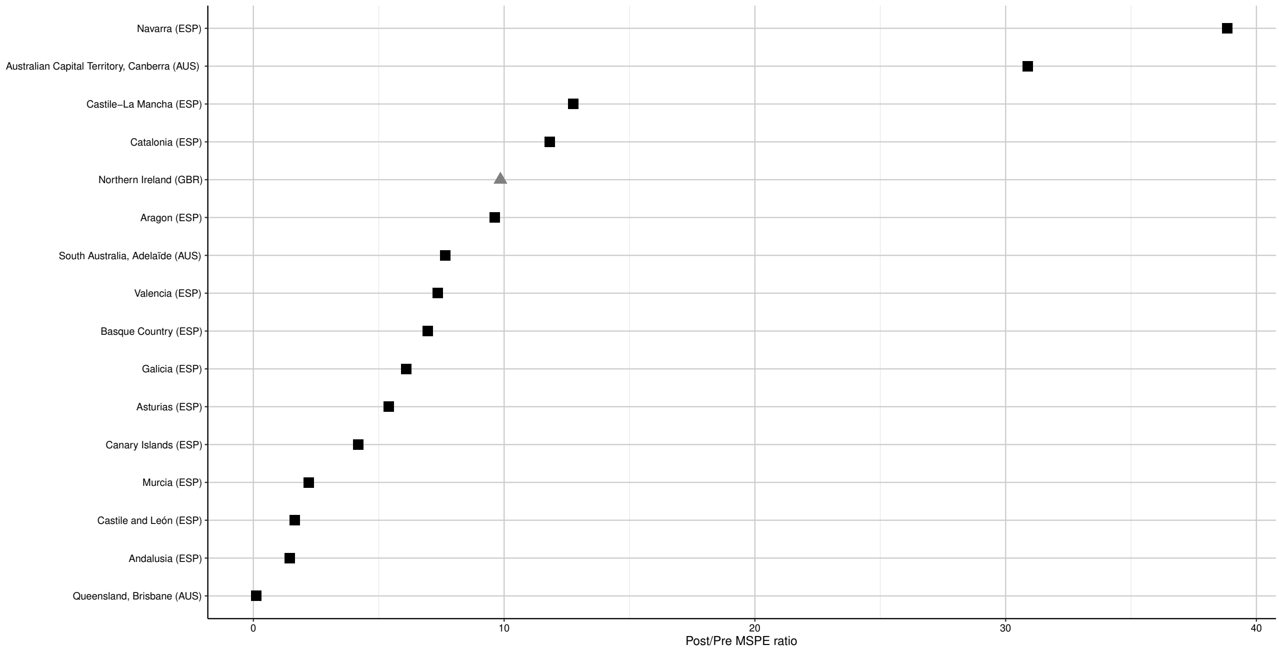

Figure A.13: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for Northern Ireland (GBR)

The p-value $= \frac{5}{16} \approx 0.3125$. The plot is generated Using the package SCtools.

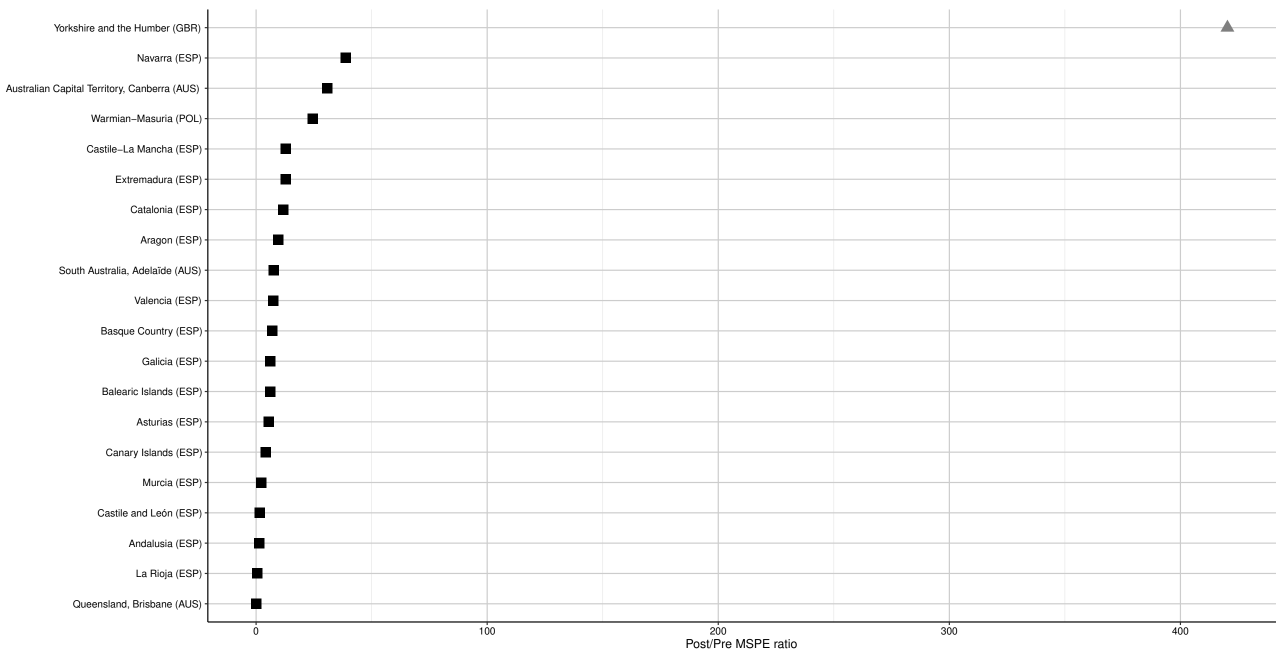

Figure A.14: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for Yorkshire and the Humber (GBR)

The p-value $= \frac{1}{20} \approx 0.05$. The plot is generated Using the package SCtools.

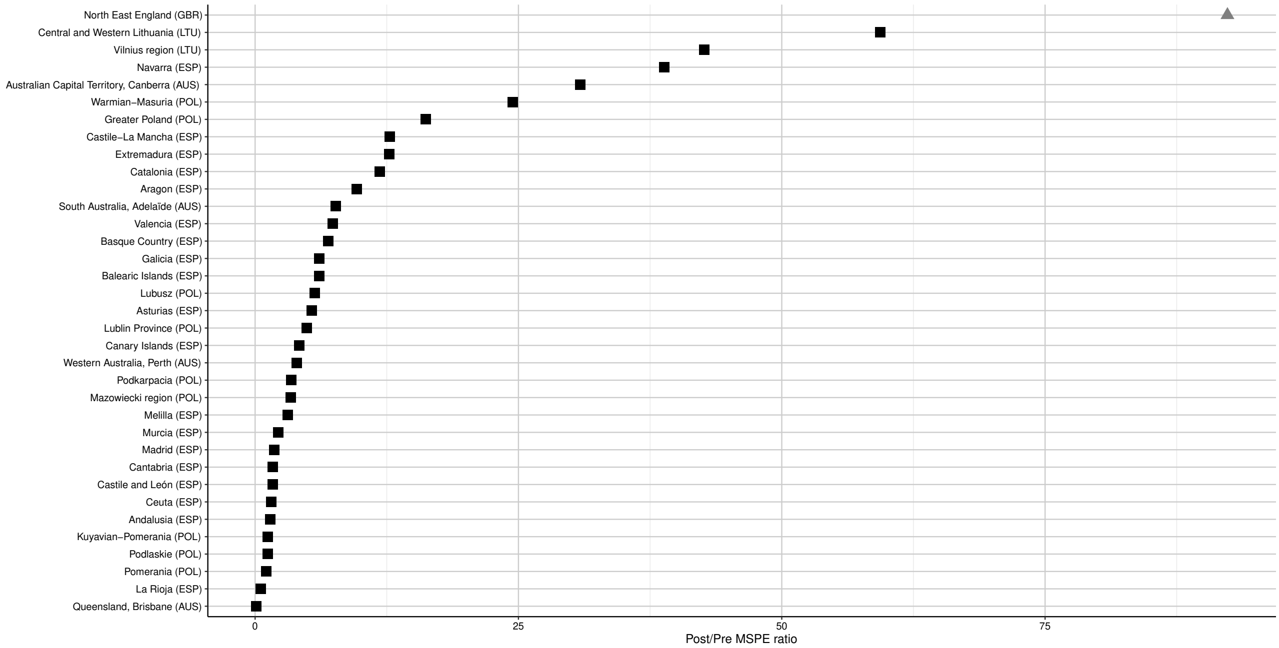

Figure A.15: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for North East of England (GBR)

The p-value $= \frac{1}{35} \approx 0.029$. The plot is generated Using the package SCtools.

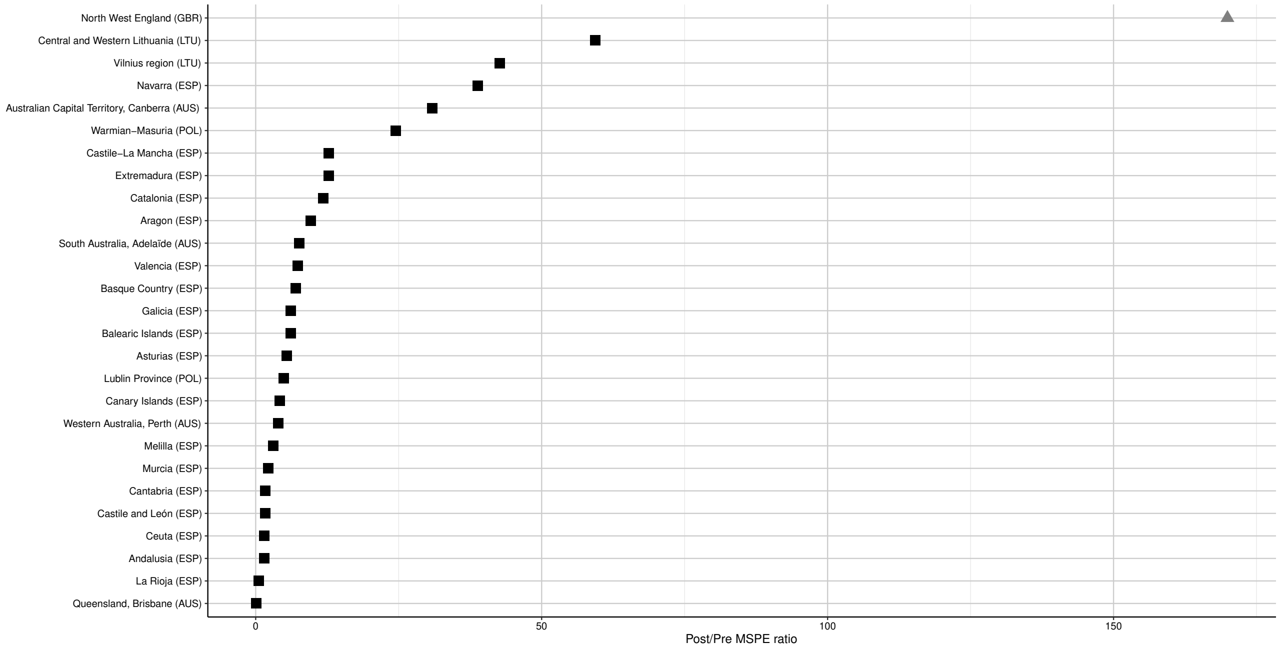

Figure A.16: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for North West of England (GBR)

The p-value $= \frac{1}{27} \approx 0.037$. The plot is generated Using the package SCtools.

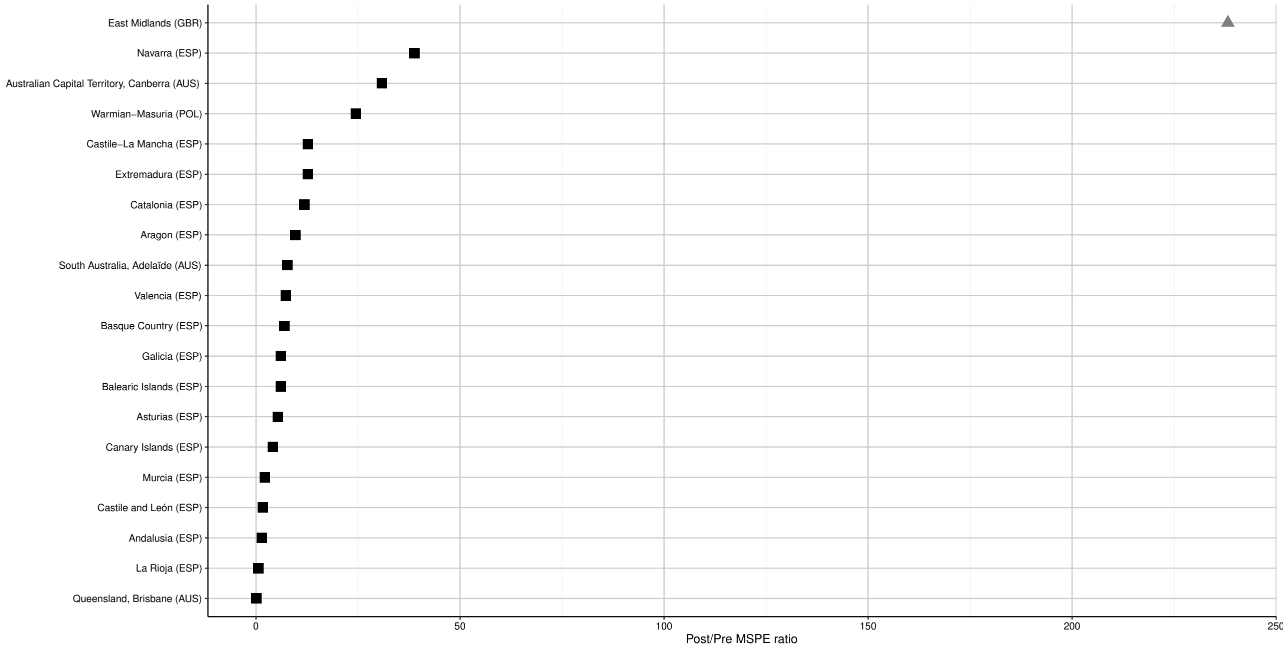

Figure A.17: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for East Midlands (GBR)

The p-value $= \frac{1}{20} \approx 0.05$. The plot is generated Using the package SCtools.

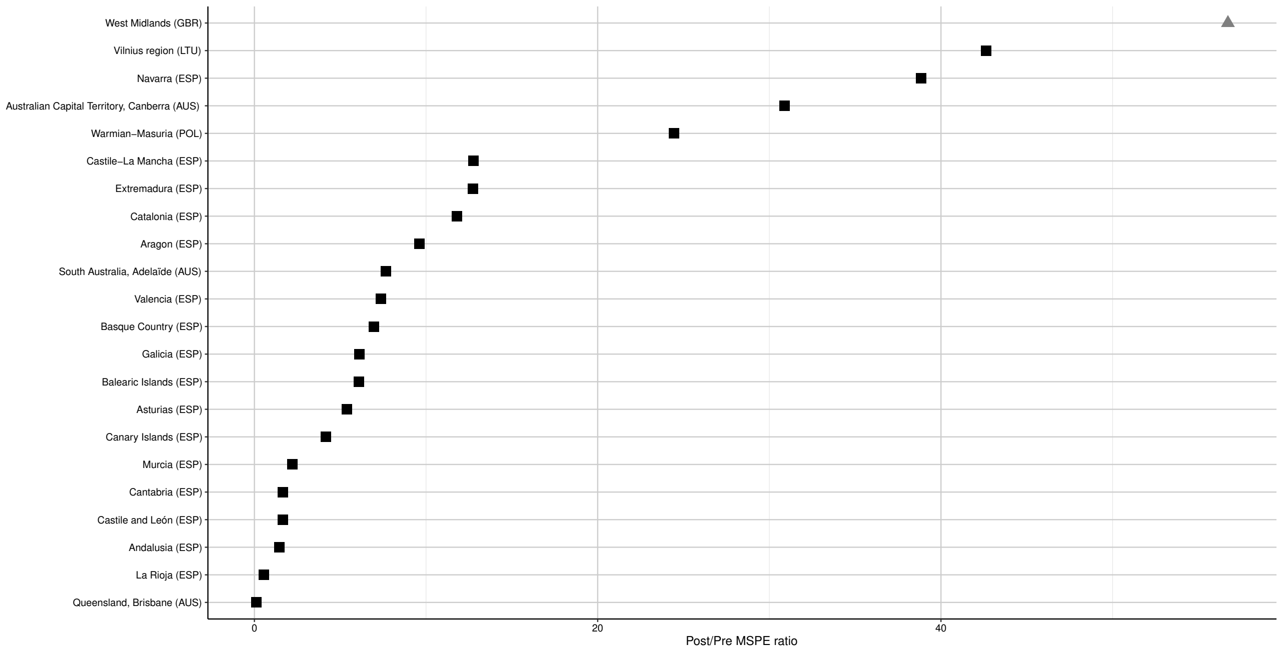

Figure A.18: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for West Midlands (GBR)

The p-value $= \frac{1}{22} \approx 0.045$. The plot is generated Using the package SCtools.

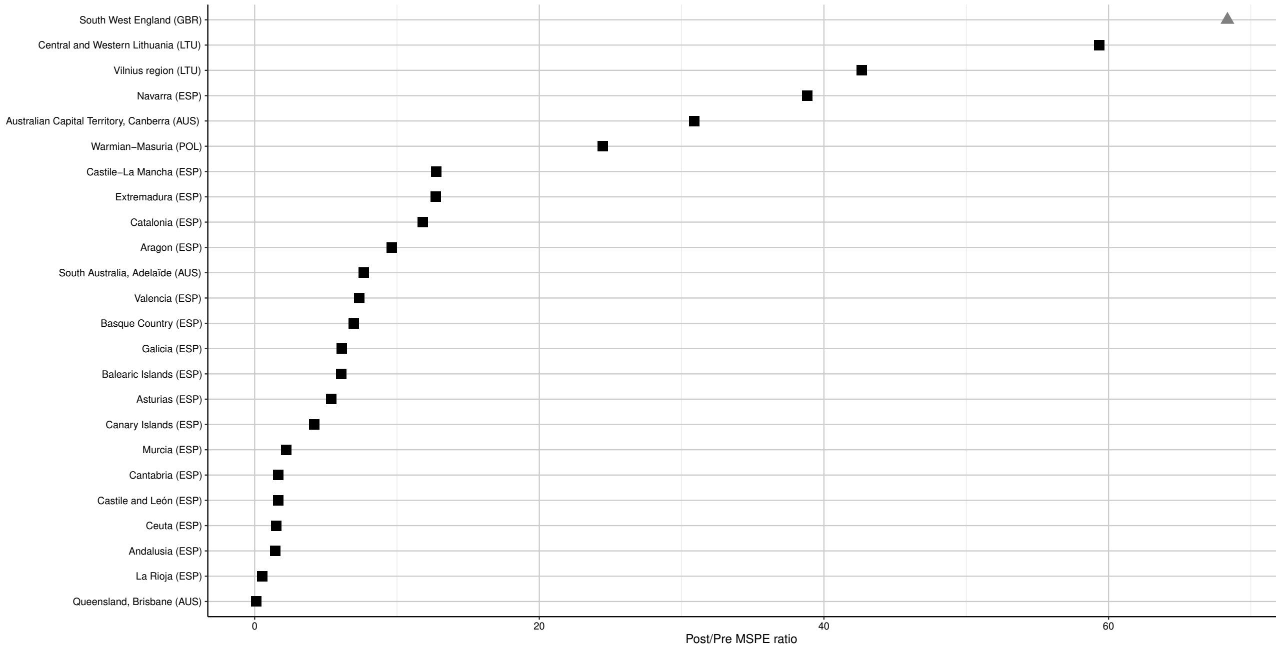

Figure A.19: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for South West of England (GBR)

The p-value $= \frac{1}{24} \approx 0.042$. The plot is generated Using the package SCtools.

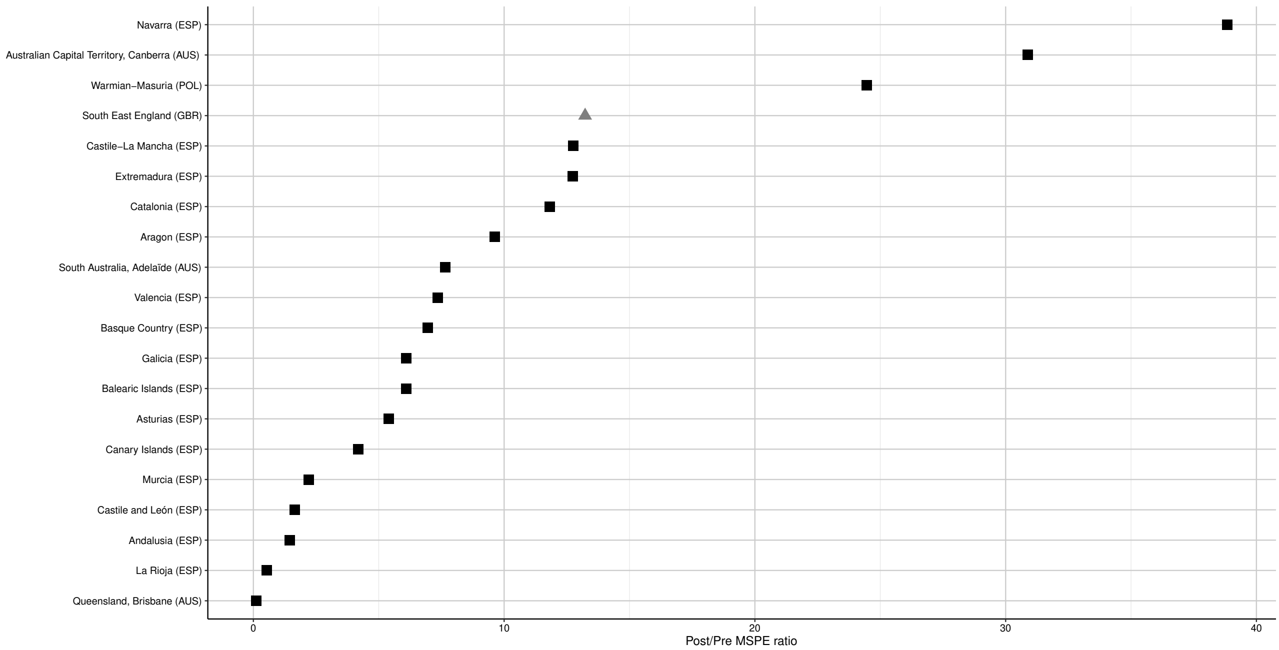

Figure A.20: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for South East of England (GBR)

The p-value $= \frac{4}{20} \approx 0.02$. The plot is generated Using the package SCtools.

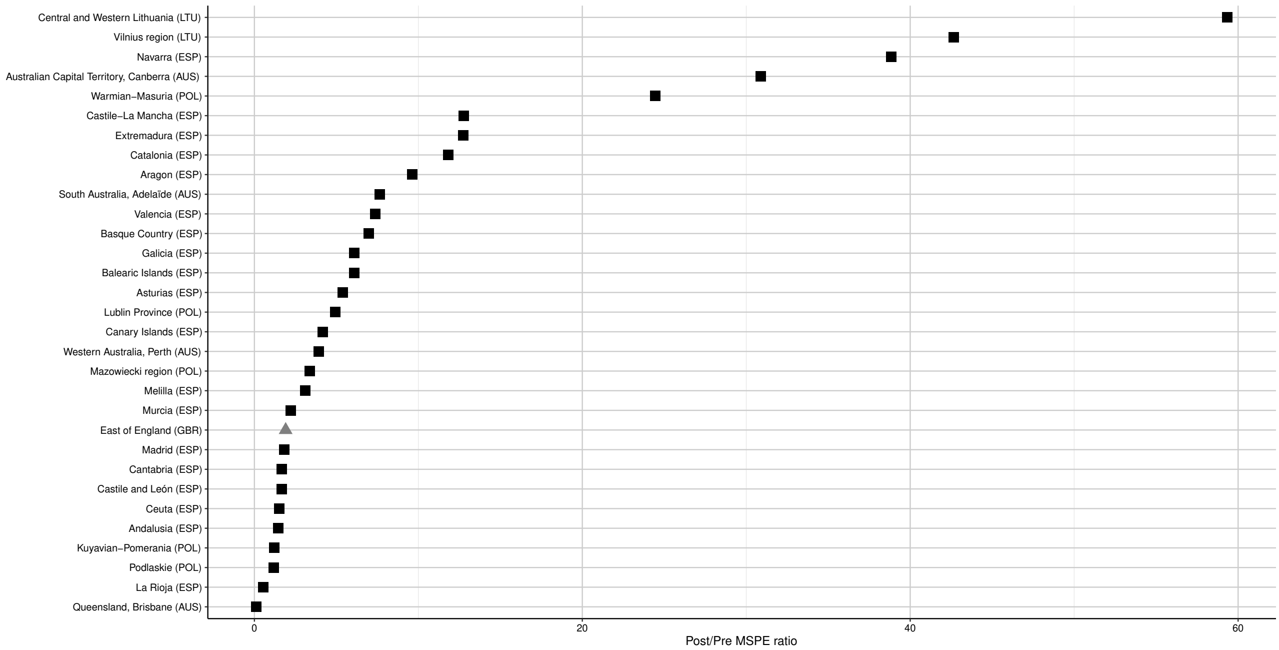

Figure A.21: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for East of England (GBR)

The p-value $= \frac{23}{32} \approx 0.72$. The plot is generated Using the package SCtools.

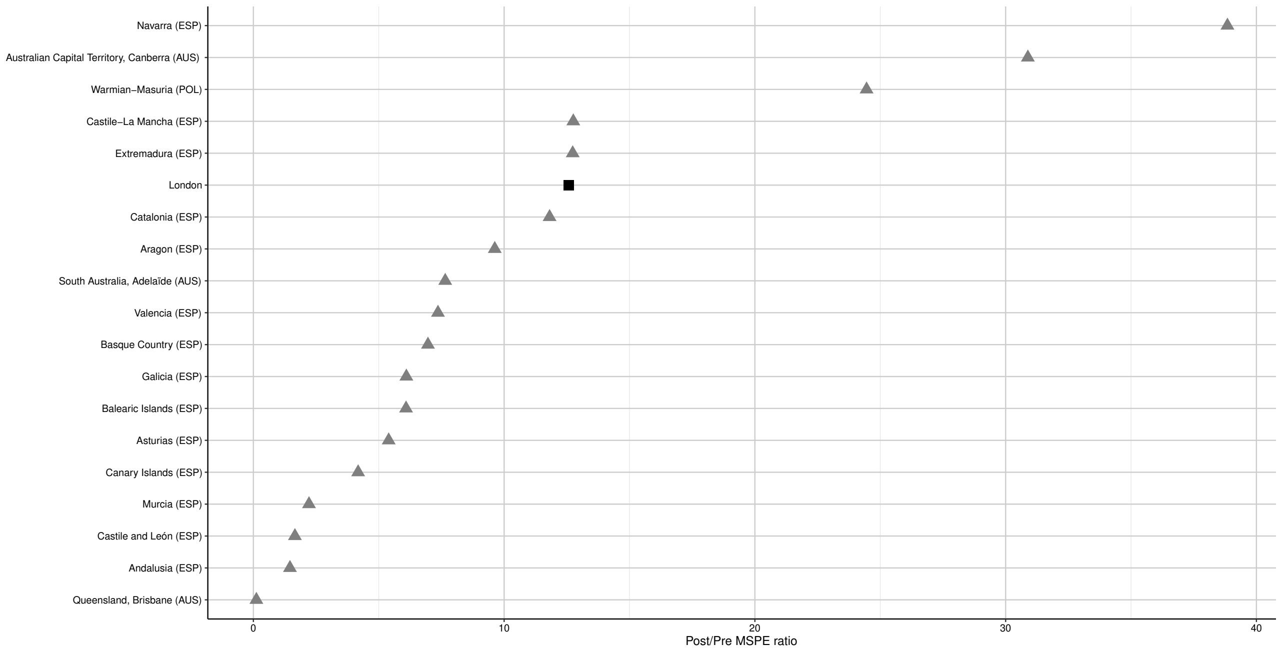

Figure A.22: Ratio of the MSE of the Post-SDLT holiday period to the Pre-SDLT holiday period for Greater London (GBR)

The p-value $= \frac{6}{20} \approx 0.3$. The plot is generated Using the package SCtools.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	3.892***	−0.001	
	(1.094)	(0.003)	
UK		−0.001	
		(0.010)	
$d_t \times UK$		0.038*	0.012
		(0.021)	(0.016)
PI		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.004***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0003
		(0.0002)	(0.001)
Interest Rate		0.001	−0.029***
		(0.002)	(0.004)
Constant	0.385	−0.073	
	(0.325)	(0.049)	
Observations	34	1,054	1,054
R^2	0.283	0.211	0.173
Adjusted R^2	0.261	0.206	0.103
Residual Std. Error	1.810 (df = 32)	0.035 (df = 1046)	
F Statistic	12.649*** (df = 1; 32)	40.056*** (df = 7; 1046)	40.560*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.13: Effect of the SDLT holiday on the Wales HPI YoY.

Estimating the effect of the SDLT holiday on the Wales HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the Wales and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^N$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	2.515^{**}	-0.001	
	(0.976)	(0.003)	
UK		-0.001	
		(0.010)	
$d_t \times$ UK		0.006	0.017
		(0.020)	(0.016)
PI		-0.002^{***}	0.002^{***}
		(0.0002)	(0.0003)
CPI		0.005^{***}	0.006^{***}
		(0.0005)	(0.001)
Employment		-0.002^{***}	0.0004
		(0.0002)	(0.001)
Interest Rate		0.001	-0.030^{***}
		(0.002)	(0.004)
Constant	0.228	-0.076	
	(0.290)	(0.050)	
Observations	34	1,054	1,054
R^2	0.172	0.208	0.173
Adjusted R^2	0.146	0.203	0.103
Residual Std. Error	1.614 (df = 32)	0.035 (df = 1046)	
F Statistic	6.641^{**} (df = 1; 32)	39.339^{***} (df = 7; 1046)	40.698^{***} (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.15: Effect of the SDLT holiday on the Scotland HPI YoY.

Estimating the effect of the SDLT holiday on the Scotland HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the Scotland and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^N$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	0.762*	−0.0001	
	(0.388)	(0.003)	
UK		0.009	
		(0.010)	
$d_t \times$ UK		−0.006	0.005
		(0.020)	(0.016)
PI		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.004***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0004
		(0.0002)	(0.001)
Interest Rate		0.001	−0.030***
		(0.002)	(0.004)
Constant	0.099	−0.071	
	(0.115)	(0.050)	
Observations	34	1,054	1,054
R^2	0.108	0.209	0.172
Adjusted R^2	0.080	0.204	0.103
Residual Std. Error	0.641 (df = 32)	0.035 (df = 1046)	
F Statistic	3.862* (df = 1; 32)	39.486*** (df = 7; 1046)	40.452*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.17: Effect of the SDLT holiday on the Northern Ireland HPI YoY.

Estimating the effect of the SDLT holiday on the Northern Ireland HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the Northern Ireland and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	3.312***	−0.001	
	(0.938)	(0.003)	
UK		0.001	
		(0.010)	
$d_t\times$ UK		0.021	0.032**
		(0.020)	(0.016)
PI		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.004***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0005
		(0.0002)	(0.001)
Interest Rate		0.001	−0.030***
		(0.002)	(0.004)
Constant	0.323	−0.072	
	(0.279)	(0.050)	
Observations	34	1,054	1,054
R^2	0.280	0.209	0.176
Adjusted R^2	0.258	0.204	0.106
Residual Std. Error	1.551 (df = 32)	0.035 (df = 1046)	
F Statistic	12.470*** (df = 1; 32)	39.597*** (df = 7; 1046)	41.409*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.19: Effect of the SDLT holiday on the Yorkshire and the Humber HPI YoY.

Estimating the effect of the SDLT holiday on the Yorkshire and the Humber HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the Yorkshire and the Humber and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	4.062***	−0.001	
	(1.291)	(0.003)	
UK		−0.019*	
		(0.010)	
$d_t\times$ UK		0.025	0.035**
		(0.020)	(0.016)
PI		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.005***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0005
		(0.0002)	(0.001)
Interest Rate		0.001	−0.030***
		(0.002)	(0.004)
Constant	0.350	−0.087*	
	(0.383)	(0.050)	
Observations	34	1,054	1,054
R^2	0.236	0.211	0.176
Adjusted R^2	0.212	0.206	0.107
Residual Std. Error	2.135 (df = 32)	0.035 (df = 1046)	
F Statistic	9.904*** (df = 1; 32)	40.019*** (df = 7; 1046)	41.618*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.21: Effect of the SDLT holiday on the North East HPI YoY.

Estimating the effect of the SDLT holiday on the North East HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the North East and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^{N}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	3.733***	−0.001	
	(1.033)	(0.003)	
UK		0.006	
		(0.010)	
$d_t\times$ UK		0.020	0.031*
		(0.020)	(0.016)
PI		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.004***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0005
		(0.0002)	(0.001)
Interest Rate		0.001	−0.030***
		(0.002)	(0.004)
Constant	0.349	−0.068	
	(0.307)	(0.050)	
Observations	34	1,054	1,054
R^2	0.290	0.210	0.175
Adjusted R^2	0.268	0.205	0.106
Residual Std. Error	1.708 (df = 32)	0.035 (df = 1046)	
F Statistic	13.069*** (df = 1; 32)	39.768*** (df = 7; 1046)	41.319*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.23: Effect of the SDLT holiday on the North West HPI YoY.

Estimating the effect of the SDLT holiday on the North West HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the North West and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	2.425***	−0.0003	
	(0.723)	(0.003)	
UK		0.012	
		(0.010)	
$d_t \times$ UK		0.010	0.021
		(0.020)	(0.016)
pi		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.004***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0004
		(0.0002)	(0.001)
Interest Rate		0.001	−0.030***
		(0.002)	(0.004)
Constant	0.243	−0.066	
	(0.215)	(0.050)	
Observations	34	1,054	1,054
R^2	0.260	0.210	0.174
Adjusted R^2	0.237	0.205	0.104
Residual Std. Error	1.195 (df = 32)	0.035 (df = 1046)	
F Statistic	11.254*** (df = 1; 32)	39.828*** (df = 7; 1046)	40.829*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.25: Effect of the SDLT holiday on the East Midlands HPI YoY.

Estimating the effect of the SDLT holiday on the East Midlands HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the East Midlands and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^N$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	1.968***	−0.0002	
	(0.566)	(0.003)	
UK		0.010	
		(0.010)	
$d_t \times$ UK		0.0001	0.011
		(0.020)	(0.016)
PI		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.004***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0004
		(0.0002)	(0.001)
Interest Rate		0.001	−0.030***
		(0.002)	(0.004)
Constant	0.171	−0.069	
	(0.168)	(0.050)	
Observations	34	1,054	1,054
R^2	0.274	0.209	0.173
Adjusted R^2	0.251	0.204	0.103
Residual Std. Error	0.937 (df = 32)	0.035 (df = 1046)	
F Statistic	12.078*** (df = 1; 32)	39.546*** (df = 7; 1046)	40.540*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.27: Effect of the SDLT holiday on the West Midlands HPI YoY.

Estimating the effect of the SDLT holiday on the West Midlands HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the West Midlands and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^N$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	2.430***	−0.001	
	(0.702)	(0.003)	
UK		−0.001	
		(0.010)	
$d_t \times$ UK		0.016	0.027*
		(0.020)	(0.016)
PI		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.004***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0004
		(0.0002)	(0.001)
Interest Rate		0.001	−0.030***
		(0.002)	(0.004)
Constant	0.190	−0.074	
	(0.209)	(0.050)	
Observations	34	1,054	1,054
R^2	0.272	0.209	0.175
Adjusted R^2	0.250	0.204	0.105
Residual Std. Error	1.161 (df = 32)	0.035 (df = 1046)	
F Statistic	11.974*** (df = 1; 32)	39.462*** (df = 7; 1046)	41.122*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.29: Effect of the SDLT holiday on the South West of England HPI YoY.

Estimating the effect of the SDLT holiday on the South West of England HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the South West of England and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	0.606	−0.001	
	(0.532)	(0.003)	
UK		−0.004	
		(0.010)	
$d_t \times$ UK		0.003	0.014
		(0.020)	(0.016)
PI		-0.002^{***}	0.002^{***}
		(0.0002)	(0.0003)
CPI		0.005^{***}	0.006^{***}
		(0.0005)	(0.001)
Employment		-0.002^{***}	0.0004
		(0.0002)	(0.001)
Interest Rate		0.001	-0.030^{***}
		(0.002)	(0.004)
Constant	−0.104	−0.079	
	(0.158)	(0.050)	
Observations	34	1,054	1,054
R^2	0.039	0.208	0.173
Adjusted R^2	0.009	0.203	0.103
Residual Std. Error	0.879 (df = 32)	0.035 (df = 1046)	
F Statistic	1.299 (df = 1; 32)	39.357^{***} (df = 7; 1046)	40.620^{***} (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.31: Effect of the SDLT holiday on the East of England HPI YoY.

Estimating the effect of the SDLT holiday on the East of England HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the East of England and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}^N_{1t}$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	0.857**	−0.001	
	(0.358)	(0.003)	
UK		−0.013	
		(0.010)	
$d_t\times$ UK		0.013	0.024
		(0.020)	(0.016)
PI		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.005***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0004
		(0.0002)	(0.001)
Interest Rate		0.001	−0.030***
		(0.002)	(0.004)
Constant	0.017	−0.085*	
	(0.106)	(0.050)	
Observations	34	1,054	1,054
R^2	0.152	0.210	0.174
Adjusted R^2	0.125	0.204	0.104
Residual Std. Error	0.592 (df = 32)	0.035 (df = 1046)	
F Statistic	5.735** (df = 1; 32)	39.642*** (df = 7; 1046)	40.955*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.32: Effect of the SDLT holiday on the South East of England HPI YoY.

Estimating the effect of the SDLT holiday on the South East of England HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the South East of England and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

	Dependent variable:		
	$Y_{1t} - \hat{Y}_{1t}^N$	Y_{it}	
	OLS	*OLS*	*FE*
	(1)	(2)	(3)
d_t	−0.096	−0.001	
	(0.344)	(0.003)	
UK		−0.027***	
		(0.010)	
$d_t\times$ UK		0.021	0.032**
		(0.020)	(0.016)
PI		−0.002***	0.002***
		(0.0002)	(0.0003)
CPI		0.005***	0.006***
		(0.0005)	(0.001)
Employment		−0.002***	0.0005
		(0.0002)	(0.001)
Interest Rate		0.001	−0.030***
		(0.002)	(0.004)
Constant	−0.032	−0.094*	
	(0.102)	(0.049)	
Observations	34	1,054	1,054
R^2	0.002	0.214	0.176
Adjusted R^2	−0.029	0.209	0.106
Residual Std. Error	0.568 (df = 32)	0.035 (df = 1046)	
F Statistic	0.078 (df = 1; 32)	40.695*** (df = 7; 1046)	41.379*** (df = 5; 971)
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A.34: Effect of the SDLT holiday on Greater London HPI YoY.

Estimating the effect of the SDLT holiday on the Greater London HPI YoY. Column (1) shows the estimation of equation 3.3.3, where the gap between the Greater London and its synthetic counterpart is regressed on a constant and a dummy variable. The standard errors reported in the parenthesis of column (1) are HAC-SE. Columns(2) and (3) display the estimation of the DiD model of equation 3.3.5, where the dependent variable is the HPI YoY is regressed on a constant, time dummy variable, UK dummy variable, the interaction of both dummy variables and a vector of economic covariates.

A.3 Chapter 4 Appendix

AP | Associated Press

Just 11 pct. of US long-term unemployed find jobs

Josh Boak, AP Economics Writer
20 March 2014

WASHINGTON (AP) -- A new study documents the bleak plight of Americans who have been unemployed for more than six months: Just 11 percent of them, on average, will ever regain steady full-time work.

The findings by three Princeton University economists show the extent to which the long-term unemployed have been shunted to the sidelines of the U.S. economy since the Great Recession. The long-term jobless number 3.8 million, or 37 percent of all unemployed Americans.

"The long-term unemployed are more than twice as likely" to stop looking for a job than to find one, according to the paper co-written by Alan Krueger, formerly President Barack Obama's chief economic adviser. "And when they exit the labor force, the long-term unemployed tend to say they no longer want a job."

During any given month from 2008 to 2012, barely more than one in 10 of the long-term unemployed had found full-time work. Their troubles were similar in states with high as well as low unemployment rates.

The analysis shows that a better predictor of hiring comes from the short-term unemployed, who are far more likely to be rehired.

Across the country, levels of short-term unemployment have essentially returned to pre-recession averages, even though the overall national unemployment rate remains historically high at 6.7 percent. The paper says that based on the number of short-term unemployed, further job gains could lead to "rising inflation and stronger real wage growth."

That's because the relative exclusion of the long-term unemployed means that employers must choose from among a limited supply of workers. That trend could push up prices.

Federal Reserve Chair Janet Yellen on Wednesday cited flat wages and low inflation running below the Fed's 2 percent target as evidence that the Fed should continue to support the recovery by keeping interest rates extremely low.

The number of people unemployed for more than six months has tripled since the recession began at the end of 2007. It peaked at 6.77 million in early 2010 and has declined to a still-high level as more unemployed workers have ended their job hunts and are no longer

Figure A.23: Topic modelling example.

Article from the Associated Press published on 20th of March 2014 with extreme negative sentiment.

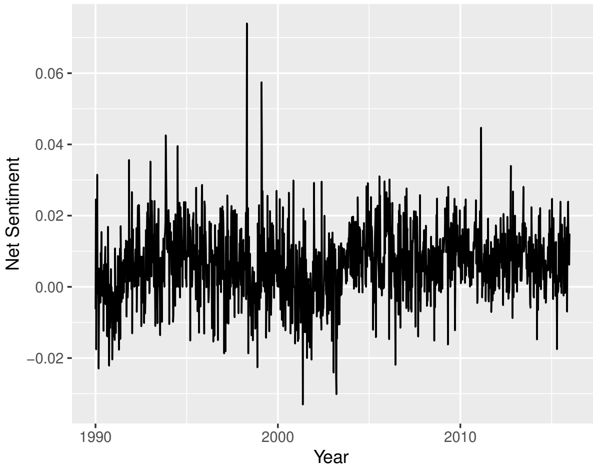

(a) Economics Topics

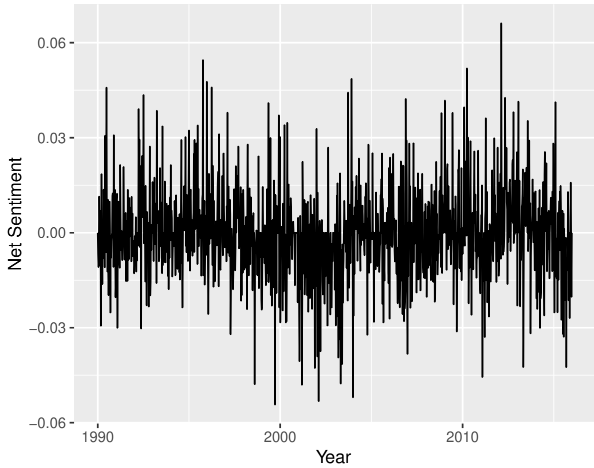

(b) Monetary Policy Topics

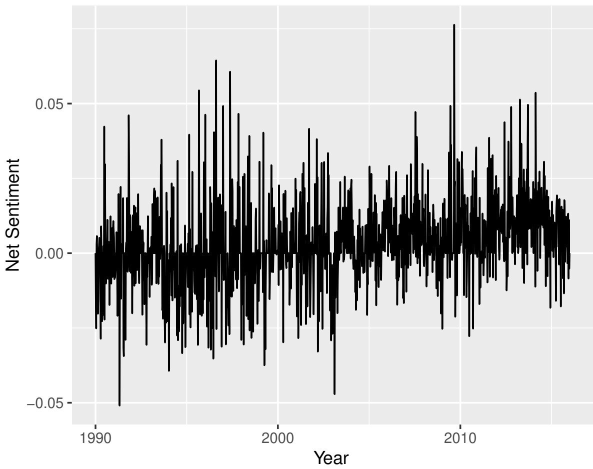

(c) Commodities Topics

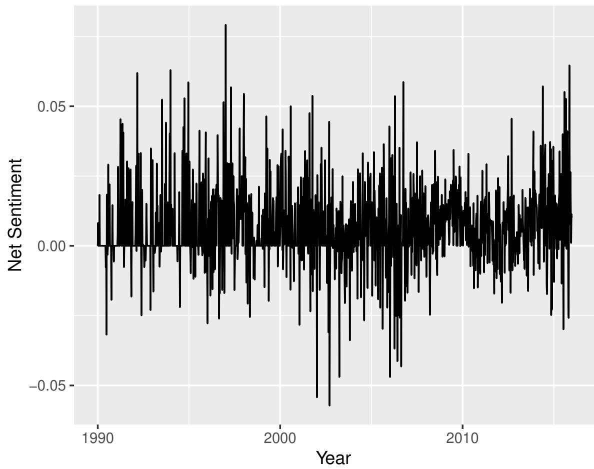

(d) Fiscal Policy Topics

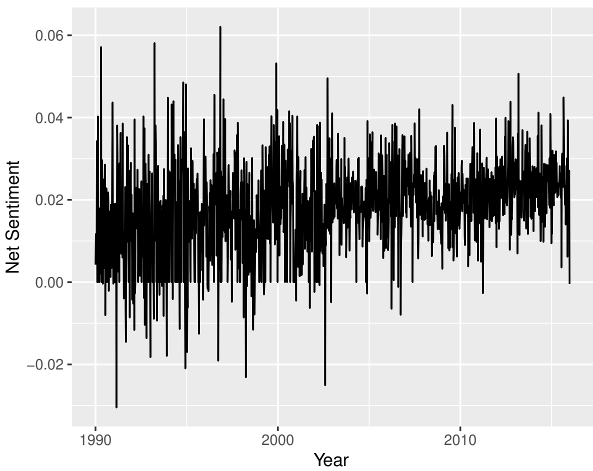

(e) Company Topics

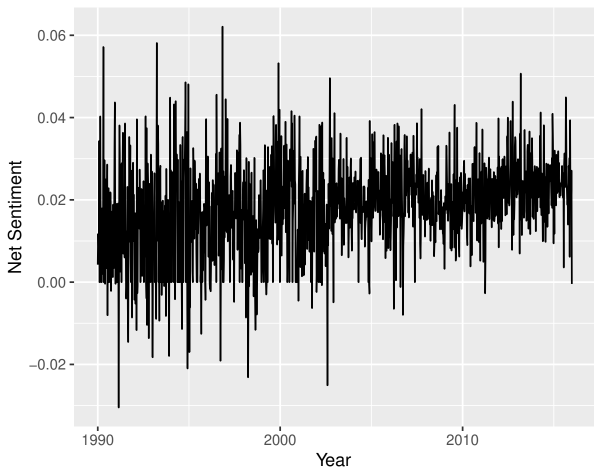

(f) Political Economy Topics

Figure A.24: Weekly net sentiment for six of the seventeen news topic.

The net sentiment is calculated as $s_{n,t} = \frac{N^{+}_{n,t} - N^{-}_{n,t}}{N^{+}_{n,t} + N^{-}_{n,t} + N^{0}_{n,t}}$. The weekly news sentiment is for the period 03/01/1990 to 28/12/2015.

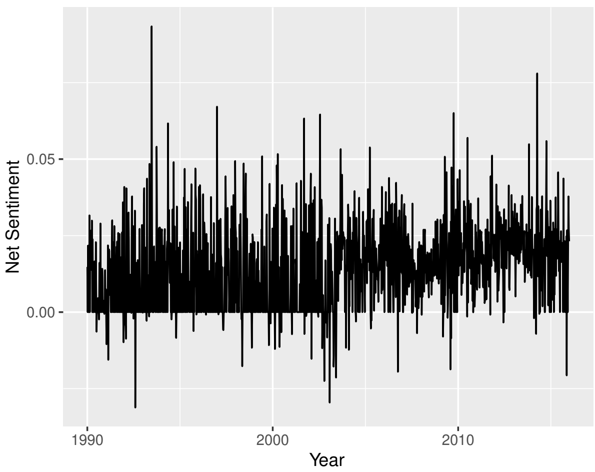

(a) Housing Market Topic

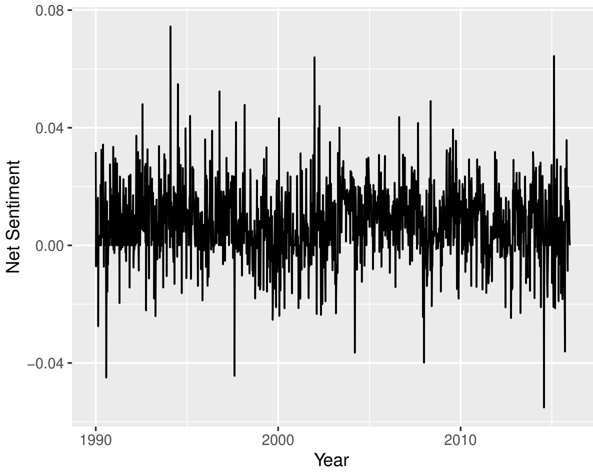

(b) Stock Market Topic

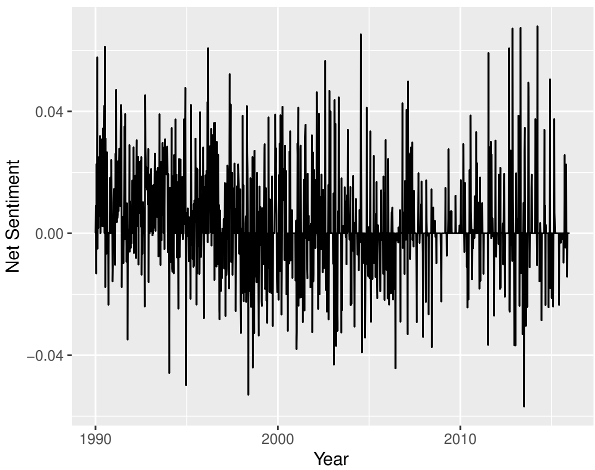

(c) Trade Topic

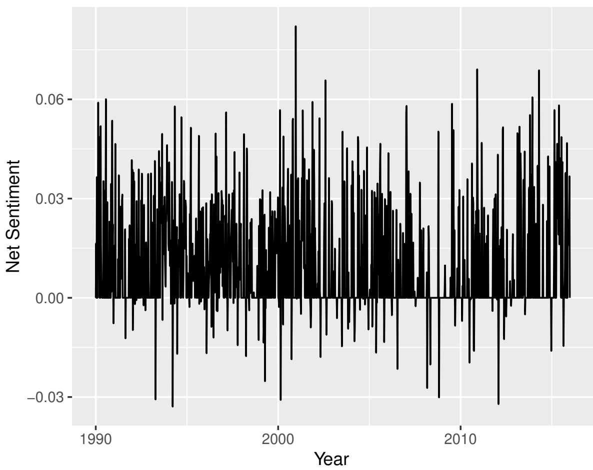

(d) Trade Policy Topic

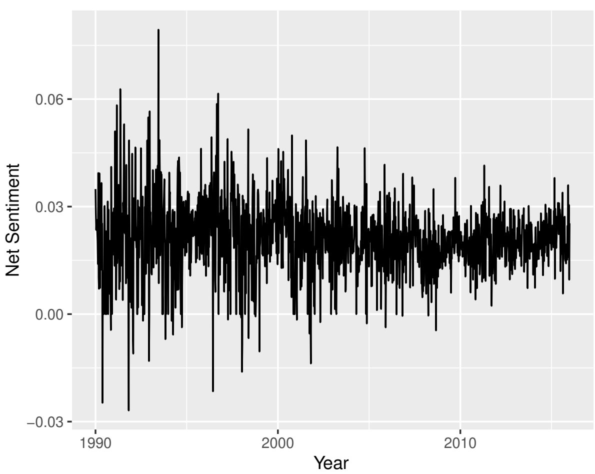

(e) Local Economy Topic

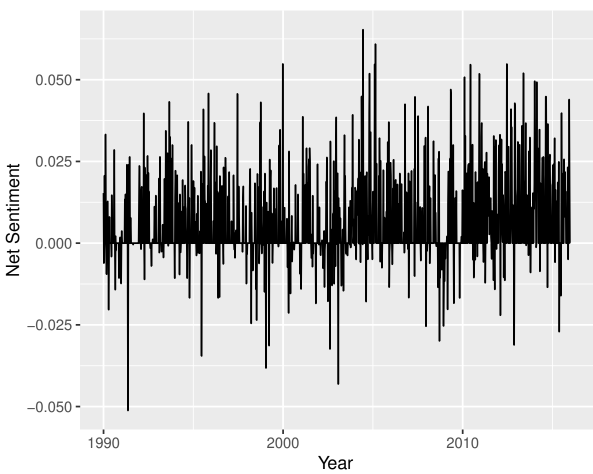

(f) Manufacturing Topic

Figure A.25: Weekly net sentiment for six of the seventeen news topic.

The net sentiment is calculated as $s_{n,t} = \frac{N^+_{n,t} - N^-_{n,t}}{N^+_{n,t} + N^-_{n,t} + N^0_{n,t}}$. The weekly news sentiment is for the period 03/01/1990 to 28/12/2015.

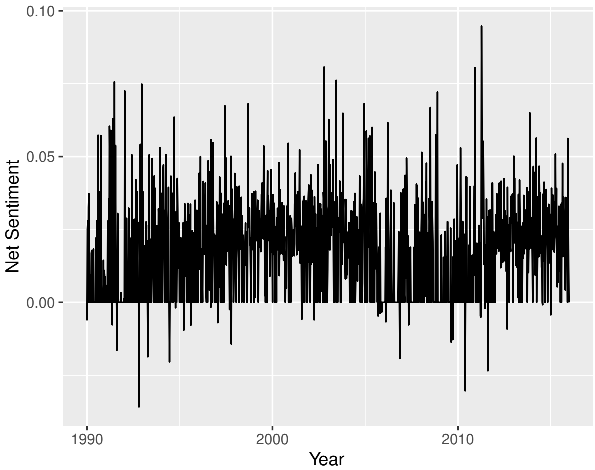

(a) Development Topic

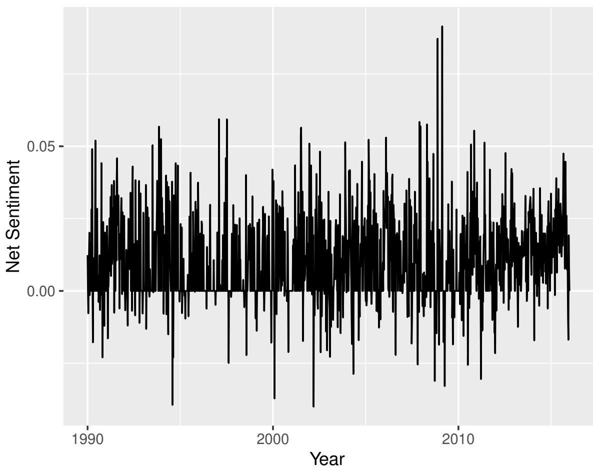

(b) Geo Political Topic

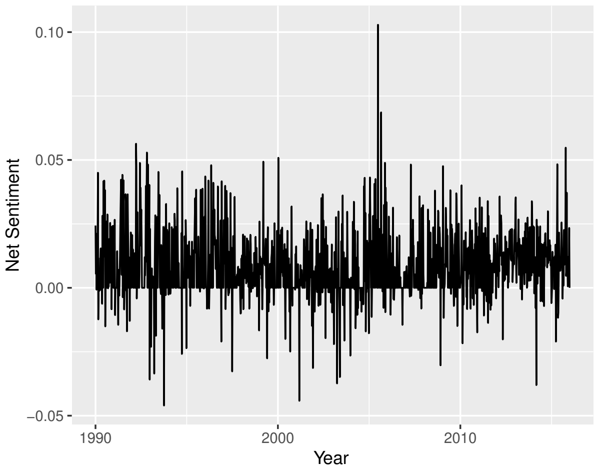

(c) Foreign Economy Topic

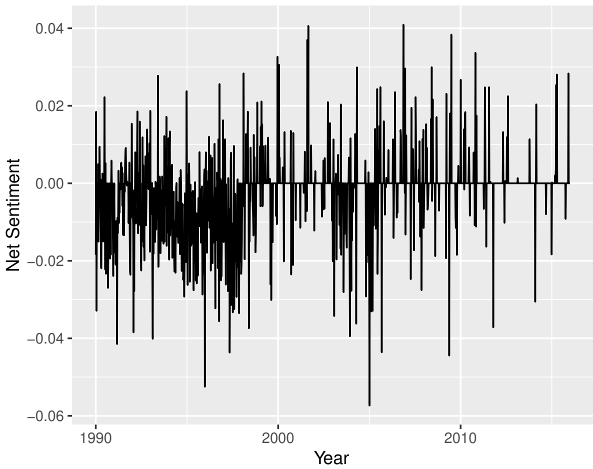

(d) Currency Topic

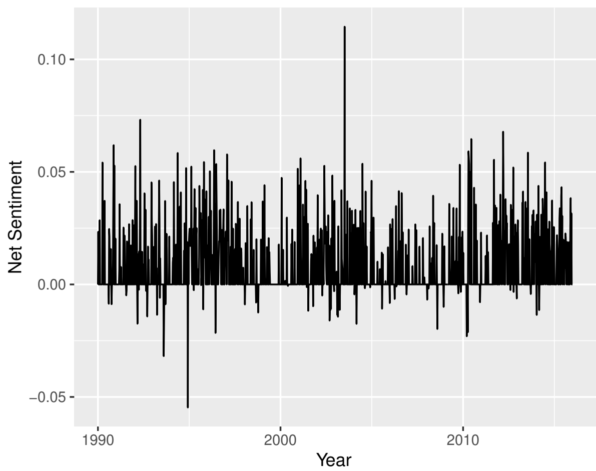

(e) Tourism Topic

Figure A.26: Weekly net sentiment for six of the seventeen news topic.

The net sentiment is calculated as $s_{n,t} = \frac{N^+_{n,t} - N^-_{n,t}}{N^+_{n,t} + N^-_{n,t} + N^0_{n,t}}$. The weekly news sentiment is for the period 03/01/1990 to 28/12/2015.

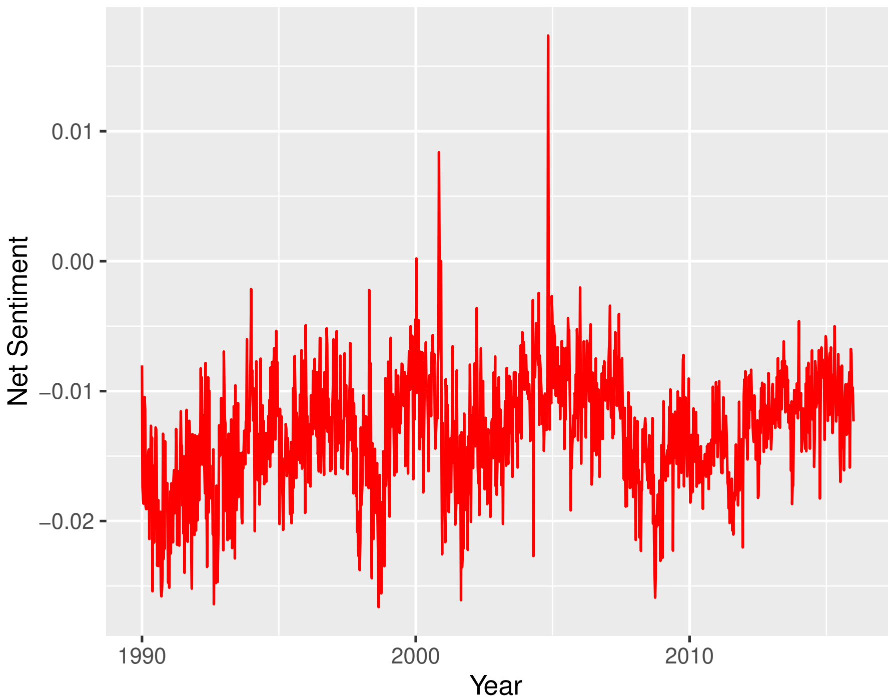

Figure A.27: News corpus net sentiment.

Loughran and McDonald (2011) Net sentiment for all the news corpus that combines 17 news topics on a weekly basis from 03/01/1990 to 28/12/2015. The net sentiment is calculated as $s_{n,t} = \frac{N^{+}_{n,t} - N^{-}_{n,t}}{N^{+}_{n,t} + N^{-}_{n,t} + N^{0}_{n,t}}$.

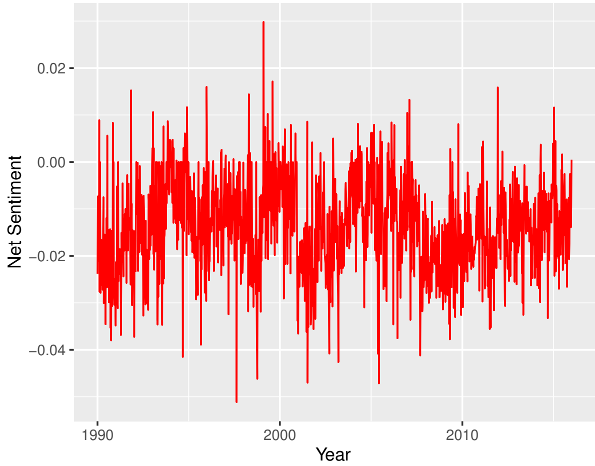

(a) Economics Topics

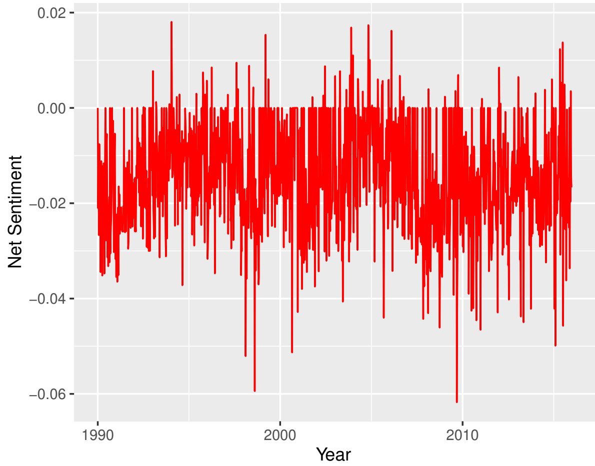

(b) Monetary Policy Topics

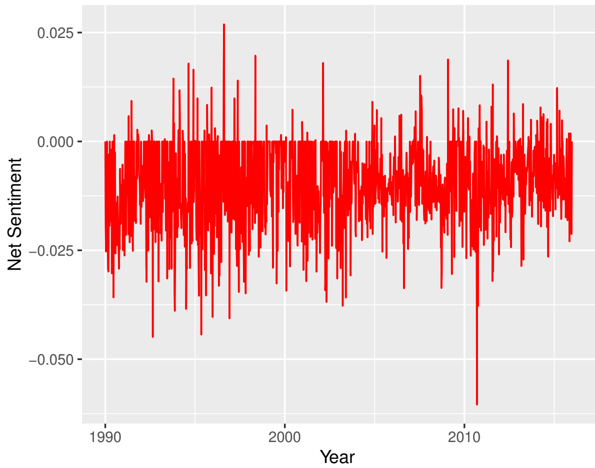

(c) Commodities Topics

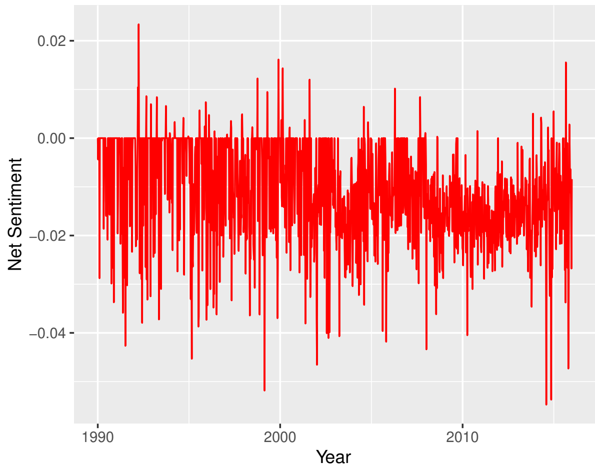

(d) Fiscal Policy Topics

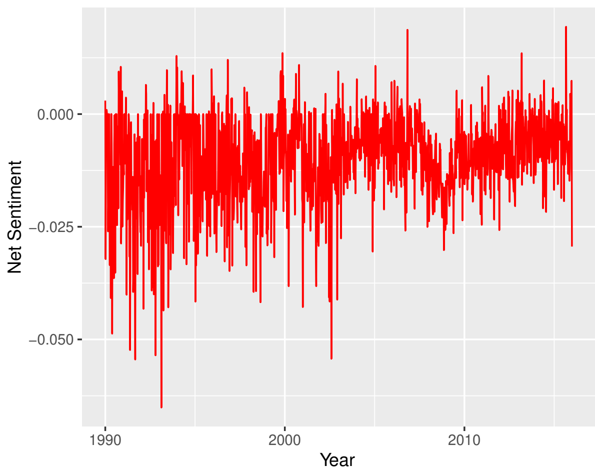

(e) Company Topics

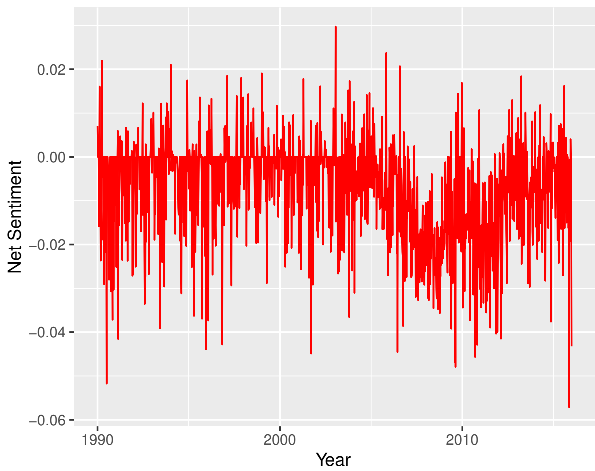

(f) Political Economy Topics

Figure A.28: Weekly net sentiment for six of the seventeen news topic using Loughran and McDonald (2011) dictionary.

The net sentiment is calculated as $s_{n,t} = \frac{N^{+}_{n,t} - N^{-}_{n,t}}{N^{+}_{n,t} + N^{-}_{n,t} + N^{0}_{n,t}}$. The weekly news sentiment is for the period 03/01/1990 to 28/12/2015.

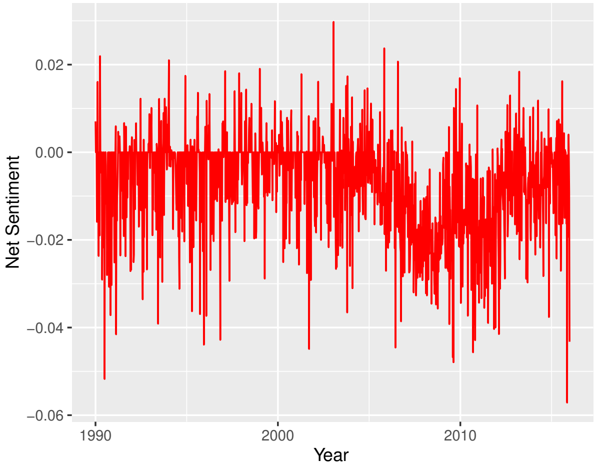

(a) Housing Market Topic

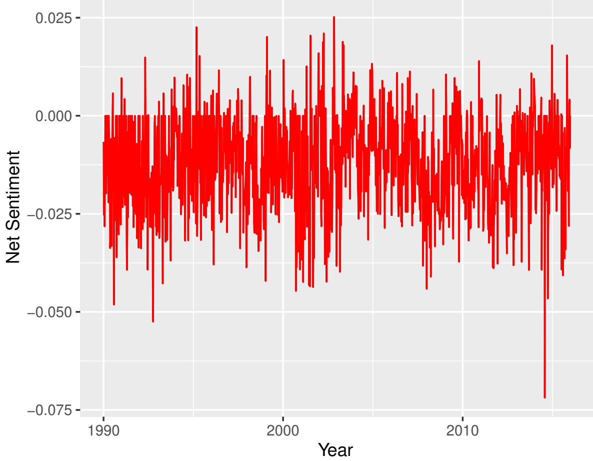

(b) Stock Market Topic

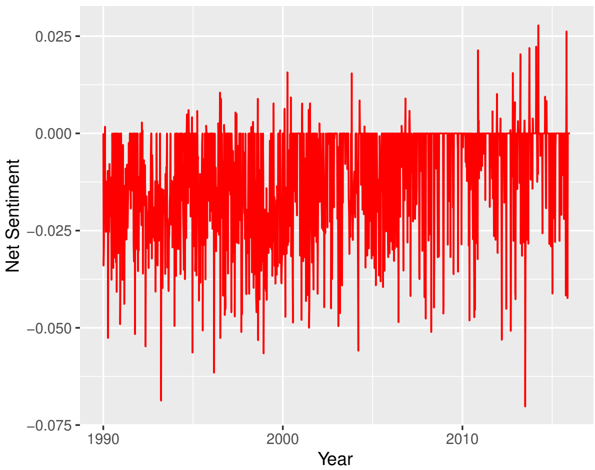

(c) Trade Topic

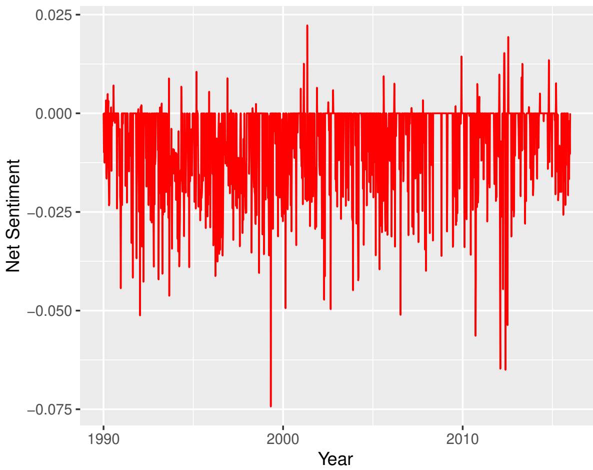

(d) Trade Policy Topic

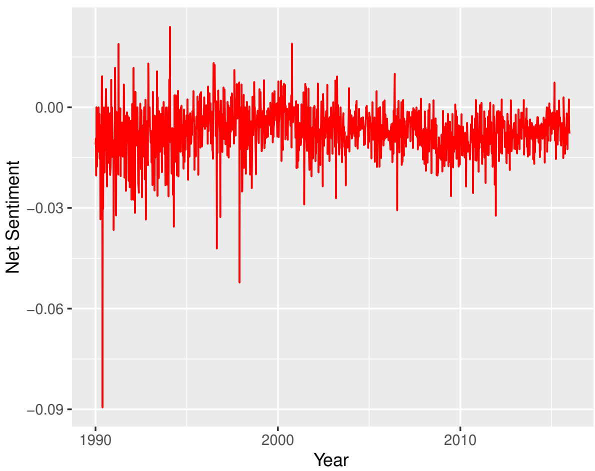

(e) Local Economy Topic

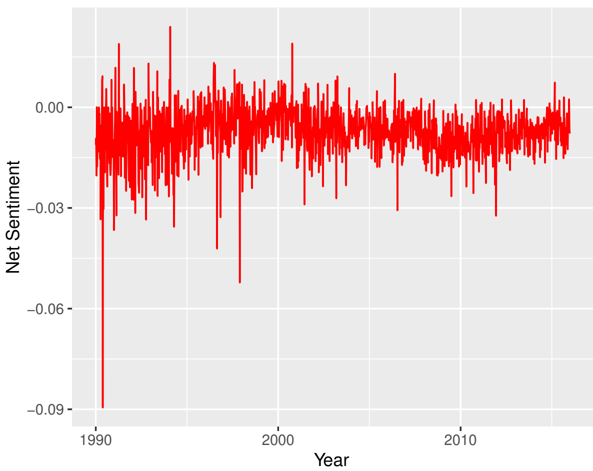

(f) Manufacturing Topic

Figure A.29: Weekly net sentiment for six of the seventeen news topic using Loughran and McDonald (2011) dictionary.

The net sentiment is calculated as $s_{n,t} = \frac{N^+_{n,t} - N^-_{n,t}}{N^+_{n,t} + N^-_{n,t} + N^0_{n,t}}$. The weekly news sentiment is for the period 03/01/1990 to 28/12/2015.

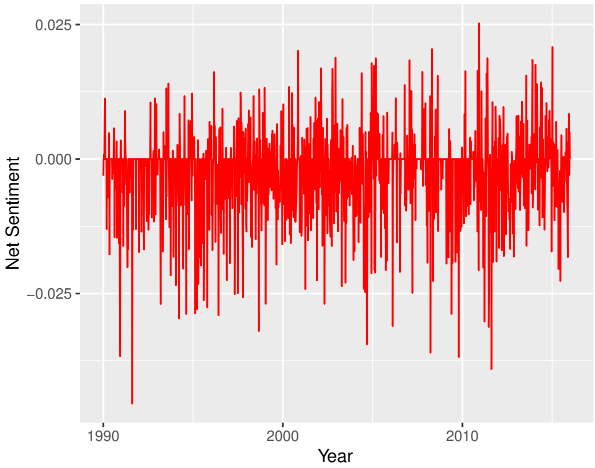

(a) Development Topic

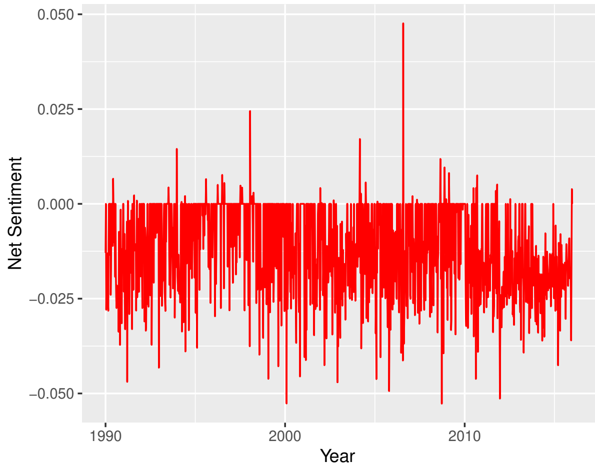

(b) Geo Political Topic

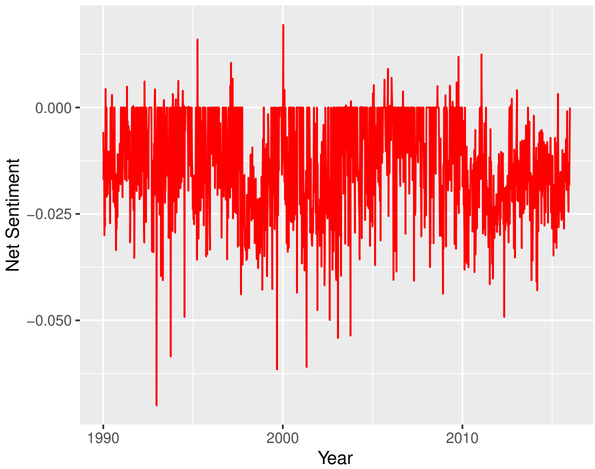

(c) Foreign Economy Topic

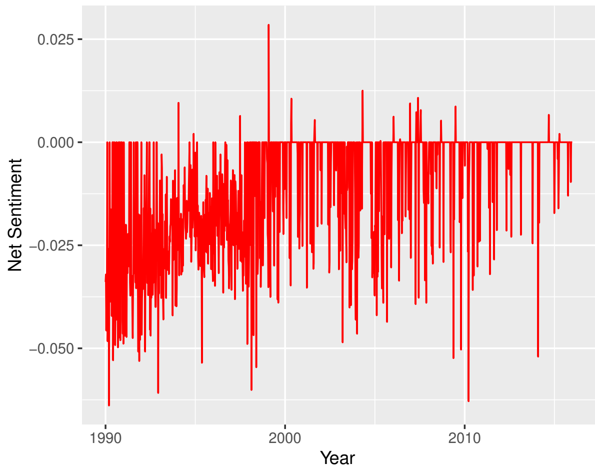

(d) Currency Topic

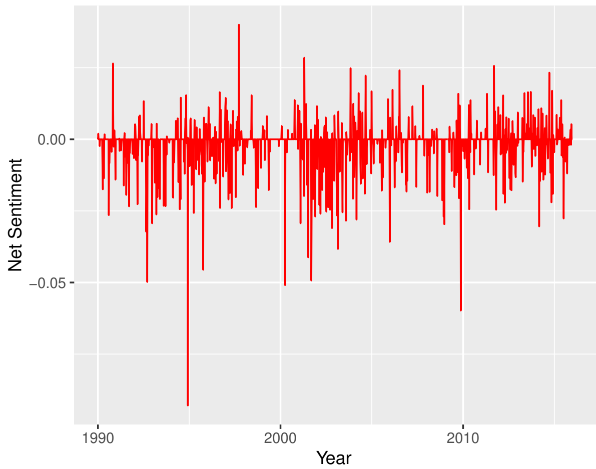

(e) Tourism Topic

Figure A.30: Weekly net sentiment for six of the seventeen news topic.

The net sentiment is calculated as $s_{n,t} = \frac{N^{+}_{n,t} - N^{-}_{n,t}}{N^{+}_{n,t} + N^{-}_{n,t} + N^{0}_{n,t}}$. The weekly news sentiment is for the period 03/01/1990 to 28/12/2015.

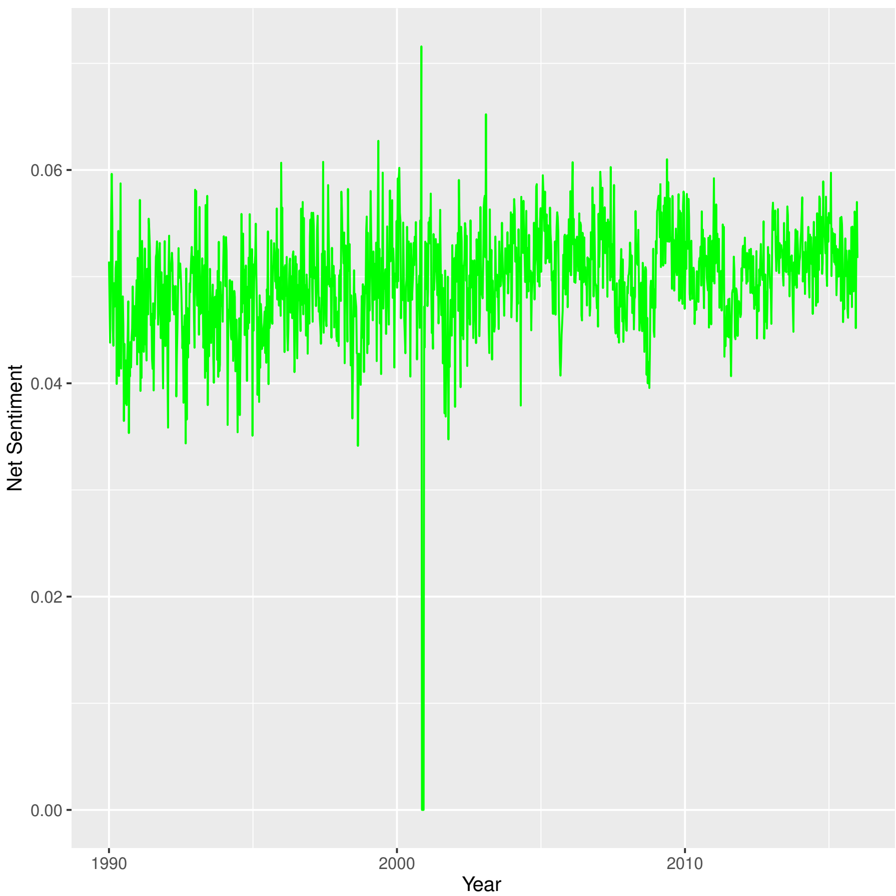

Figure A.31: News corpus net sentiment.

Baccianella et al. (2010) Net sentiment for all the news corpus that combines 17 news topics on a weekly basis from 03/01/1990 to 28/12/2015. This dictionary give different weights to different words based on their degree of positivity and negativity.

Figure A.32: Heat map correlation matrix for the news sentiment indices.

The Pearson correlation coefficient is written inside the squares. Correlation Coefficient is calculated as $r = \frac{\sum(x_i-\bar{x})(y_i-\bar{y})}{\sqrt{\sum(x_i-\bar{x})^2\sum(y_i-\bar{y})^2}}$.

	GARCH	Monetary Policy	Fiscal Policy	Politics
Mean equation				
μ	0.117	0.176	-0.291^{***}	0.083
	(0.300)	(0.066)	(0.073)	(0.303)
Variance equation				
α_0	5.729^{***}	5.536^{***}	5.843^{***}	6.266^{***}
	(1.954)	(1.930)	(1.300)	(2.129)
α_1	0.632^{**}	0.634^{***}	0.338^{***}	0.648^{**}
	(0.259)	(0.247)	(0.085)	(0.260)
β_1	0.116	0.113	0.289^{***}	0.104
	(0.127)	(0.141)	(0.022)	(0.127)
ψ		-0.324^{**}	0.600^{***}	-0.125^{***}
		(0.121)	(0.108)	(0.066)

Table A.35: Estimation of the GARCH models during the recession period.

I use the Berndt-Hall-Hall-Hausman (BHHH) optimisation algorithm to estimate the parameters of the LLF. The Bollerslev-Wooldridge standard errors (sandwich estimator) are reported in the parentheses. The significant levels are $^{***}p < 0.01,^{**} p < 0.05,^{*} p < 0.1$ respectively. The GARCH(1,1) and GARCH(1,1)-X models are estimated during the recession period from December 2007 to June 2009. Coefficients in the colour green indicate the statistical significance of the topical news sentiment variables in the GARCH-X equations.

	GARCH	Monetary Policy	Company	Currency	Construction	Tourism
Mean equation						
μ	0.317^{**}	0.307^{**}	0.356^{***}	0.313^{**}	0.217^{**}	0.324^{***}
	(0.101)	(0.010)	(0.102)	(0.100)	(0.107)	(0.097)
Variance equation						
α_0	0.499^{**}	0.507^{**}	0.212	0.426^{**}	3.787^{**}	0.653^{**}
	(0.223)	(0.209)	(0.206)	(0.19)	(1.385)	(0.241)
α_1	0.220^{**}	0.192^{***}	0.228^{**}	0.03^{***}	0.141	0.224^{***}
	(0.07)	(0.063)	0.073	(0.067)	(0.089)	(0.066)
β_1	0.683^{***}	0.714^{***}	0.681^{***}	0.714^{***}	0.0393^{*}	0.654^{***}
	(0.09)	(0.080)	(0.085)	(0.081)	(0.219)	(0.089)
ψ		-0.037^{***}	0.034^{**}	-0.152^{***}	-0.089^{***}	-0.054^{***}
		(0.013)	(0.017)	(0.048)	(0.013)	(0.011)

Table A.36: Estimation of the GARCH models during the post recession period.

Estimation of the GARCH models. I use the Berndt-Hall-Hall-Hausman (BHHH) optimisation algorithm to estimate the parameters of the LLF. The Bollerslev-Wooldridge standard errors (sandwich estimator) are reported in the parentheses. The significant levels are $^{***}p < 0.01,^{**} p < 0.05,^{*} p < 0.1$ respectively. The GARCH(1,1) and GARCH(1,1)-X models are estimated during the post recession period from July 2009 to December 2015. Coefficients in the colour green indicate the statistical significance of the topical news sentiment variables in the GARCH-X equations.

	GJR-GARCH	Monetary Policy	Trade	Currency
Mean equation				
μ	-0.399	-0.407	-0.393	-1.131^{***}
	(0.339)	(0.335)	(0.335)	(0.108)
Variance equation				
α_0	6.033^{**}	5.967^{***}	5.947^{***}	8.096^{***}
	(2.761)	(1.666)	(1.674)	(2.18)
α_1	0.072	0.056	0.06	-0.076
	(0.225)	(0.075)	(0.076)	(0.064)
γ	0.959^{*}	0.918^{**}	0.880^{*}	0.800
	(0.507)	(0.444)	(0.456)	(0.547)
β_1	0.143	0.149	0.148	0.331
	(0.235)	(0.141)	(0.146)	(0.203)
ψ		-0.025^{*}	1.278^{*}	0.813^{***}
		(0.770)	(0.750)	(0.125)

Table A.37: Estimation of the GJR-GARCH models during the recession period.

Estimation of the GJR-GARCH models. I use the Berndt-Hall-Hall-Hausman (BHHH) optimisation algorithm to estimate the parameters of the LLF. The Bollerslev-Wooldridge standard errors (sandwich estimator) are reported in the parentheses. The significant levels are $^{***}p < 0.01,^{**} p < 0.05,^{*} p < 0.1$ respectively. The GJR-GARCH(1,1) and GARCH(1,1)-X models are estimated during the recession period from December 2007 to June 2009. Coefficients in the colour green indicate the statistical significance of the topical news sentiment variables in the GJR-GARCH-X equations.

	GJR-GARCH	Monetary Policy	Commodities	Currency	Tourism
Mean equation					
μ	0.157^{*}	0.170^{*}	0.124	0.162^{*}	0.18^{**}
	(0.095)	(0.092	(0.093)	(0.094)	(0.093)
Variance equation					
α_0	0.371^{**}	0.410^{**}	0.445^{**}	0.337^{**}	0.519^{**}
	(0.158)	(0.145)	(0.157)	(0.140)	(0.179)
α_1	-0.01	-0.025	-0.014	-0.021	-0.005
	(0.052)	(0.052)	0.048	(0.048)	(0.062)
γ	0.346^{**}	0.339^{***}	0.325^{***}	0.333^{**}	0.354^{**}
	(0.107)	(0.103)	(0.091)	(0.097)	(0.116)
β_1	0.763^{***}	0.781^{***}	0.776^{***}	0.787^{***}	0.729^{***}
	(0.068)	(0.061)	(0.063)	(0.063)	(0.067)
ψ		-0.041^{***}	-0.057^{**}	-0.165^{***}	-0.052^{***}
		(0.100)	(0.017)	(0.047)	(0.010)

Table A.38: Estimation of the GJR-GARCH models during the post recession period.

Estimation of the GJR-GARCH models. I use the Berndt-Hall-Hall-Hausman (BHHH) optimisation algorithm to estimate the parameters of the LLF. The Bollerslev-Wooldridge standard errors (sandwich estimator) are reported in the parentheses. The significant levels are $^{***}p < 0.01,^{**} p < 0.05,^{*} p < 0.1$ respectively. The GJR-GARCH(1,1) and GARCH(1,1)-X models are estimated during the post recession period from July 2009 to December 2015.. Coefficients in the colour green indicate the statistical significance of the topical news sentiment variables in the GJR-GARCH-X equations.

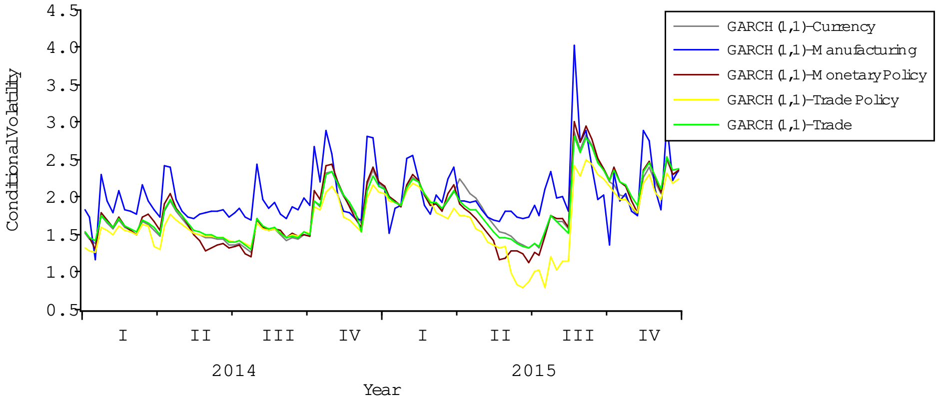

(a) GARCH(1,1)-X Conditional Volatility.

(b) GJR-GARCH(1,1)-X Conditional Volatility

Figure A.33: One step ahead forecasts for the conditional volatility for the GJR-GARCH(1,1)-X.

One step ahead forecasts for the conditional volatility, $\sqrt{h_t}$ for the GJR-GARCH(1,1)-X with the variables with significant in-sample coefficient for ψ.

A.3.1 GARCH Models Estimation Procedure

For the GARCH(p,q) models and the GJR-GARCH(p,q,r) with the exogenous variables discussed above, the vector of the unknown parameters θ of the GARCH models is estimated using the maximum likelihood estimation (MLE). The function of the parameters set is called the log-Likelihood Function (LLF)

$$L(\theta) = \sum_{t=1}^{T} -\frac{1}{2}\left(\log(2\pi) + \log(h_t) + \frac{(r_t - \mu)^2}{h_t}\right), \tag{A.3.1}$$

where T is the number of observations. A starting value for the parameters of the GARCH model $(\mu, \alpha_0, \alpha_1, \beta)$ should be chosen. The usual solution to parameter initialisation is to set parameter values in the conditional mean equation to that estimated using a "first pass" OLS regression (Brooks et al. 2001). Bollerslev (1986) suggests the use of the sample variance as an initialisation of the conditional variance.

$$h_t = \hat{u}_t^2 = \frac{\sum_{t=1}^{T} u_t^2}{T}.$$

Once the LLF is initialised, it can be maximised using numerical optimisation techniques. As in Awartani and Corradi (2005), I use the Berndt-Hall-Hall-Hausman (BHHH) algorithm for optimisation. This algorithm employs only the first derivative of the LLF with respect to the parameter values at each iteration i and an approximation to the second derivatives formed by the sum of outer product of the gradient (OPG) vectors. Although the approximation may be poor when the LLF is far from its maximum value, the BHHH algorithm can be modified by employing the Marquardt correction which adds a correction matrix to the sum of the Outer Product of the Gradient (OPG) vectors, the correction pushes the coefficient estimates more quickly to their optimal values.

$$\mathbf{B}(\hat{\theta}_i) = \sum_{t=1}^{T} \mathbf{G}_t(\hat{\theta}_i)\mathbf{G}'_t(\hat{\theta}_i) - aI, \tag{A.3.2}$$

where $\mathbf{G}_t(\hat{\theta})$ is the log-likelihood function for each observation, I is an identity matrix and a is a positive number chosen by the algorithm. This optimisation algorithm was also used by Charles and Darné (2019) when estimating the volatility for Bitcoin returns.

It is critical to know the distribution of η_t. From equation 4.3.3, η_t can be written as:

$$\eta_t = \frac{u_t}{\sqrt{h_t}}.$$

In other words, η_t is equivalent to the standardised residual. I test for the normality of η_t in a GARCH(1,1). The kurtosis is equal to 4.067 and the Jarque-Bera test for normality has a critical value of 97.343 and a p-value of 0.000, indicating that η_t is not normally distributed. Ignoring that η_t is not normally distributed will result in the parameter estimates being consistent, but inefficient. According to Campbell et al. (1997), one way to handle the problem of conditional non-normality is to use a different variance-covariance matrix estimator that is robust to non-normality such as that of Bollerslev and Wooldridge (1992). This estimation is known as Maximum Likelihood Estimation with Bollerslev-Wooldridge standard errors or Quasi-Maximum Likelihood (QML) "sandwich" estimator.[1] The covariance matrix is estimated using the matrix of the second order derivative of the LLF (the Hessian matrix, $\mathbf{H}(.)$) and the matrix of the outer products of the first order derivatives of the LLF ($\mathbf{B}(.)$):

$$\mathbf{Q}(\hat{\theta}) = \mathbf{H}(\hat{\theta}_{QML})^{-1}\mathbf{B}(\hat{\theta}_{QML})\mathbf{H}(\hat{\theta}_{QML})^{-1}, \tag{A.3.3}$$

where $\hat{\theta}_{QML}$ is the Quasi Maximum Likelihood estimator of θ and the coefficient standard errors as the robust Bollerslev-Wooldridge standard errors.

[1]Another way to deal with the conditional non-normality is to use an alternative distribution assumption such as the Student-t or Skewed Student-t distributions.

Bibliography

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, (2010). “Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program”. *Journal of the American Statistical Association* 105.490, pp. 493–505.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, (2015). “Comparative politics and the synthetic control method”. *American Journal of Political Science* 59.2, pp. 495–510.

Abadie, Alberto and Javier Gardeazabal, (Mar. 2003). “The Economic Costs of Conflict: A Case Study of the Basque Country”. *American Economic Review* 93.1, pp. 113–132.

Adämmer, Philipp and Rainer A Schüssler, (2020). “Forecasting the equity premium: mind the news!” *Review of Finance* 24.6, pp. 1313–1355.

Antweiler, Werner and Murray Z Frank, (2004). “Is all that talk just noise? The information content of internet stock message boards”. *The Journal of Finance* 59.3, pp. 1259–1294.

Ardia, David, Keven Bluteau, and Kris Boudt, (2019). “Questioning the news about economic growth: Sparse forecasting using thousands of news-based sentiment values”. *International Journal of Forecasting* 35.4, pp. 1370–1386.

Arnold, Mark, (2020). *The Fiscal Impacts of a Permanent Stamp Duty Holiday.* Tech. rep. Centre for Economics and Business Research.

Ashworth, John and Simon C Parker, (1997). “Modelling regional house prices in the UK”. *Scottish Journal of Political Economy* 44.3, pp. 225–246.

Athey, Susan and Guido W. Imbens, (May 2017). “The State of Applied Econometrics: Causality and Policy Evaluation”. *Journal of Economic Perspectives* 31.2, pp. 3–32.

Awartani, Basel MA and Valentina Corradi, (2005). "Predicting the volatility of the S&P-500 stock index via GARCH models: the role of asymmetries". *International Journal of Forecasting* 21.1, pp. 167–183.

Baccianella, Stefano, Andrea Esuli, and Fabrizio Sebastiani, (May 2010). "SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining". *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*. Valletta, Malta: European Language Resources Association (ELRA).

Bates, John M and Clive WJ Granger, (1969). "The combination of forecasts". *Journal of the Operational Research Society* 20.4, pp. 451–468.

Battu, Harminder, Ada Ma, and Euan Phimister, (2008). "Housing tenure, job mobility and unemployment in the UK". *The Economic Journal* 118.527, pp. 311–328.

Bayer, Christian, Benjamin Born, Ralph Luetticke, and Gernot J Müller, (2020). "The Coronavirus Stimulus Package: How large is the transfer multiplier?" *CEPR Discussion Paper No. DP14600*.

Ben-Rephael, Azi, Zhi Da, and Ryan D Israelsen, (2017). "It depends on where you search: Institutional investor attention and underreaction to news". *The Review of Financial Studies* 30.9, pp. 3009–3047.

Besley, Timothy, Neil Meads, and Paolo Surico, (2014). "The incidence of transaction taxes: Evidence from a stamp duty holiday". *Journal of Public Economics* 119, pp. 61–70.

Best, Michael Carlos and Henrik Jacobsen Kleven, (2018). "Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK". *The Review of Economic Studies* 85.1, pp. 157–193.

BFC, (2019). *The Impact of No Deal Brexit for British Fashion Industry*. Accessed on 26th May 2023. URL: https://www.britishfashioncouncil.co.uk/pressreleases/The-Impact-of-No-Deal-Brexit-for-British-Fashion-Industry-.

Bilgel, Fırat and Brian Galle, (2015). "Financial incentives for kidney donation: A comparative case study using synthetic controls". *Journal of Health Economics* 43, pp. 103–117.

Billmeier, Andreas and Tommaso Nannicini, (2013). "Assessing Economic Liberalization Episodes: A Synthetic Control Approach". *The Review of Economics and Statistics* 95.3, pp. 983–1001.

Blei, David M, (2012). "Probabilistic topic models". *Communications of the ACM* 55.4, pp. 77–84.

Blei, David M, Andrew Y Ng, and Michael I Jordan, (2003). "Latent dirichlet allocation". *The Journal of Machine Learning Research* 3, pp. 993–1022.

Bloom, Nicholas, (2009). "The impact of uncertainty shocks". *Econometrica* 77.3, pp. 623–685.

Bloom, Nicholas, Philip Bunn, Scarlet Chen, Paul Mizen, Pawel Smietanka, and Gregory Thwaites, (2019). "The impact of Brexit on UK firms".

Bloom, Nicholas, Philip Bunn, Paul Mizen, Pawel Smietanka, Greg Thwaites, and Garry Young, (2017). "Tracking the views of British businesses: evidence form the Decision Maker Panel". *Bank of England Quarterly Bulletin* 57.2, pp. 110–120.

Bloomberg, (Aug. 6, 2021). "Work From Home Is Driving U.K. Property Price Surge, BOE Says". *Bloomberg*. Accessed on 20th April 2022. URL: https://www.bloomberg.com/news/articles/2021-08-06/work-from-home-is-driving-u-k-property-price-surge-boe-says?leadSource=uverify%5C%20wall.

Boesel, Molly, Shu Chen, and Frank E Nothaft, (2021). "Housing preferences during the pandemic: effect on home price, rent, and inflation measurement". *Business Economics* 56.4, pp. 200–211.

Bollerslev, Tim, (1986). "Generalized autoregressive conditional heteroskedasticity". *Journal of Econometrics* 31.3, pp. 307–327.

Bollerslev, Tim and Jeffrey M Wooldridge, (1992). "Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances". *Econometric Reviews* 11.2, pp. 143–172.

Born, Benjamin, Gernot J Müller, Moritz Schularick, and Petr Sedláček, (2019). "The costs of economic nationalism: evidence from the Brexit experiment". *The Economic Journal* 129.623, pp. 2722–2744.

Botosaru, Irene and Bruno Ferman, (2019). "On the role of covariates in the synthetic control method". *The Econometrics Journal* 22.2, pp. 117–130.

Breinlich, Holger, Elsa Leromain, Dennis Novy, and Thomas Sampson, (Dec. 2019). *Exchange rates and consumer prices: evidence from Brexit*. CEP Discussion Papers dp1667. Centre for Economic Performance, LSE.

Breinlich, Holger, Elsa Leromain, Dennis Novy, and Thomas Sampson, (2020). "Voting with their money: Brexit and outward investment by UK firms". *European Economic Review* 124, p. 103400.

Brooks, Chris, Simon P Burke, and Gita Persand, (2001). "Benchmarks and the accuracy of GARCH model estimation". *International Journal of Forecasting* 17, pp. 45–56.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, (1997). *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.

Cao, Juan, Tian Xia, Jintao Li, Yongdong Zhang, and Sheng Tang, (2009). "A density-based method for adaptive LDA model selection". *Neurocomputing* 72.7, pp. 1775–1781.

Carlton, Dennis W, (1981). "The spatial effects of a tax on housing and land". *Regional Science and Urban Economics* 11.4, pp. 509–527.

Casadei, Patrizia and Simona Iammarino, (2021). "Trade policy shocks in the UK textile and apparel value chain: Firm perceptions of Brexit uncertainty". *Journal of International Business Policy* 4, pp. 262–285.

Cebula, Richard, (2009). "Are property taxes capitalized into housing prices in Savannah, Georgia? An investigation of the market mechanism". *Journal of Housing Research* 18.1, pp. 63–75.

Charles, Amélie and Olivier Darné, (2019). "The accuracy of asymmetric GARCH model estimation". *International Economics* 157, pp. 179–202.

Chemical Industries Association, (2018). *The Chemical Industry – Brexit priorities for UK growth*. Accessed on 20th May 2023. URL: https://www.cia.org.uk/Portals/0/Brexit%5C%20update_August%5C%2018.pdf.

Coombs, Christopher K, Nikias Sarafoglou, and William Crosby, (2012). "Property taxation, capitalization, and the economic implications of raising property taxes". *International Advances in Economic Research* 18.2, pp. 151–161.

Da, Zhi, Joseph Engelberg, and Pengjie Gao, (2011). "In search of attention". *The Journal of Finance* 66.5, pp. 1461–1499.

Davenport, Alexander and Peter Levell, (Jan. 2022). "Brexit and labour market inequalities: potential spatial and occupational impacts". *Oxford Review of Economic Policy* 38.1, pp. 50–67.

Dhingra, Swati, Hanwei Huang, Gianmarco Ottaviano, João Paulo Pessoa, Thomas Sampson, and John Van Reenen, (2017). "The costs and benefits of leaving the EU: trade effects". *Economic Policy* 32.92, pp. 651–705.

Diebold, Francis X and Robert S Mariano, (2002). "Comparing predictive accuracy". *Journal of Business & Economic Statistics* 20.1, pp. 134–144.

Engle, Robert F, (1982). "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation". *Econometrica*, pp. 987–1007.

Engle, Robert F and Andrew J Patton, (2007). "What good is a volatility model?" *Forecasting Volatility in the Financial Markets*. Elsevier, pp. 47–63.

Erken, Hugo, Raphie Hayat, Carlijn Prins, Marijn Heijmerikx, and Inge de Vreede, (2018). "Measuring The Permanent Costs of Brexit". *National Institute Economic Review* 244.1, R46–R55.

European-Commission, (2016). *European construction sector observatory*. Tech. rep. Country Profile Spain.

Farid, Moatazbellah, (2021a). "Land Tax Holiday and House Prices: Evidence from the UK". URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3949500 (visited on 10/25/2021).

Farid, Moatazbellah, (2021b). "The Winners and Losers of Brexit: A UK Manufacturing Sector Analysis". URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3765122 (visited on 01/13/2021).

Ferman, Bruno, Cristine Pinto, and Vitor Possebom, (2020). "Cherry picking with synthetic controls". *Journal of Policy Analysis and Management* 39.2, pp. 510–532.

Fetzer, Thiemo et al., (2021). “Subsidizing the Spread of Covid19: Evidence from the UK’s Eat-out To-help-out Scheme”. *The Economic Journal* 132.643, pp. 1200–1217.

Fetzer, Thiemo and Shizhou Wang, (2020). *Measuring the Regional Economic Cost of Brexit: Evidence up to 2019*. The Warwick Economics Research Paper Series (TWERPS) 1280. University of Warwick, Department of Economics.

Financial Times, (June 2018a). *Airbus future in UK threatened by ‘negative’ Brexit impact*. Financial Times. URL: https://www.ft.com/content/595220cc-759e-11e8-b326-75a27d27ea5f (visited on 05/25/2023).

Financial Times, (June 2018b). *Calls for changes to levy scheme after apprenticeship starts tumble*. Accessed on 27th May 2023. Financial Times. URL: https://www.ft.com/content/f8cfcade-6fac-11e8-92d3-6c13e5c92914.

Financial Times, (June 3, 2019). “Tory talk of no-deal Brexit costs jobs, warns industry”. *Financial Times*. URL: https://www.ft.com/content/49a0ebf0-8393-11e9-9935-ad75bb96c849.

Garcia, Diego, (2013). “Sentiment during Recessions”. *The Journal of Finance* 68.3, pp. 1267–1300.

Gardeazabal, Javier and Ainhoa Vega-Bayo, (2017). “An Empirical Comparison Between the Synthetic Control Method and HSIAO et al.’s Panel Data Approach to Program Evaluation”. *Journal of Applied Econometrics* 32.5, pp. 983–1002.

Gasiorek, Michael, Ilona Serwicka, and Alasdair Smith, (2019). “Which manufacturing industries and sectors are most vulnerable to Brexit?” *The World Economy* 42.1, pp. 21–56.

Gilchrist, David, Thomas Emery, Nuno Garoupa, and Rok Spruk, (2023). “Synthetic Control Method: A tool for comparative case studies in economic history”. *Journal of Economic Surveys* 37.2, pp. 409–445.

Glosten, Lawrence R, Ravi Jagannathan, and David E Runkle, (1993). “On the relation between the expected value and the volatility of the nominal excess return on stocks”. *The Journal of Finance* 48.5, pp. 1779–1801.

Griffiths, Thomas L and Mark Steyvers, (2004). “Finding scientific topics”. *Proceedings of the National Academy of Sciences* 101.suppl 1, pp. 5228–5235.

Grün, Bettina and Kurt Hornik, (2011). "topicmodels: An R Package for Fitting Topic Models". *Journal of Statistical Software* 40.13, pp. 1–30.

Grylls, George, (2020). *Brexit isn't done: what next for the UK chemicals industry?* New state man.

Hancké, Bob, (2018). *Made in the UK: Brexit and manufacturing revisited*. LSE European Politics and Policy (EUROPP) Blog. Accessed on 16th May 2023. URL: https://blogs.lse.ac.uk/europpblog/2018/07/06/made-in-the-uk-brexit-and-manufacturing-revisited/.

Hansen, Peter R. and Asger Lunde, (2005). "A forecast comparison of volatility models: does anything beat a GARCH(1,1)?" *Journal of Applied Econometrics* 20.7, pp. 873–889.

Harvey, David, Stephen Leybourne, and Paul Newbold, (1997). "Testing the equality of prediction mean squared errors". *International Journal of Forecasting* 13.2, pp. 281–291.

Highways England, (Nov. 2017). *£1.5bn A14 Cambridge to Huntingdon upgrade marks its first year of construction*. Accessed on 20th May 2023. URL: https://www.gov.uk/government/news/15bn-a14-cambridge-to-huntingdon-upgrade-marks-its-first-year-of-construction.

Hilber, Christian AL and Teemu Lyytikäinen, (2017a). *Stamp duty, mobility and the UK housing crisis*. URL: https://ideas.repec.org/p/cep/cepcnp/516.html.

Hilber, Christian AL and Teemu Lyytikäinen, (2017b). "Transfer taxes and household mobility: Distortion on the housing or labor market?" *Journal of Urban Economics* 101, pp. 57–73.

HM Government, (2017). *Industrial Strategy Building a Britain fit for the future*. Secretary of State for Business, Energy and Industrial Strategy.

House of Commons Business, Energy and Industrial Strategy Committe, (2018). *The Impact of Brexit on the Automotive Sector: Fifth Report of Session 2017-19*. HC 379. House of Commons. URL: https://publications.parliament.uk/pa/cm201719/cmselect/cmbeis/379/379.pdf.

Hsu, Yen-Ju, Yang-Cheng Lu, and J Jimmy Yang, (2021). "News sentiment and stock market volatility". *Review of Quantitative Finance and Accounting* 57.3, pp. 1093–1122.

IBIS World, (2020a). *Clothing Manufacturing in the UK*. IBIS world.

IBIS World, (2020b). *Motor Vehicle Manufacturing in the UK*. IBIS world.

Johnson, Paul, (2023). *Stamping on stamp duty would free empty nesters to fly their coops [Comment]*. Accessed: 9 June 2023. URL: https://ifs.org.uk/articles/stamping-stamp-duty-would-free-empty-nesters-fly-their-coops.

Kaplan, Greg, Benjamin Moll, and Giovanni L Violante, (2020). *The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the US*. Tech. rep. National Bureau of Economic Research.

Kierzenkowski, Rafal, Nigel Pain, Elena Rusticelli, and Sanne Zwart, (2016). "The economic consequences of Brexit: A taxing decision".

Kopczuk, Wojciech and David Munroe, (2015). "Mansion tax: The effect of transfer taxes on the residential real estate market". *American Economic Journal: Economic Policy* 7.2, pp. 214–57.

KPMG, (2017a). *Brexit: The Impact on Sectors*. Tech. rep. KPMG. URL: https://assets.kpmg.com/content/dam/kpmg/uk/pdf/2017/03/brexit-the-sector-impact.pdf.

KPMG, (2017b). *The Brexit Effect on EU Nationals*. Tech. rep. KPMG. URL: https://assets.kpmg.com/content/dam/kpmg/uk/pdf/2017/08/the-brexit-effect-on-eu-nationals.pdf.

Larsen, Vegard Høghaug and Leif Anders Thorsrud, (2019). "The value of news for economic developments". *Journal of Econometrics* 210.1, pp. 203–218.

Larsen, Vegard Høghaug and Leif Anders Thorsrud, (2022). "Asset returns, news topics, and media effects". *The Scandinavian Journal of Economics* 124.3, pp. 838–868.

Lawless, Martina and Edgar LW Morgenroth, (2019). "The product and sector level impact of a hard Brexit across the EU". *Contemporary Social Science* 14.2, pp. 189–207.

Liang, Chao, Linchun Tang, Yan Li, and Yu Wei, (2020). "Which sentiment index is more informative to forecast stock market volatility? Evidence from China". *International Review of Financial Analysis* 71, p. 101552.

Liang, Qin, Chunchun Hu, and Si Chen, (2021). "Evaluation of the Optimal Topic Classification for Social Media Data Combined with Text Semantics: A Case Study of Public Opinion Analysis Related to COVID-19 with Microblogs". *ISPRS International Journal of Geo-Information* 10.12.

Liu, Shimeng, (2015). "Spillovers from universities: Evidence from the land-grant program". *Journal of Urban Economics* 87, pp. 25–41.

Liu, Sitian and Yichen Su, (2021). "The impact of the Covid-19 pandemic on the demand for density: Evidence from the US housing market". *Economics letters* 207, p. 110010.

Loughran, Tim and Bill McDonald, (2011). "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks". *The Journal of Finance* 66.1, pp. 35–65.

Manela, Asaf and Alan Moreira, (2017). "News implied volatility and disaster concerns". *Journal of Financial Economics* 123.1, pp. 137–162.

Minford, Patrick, (2015). *Evaluating European trading arrangements*. eng. Cardiff Economics Working Papers E2015/17. Cardiff. URL: http://hdl.handle.net/10419/174106.

Morton, Alex, (2017). *Stamping Down Why Cutting Residential Stamp Duty is Easier than You Think*. Tech. rep. Centre for Policy Studies. URL: https://cps.org.uk/wp-content/uploads/2021/07/191026121449-CPSSDLT.pdf.

Nikita, M and N Chaney, (2020). "Tuning of the Latent Dirichlet Allocation Models Parameters [R Package Ldatuning Version 1.0. 2]". *Comprehensive R Archive Network (CRAN)*.

Nowotarski, Jakub, Eran Raviv, Stefan Trück, and Rafa⟨*Weron*, (2014). "An empirical comparison of alternative schemes for combining electricity spot price forecasts". *Energy Economics* 46, pp. 395–412.

Oates, Wallace E, (1969). "The effects of property taxes and local public spending on property values: An empirical study of tax capitalization and the Tiebout hypothesis". *Journal of Political Economy* 77.6, pp. 957–971.

ONS, (Jan. 2021). *Construction statistics, Great Britain: 2019*. Accessed on 5th June 2023. URL: https://www.ons.gov.uk/businessindustryandtrade/constructionindustry/articles/constructionstatistics/2019#:~:text=The%5C%20number%5C%20of%5C%20permanent%5C%20new,4.3%5C%25%5C%20to%5C%206%5C%2C070%5C%20in%5C%20Wales..

Opatrny, Matej, (2020). "The impact of the Brexit vote on UK financial markets: a synthetic control method approach". *Empirica*, pp. 1–29.

Oswald, Andrew, (1996). "A conjecture on the explanation for high unemployment in the industrialised nations: Part 1". *Warwick Economic Research Papers*.

Papyrakis, Elissaios, Lorenzo Pellegrini, and Luca Tasciotti, (2023). "Impacts of the Brexit referendum on UK employment: a synthetic control method approach". *Applied Economics Letters* 30.11, pp. 1407–1410.

Patton, Andrew J, (2011). "Volatility forecast comparison using imperfect volatility proxies". *Journal of Econometrics* 160.1. Realized Volatility, pp. 246–256.

Patton, Andrew J and Kevin Sheppard, (2009). "Evaluating volatility and correlation forecasts". *Handbook of Financial Time Series*. Springer, pp. 801–838.

Pellegrini, Lorenzo, Luca Tasciotti, and Andrea Spartaco, (2021). "A regional resource curse? A synthetic-control approach to oil extraction in Basilicata, Italy". *Ecological Economics* 185, p. 107041.

Peng, Lin and Wei Xiong, (2006). "Investor attention, overconfidence and category learning". *Journal of Financial Economics* 80.3, pp. 563–602.

Politico, (Apr. 2018). *1 in 5 UK manufacturing jobs at risk because of Brexit: study*. Accessed on 27th May 2023. URL: https://www.politico.eu/article/brexit-manufacturing-job-losses-chartered-institute-procurement-supply-cips/#:~:text=More%5C%20than%5C%20one%5C%20in%5C%20five,by%5C%20the%5C%20Observer%5C%20on%5C%20Sunday..

Puzanova, Yuliia and Mustafa Hakan Eratalay, (2021). "Effect of real estate news sentiments on the stock returns of Swedbank and SEB Bank". *Econometric Research in Finance* 6.2, pp. 77–117.

Ramani, Arjun and Nicholas Bloom, (May 2021). *The Donut Effect of Covid-19 on Cities.* Working Paper 28876. National Bureau of Economic Research.

Rapach, David E, Jack K Strauss, and Guofu Zhou, (2010). “Out-of-sample equity premium prediction: Combination forecasts and links to the real economy”. *The Review of Financial Studies* 23.2, pp. 821–862.

Rappaport, Jordan et al., (2021). “Hybrid Officing Will Shift Where People and Businesses Decide to Locate”. *Economic Bulletin* February 2, 2021, pp. 1–4.

Rosenthal, Stuart S., William C. Strange, and Joaquin A. Urrego, (2022). “JUE insight: Are city centers losing their appeal? Commercial real estate, urban spatial structure, and COVID-19”. *Journal of Urban Economics* 127. JUE Insights: COVID-19 and Cities, p. 103381.

Savills, (2021). *London Supply Update Q1 2021.* URL: https://www.savills.co.uk/research_articles/229130/314488-0 (visited on 05/25/2022).

Scanlon, Kath, Christine Whitehead, and Fanny Blanc, (2017). *A taxing question: Is Stamp Duty Land Tax suffocating the English housing market.* Tech. rep. LSE London Report for Family Building Society.

Scanlon, Kath, Christine Whitehead, and Fanny Blanc, (2021). *the stamp duty holiday.* Tech. rep. LSE London Report for Family Building Society.

Seely, A and M Keep, (2017). *Stamp duty land tax on residential property.* House of Commons Library briefing paper number 07050.

Sheppard, Kevin, (2019). *Univariate volatility modeling.* https://www.kevinsheppard.com/files/teaching/mfe/notes/financial-econometrics-2020-2021-chapter-7.pdf.

Shiller, Robert C, (2000). “Irrational exuberance”. *Philosophy and Public Policy Quarterly* 20.1, pp. 18–23.

Smales, Lee A, (2015). “Time-variation in the impact of news sentiment”. *International Review of Financial Analysis* 37, pp. 40–50.

Smales, Lee A, (2016). “Time-varying relationship of news sentiment, implied volatility and stock returns”. *Applied Economics* 48.51, pp. 4942–4960.

SMMT, (2014). *The UK Automotive Industry and the EU: An economic assessment of the interaction of the UK's Automotive Industry with the European Union*. Accessed on 23th May 2023. URL: https://www.smmt.co.uk/wp-content/uploads/sites/2/SMMT-KPMG-EU-Report.pdf.

Stearns, Jenna, (2015). "The effects of paid maternity leave: Evidence from Temporary Disability Insurance". *Journal of Health Economics* 43, pp. 85–102.

Tetlock, Paul C, (2007). "Giving content to investor sentiment: The role of media in the stock market". *The Journal of Finance* 62.3, pp. 1139–1168.

Tetlow, Gemma and Alex Stojanovic, (2018). *Understanding the economic impact of Brexit*. Tech. rep. Institute for Government, pp. 2–76.

The Guardian, (Mar. 2017). *Brexit threatens UK car industry, says BMW boss*. The Guardian. URL: https://www.theguardian.com/business/2017/mar/03/brexit-uk-car-industry-mini-britain-eu (visited on 05/25/2023).

The Guardian, (Mar. 2018). *Industry planning jobs cull to offset Brexit*. URL: https://www.theguardian.com/politics/2018/mar/31/manufacturing-firms-plan-job-cuts-after-brexit (visited on 05/16/2023).

The Guardian, (Oct. 2021). *UK Car Industry Says Brexit Rules are Denting Competitiveness*. Accessed on 12th May 2023. The Guardian. URL: https://www.theguardian.com/business/2021/oct/12/uk-car-industry-brexit-rules-smmt-energy-prices-electric-vehicle#:~:text=Cars%5C%20remain%5C%20the%5C%20UK%5C%27s%5C%20No,sector%5C%27s%5C%20trade%5C%20body%5C%20has%5C%20said..

The New York Times, (Jan. 2021). *For One British Industry, Brexit's Red Tape Is Just Beginning*. Accessed on 29th May 2023. New York Times. URL: https://www.nytimes.com/2021/01/18/business/brexit-chemical-industry-uk.html.

The UK in a Changing Europe, (2020). *Manufacturing and Brexit*. Technical Report. The UK in a Changing Europe.

Uhl, Matthias W, (2014). "Reuters sentiment and stock returns". *Journal of Behavioral Finance* 15.4, pp. 287–298.

UKFT, (2019). *No-deal Brexit could cost UK fashion and textile industry £870m.* Accessed on 26th May 2023. URL: https://www.ukft.org/no-deal-brexit-could-cost-uk-fashion-and-textile-industry-870m/.

UKFT, (2021). *Brexit bureaucracy hits UK fashion and textile companies hard.* Accessed on 26th May 2023. URL: https://www.ukft.org/brexit-bureaucracy-survey-july-2021/.

Van Leuvensteijn, Michiel and Pierre Koning, (2004). "The effect of home-ownership on labor mobility in the Netherlands". *Journal of Urban Economics* 55.3, pp. 580–596.

Van Ommeren, Jos and Michiel Van Leuvensteijn, (2005). "New evidence of the effect of transaction costs on residential mobility". *Journal of Regional Science* 45.4, pp. 681–702.

Villar, Paul Fenton and Elissaios Papyrakis, (2017). "Evaluating the impact of the Extractive Industries Transparency Initiative (EITI) on corruption in Zambia". *The Extractive Industries and Society* 4.4, pp. 795–805.

Weiss, Christoph E., Eran Raviv, and Gernot Roetzer, (2018). "Forecast Combinations in R using the ForecastComb Package". *The R Journal* 10.2, pp. 262–281.

Wells, James and Mark Stephens, (2017). *The impact of sterling depreciation on prices and turnover in the UK manufacturing sector: 2017.* Tech. rep. Office of National Statistics.

Zakoian, Jean-Michel, (1994). "Threshold heteroskedastic models". *Journal of Economic Dynamics and Control* 18.5, pp. 931–955.

Zhuang, Hong, Miao Grace Wang, Imre Ersoy, and Mesut Eren, (2023). "Does joining the European monetary union improve labor productivity? A synthetic control approach". *Journal of Productivity Analysis* 59, pp. 287–306.

Zoopla, (2020a). *Revealed: the top 10 search terms most used by homehunters in 2020.* Accessed on 20th April 2023. URL: https://www.zoopla.co.uk/discover/property-news/top-10-search-terms-most-used-by-homehunters-in-2020/.

Zoopla, (2020b). *Top 10 rural hotspots: the country locations buyers are flocking to during lockdown.* Accessed on 12th April 2023. URL: https://www.zoopla.co.

uk/discover/property-news/top-10-rural-hotspots-lockdown-ryedale-yorkshire-cotswolds-kent-where-buyers-are-flocking-to/.