

Do Labels Polarise? Theory and Evidence from the Brexit Referendum

Su-Min Lee Alexandru Savu

Abstract

Why has geographical political polarisation increased in recent times? We propose a theoretical social learning mechanism whereby policy preferences become more homogeneous within geographical units, yet increasingly heterogeneous between units over time as voters become better informed on the views of those in their vicinity. To study our model's predictions, we exploit the delayed implementation of Brexit and its salience in the elections following the 2016 referendum. Analysing constituency-level longitudinal-data, we find that voters updated their Brexit views after observing the referendum's local results, and acted upon their new beliefs in the following elections. We document a two percentage-point relative decrease in the (anti-Brexit) Liberal Democrat vote share in constituencies where Leave narrowly won, mirrored by an increase for the Conservatives. Our findings have implications for how group-based identities form more broadly.

Reference Details

CWPE	2227
Published	11 April 2022
Key Words	Elections, Brexit, Local Contextual Effects, Information, Social Learning, Political Attitudes
JEL Codes	D71, D72
Website	www.econ.cam.ac.uk/cwpe

Do Labels Polarise? Theory and Evidence from the Brexit Referendum *

Su-Min Lee and Alexandru Savu

April 11, 2022

Abstract

Why has geographical political polarisation increased in recent times? We propose a theoretical social learning mechanism whereby policy preferences become more homogeneous within geographical units, yet increasingly heterogeneous between units over time as voters become better informed on the views of those in their vicinity. To study our model's predictions, we exploit the delayed implementation of Brexit and its salience in the elections following the 2016 referendum. Analysing constituency-level longitudinal-data, we find that voters updated their Brexit views *after* observing the referendum's local results, and acted upon their new beliefs in the following elections. We document a two percentage-point relative decrease in the (anti-Brexit) Liberal Democrat vote share in constituencies where Leave narrowly won, mirrored by an increase for the Conservatives. Our findings have implications for how group-based identities form more broadly.

Keywords: Elections, Brexit, Local Contextual Effects, Information, Social Learning, Political Attitudes

*We wish to thank Toke Aidt for his guidance and advice. Contact: smlee@londoneconomics.co.uk (SML), ams269@cam.ac.uk (AS); Christ's College and Faculty of Economics, University of Cambridge, Austin Robinson Building, Sidgwick Ave, Cambridge CB3 9DD (AS). Alexandru Savu acknowledges support from the Gates Cambridge Trust.

1 Introduction

Since geography forms the basis of political representation in most democratic systems (Cutler 2007), it comes as no surprise that political views are strongly spatially-clustered (Cho and Rudolph 2008). An explanation comes from a well-established literature on contextual effects, which asserts that one's beliefs are causally influenced by the political preferences of those in one's vicinity¹ (Cox 1969; Books and Prysby 1988; Cantoni and Pons 2021). According to this argument, people assimilate to their environment's political orientation² (Gallego et al. 2016), thus shaping a non-random geographical distribution of beliefs.

Less understood, however, is why the spatial clustering of political preferences appears to have risen in a remarkable manner recently in certain democracies (Brown and Enos 2021). In the United States, for instance, Kaplan et al. (2020) show that geographical political divisions have "increased dramatically" in the last seventy years. Echoing this observation, it is not uncommon to hear about "red areas getting redder, while the blue areas become bluer" in the popular media³, highlighting increasing geographical polarisation (Johnston et al. 2020). Complementing these insights, in the United Kingdom, an emerging literature on the country's withdrawal from the European Union asserts that Brexit may have been both the result, as well as a cause of rising political polarisation, both geographical and affective (Ford and Goodwin 2017; Jennings and Stoker 2019; Axe-Browne and Hansen 2020; Hobolt et al. 2020; Pickard et al. 2022).

Hence, it is natural to ask: *What factors exacerbate geographical political polarisation?* In this paper, we add to the debate by investigating whether polarisation partially⁴ results from more readily available information on *local* political preferences, and from the manner in which this information is communicated to the electorate.

Expanding upon the contextual effects literature, we investigate the following core mechanism: when deciding whether or not to support a policy, voters internalize the preferences of their local geographic group before making their choice. Intuitively, by

¹Although the focus of this article is geographical, and therefore "vicinity" may intuitively be interpreted in a physical sense, group-based identity is a broader notion that often encompasses other dimensions (e.g. ethnicity, Branton and Jones 2005).

²Through mechanisms such as community interactions (Huckfeldt and Sprague 1987) or strategic sorting (Gimpel and Hui 2015).

³E.g. <https://lat.ms/2SgvDao>.

⁴Importantly, we do not argue that local information is the *sole* determinant of polarisation. Indeed, existing work (Cho and Rudolph 2008; Gallego et al. 2016) has identified both individual-level characteristics such as income, education and race, and district-level variables such as inequality as important explanatory variables for the sorting and assimilation processes shaping the spatial clustering of preferences. Rather, our objective is precisely to *isolate* the effect of local information from the effects of these other factors. We do so empirically by looking at the local results of the 2016 Brexit Referendum as an information shock *orthogonal* to other influences.

learning about their community's position, voters' uncertainty in their assessment of the policy's expected effects is reduced (Bala and Goyal 1998). Then, if the voters find that those in their vicinity largely support (oppose) a certain policy, the likelihood that they themselves will see the policy as favourable increases (decreases), ultimately translating to increased within-community preference homogeneity and larger between-communities heterogeneity - the definition for geographic polarisation we employ throughout (as in Kaplan et al. 2020).

Therefore, as information on contextual preferences becomes increasingly accessible over time - for instance, due to better communication means (Zhuravskaya et al. 2020), or (our focus) the presentation of local policy support numbers in the media (Levendusky 2013) - more social learning takes place, resulting in higher rates of assimilation, and ultimately a *ceteris paribus* increase in the spatial-clustering of policy preferences.⁵

Of course, we are not the first to argue that people use contextual information to guide their decisions. In fact, social learning is at the core of most contextual influences narratives and the associated assimilation process (Converse 1964). To this point, Burbank (1997) explains that individuals might use their perception of others' political views as "shortcuts" in forming their own policy stances, while Bala and Goyal (1998) argue that when payoffs are unknown, people use the experiences of their neighbours to make choices. More recently, Cutler (2007) similarly argues that citizens will take local beliefs into account when forming their own political attitudes. Despite these views, however, little attention has been given to assessing how new information on local preferences causally influences political choices and exacerbates geographical polarisation.

We empirically assess the effects of contextual information signals on policy preferences by investigating how learning about the results of the 2016 United Kingdom European Union Membership Referendum (henceforth, "the referendum") affected *subsequent* local-level voting patterns.

Besides its topical importance, this setting is appropriate for assessing our theoretical channel for several reasons. First, a prerequisite of the social learning mechanism is for the investigated shock to be informative. Given that the referendum was the first of its kind in forty years (Evans and Menon 2017), we expect its results to provide novel information on local preferences, leading to voters updating their beliefs concerning their reference group's Brexit stance.

Second, to test whether a contemporaneous contextual information signal shapes future behaviour, we have to be able to track support for the policy after the shock occurs. While a second referendum did not take place, we make progress by exploiting two additional appealing features of the Brexit process: *i*) even after the referendum took place, uncer-

⁵Section 3 presents a simple model detailing this mechanism.

tainty concerning whether Brexit itself would ultimately be carried out (and in what form - either a "softer" or "harsher" Brexit) continued to exist in subsequent years (Graziano et al. 2021), and *ii*) the way in which Brexit would be implemented was perceived to depend on which parties wielded power (Mellon et al. 2018), with clear differences existing between the positions of two parties: the unambiguously anti-Brexit Liberal Democrats and the largely⁶ pro-Brexit Conservative Party.

In our empirical analysis, because of their stark anti-Brexit position and the manifesto the Liberal Democrats put forth in the post-referendum elections⁷, we use the vote share obtained by this party to operationalize *post-referendum local opposition to Brexit* as our main dependent variable.

Then, in light of the mechanism wherein people internalize the views of those in their vicinity when deciding whether to support Brexit, we expect support for the Liberal Democrats to decrease (increase) *after learning about the referendum's results* in areas where Leave (Remain) won majority support, mirrored by an increase (decrease) for the Conservatives. This is our main hypothesis.

That said, simply finding a correlation between the referendum's local results and vote shares is unlikely to capture a causal relationship, due to potential issues of reverse causality or omitted variables (Gallego et al. 2016). To address this fundamental issue, we proceed in a novel manner: rather than relying on a continuous measure for the local referendum results (Hanretty 2017), we construct a dichotomous treatment variable equal to one in areas where Leave support was above fifty percent (zero otherwise).

Henceforth, we refer to treated (control) constituencies as Leave (Remain)-labelled, broadly in line with Hobolt et al. (2020).

We assume⁸ that a discontinuous change in the information contained by the local referendum outcome takes place at the fifty percent threshold. Concretely, in Leave-labelled units, voters were informed that a "majority", as opposed to a "minority" of people in their geographic reference group supports Brexit, which impacted both the way in which the signal was disseminated by media outlets, and the manner in which it was internalized by the electorate.

Then, to empirically test our hypothesis, we use longitudinal constituency-level electoral data (Allison 1994) and contrast how voting patterns changed *after* learning about the referendum's results when comparing Leave and Remain-labelled areas. To increase

⁶Although the Conservatives adopted a tough stance on most Brexit-related issues, ambiguity existed in their messaging (Hobolt and Rodon 2017). That said, analysing the impact of social learning on the performance of the Conservatives following the referendum is still informative as they were largely perceived as the parliamentary party most capable and willing to deliver Brexit *relatively* to other parties. We discuss this further below.

⁷Section 2.

⁸And provide evidence in Section 5.2 corroborating this assumption.

the causal reading of our estimates (Bertrand et al. 2004), we restrict our analysis on constituencies tightly clustered around the fifty-percent Leave threshold.⁹

We find the following: First, consistent with the local contextual learning hypothesis, we find that support for the Liberal Democrats in the 2017 general elections fell significantly in Leave-labelled constituencies relative to the trend captured in Remain-labelled ones. Quantitatively, the estimate retrieved in our preferred specification suggests an average decrease of two percentage points, a politically meaningful impact considering the party’s overall score of seven percent. Further in line with the theory, we show that the Conservatives gained ground. All these findings are reassuringly qualitatively insensitive to alterations in the sample selection procedure.

We also conduct several supplementary analyses to verify the consistency and interpretation of our estimates. To corroborate the postulate wherein a discontinuous narrative change takes place at the fifty percent Leave-share threshold, we run placebo-evaluations where the classification of constituencies as treated and control is done at other percentages. In this analysis, we are unable to find any systematic effects. Second, we test for the existence of pre-referendum parallel trends in voting patterns between Leave and Remain-labelled units, and find no evidence in this sense when zooming in on constituencies around the fifty-percent result, corroborating the common trends assumption made in the main evaluation.

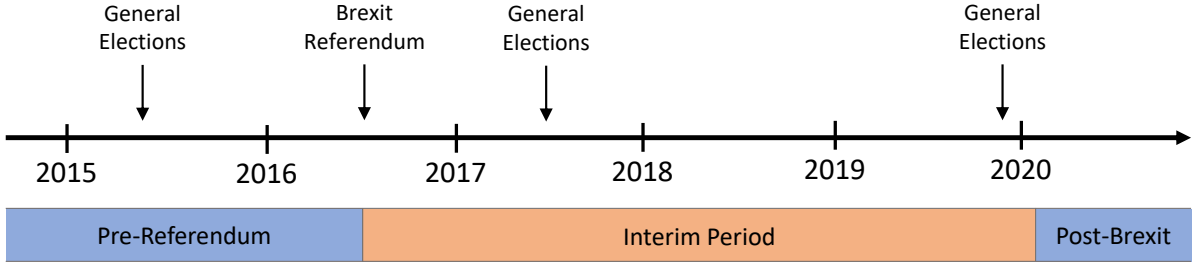
2 Political Context

In this section, we motivate our use of the 2016 referendum and its aftermath to assess whether geographical polarisation is causally exacerbated by voters learning about the policy preferences of their geographical reference group. We do not provide an exhaustive overview of Brexit, but rather explain why it is an appealing setting that allows us to study the local social learning mechanism.

The first appealing feature is Brexit’s delayed implementation. The Brexit referendum was held on 23 June 2016, with 17.4 million people (51.9%) voting in favour of leaving the EU, and 16.1 million (48.1%) voting to remain (Becker et al. 2017). Nonetheless, the the UK’s actual departure was disrupted with several delays. The official withdrawal process began on 29 March 2017, but two extensions to the allocated two-year negotiations period were granted. The first was to 31 October 2019 (from the original deadline of 29 March 2019) and then to 31 January 2020, with the latter being the end of the UK’s EU membership.

⁹We detail our empirical strategy and the assumptions required for consistency in Section 4.2.

Figure 1: Brexit Timeline



Note: Chronologically, we show the pre-referendum 7 May 2015 general elections, the 23 June 2016 Brexit referendum, the post-referendum 8 June 2017 and 12 December 2019 general elections, and the 31 January 2020 official Brexit date.

Therefore, a substantial time-lag exists between the initial referendum (which revealed the stance of voters on the issue) and the moment Brexit was carried out, a time-lag characterized by substantial uncertainty concerning the form in which (and whether at all) Brexit would occur (Graziano et al. 2021).

We exploit the delay and uncertainty in Brexit's implementation to assess how local contextual influences shape the voters' views on the policy. In contrast to other instances where voters may have only one opportunity to vote on a policy with its implementation following shortly afterwards, voters in the UK had further opportunities to vote based on their Brexit views *before* the UK left the EU, but *after learning about the referendum's outcome* - these opportunities were the general elections organized in the interim as shown in Figure 1.

In line with the local social learning mechanism, we can interpret the initial referendum result as an information signal on local preferences that voters learn about before the subsequent elections, and use the results of these elections to measure updated Brexit beliefs.

This leads to the second appealing feature: the policies of the Conservatives and the Liberal Democrats were strongly linked to the issue of leaving the EU. In 2017 (and again in 2019), the Conservative party campaigned largely in favour of Brexit¹⁰, whereas the Liberal Democrats adopted a stark anti-Brexit stance, campaigned for a second referendum in 2017¹¹. In fact, in the first chapter of their 2017 manifesto, entitled "Protect Britain's Place in Europe", the Liberal Democrats explicitly state that "*every vote for the Liberal Democrats in this election is a vote to give the final say to the British people*". We therefore proxy local anti-Brexit views following the referendum using the Liberal Democrat vote share in our empirical analysis.

¹⁰Conservative Party Manifesto in 2017 <https://tinyurl.com/y3k6xgry> (and again in 2019 <https://tinyurl.com/rc2ltrw>, which promised to 'Get Brexit Done' on the front cover.)

¹¹Liberal Democrat Manifesto in 2017 <https://tinyurl.com/y9cdeotq>

A final appealing feature is the communication of local referendum results in the media. The EU referendum result was unexpected and the margin of victory was small (Evans and Menon 2017). As a result, there was considerable interest in which areas of the country contributed to the overall result, and maps of UK voting areas "labelling" them as Remain or Leave appeared in many leading newspapers and online news media. Figure 2 shows such a map produced by the BBC shortly after the referendum.¹²

As further discussed below, the media portrayal of the results has important implications for measuring the local social learning effect empirically since the discontinuity in the labelling of constituencies as 'Leave' or 'Remain' allows for the comparison of two similar groups that may differ solely by their media portrayal¹³, and therefore constitutes an important means of identifying the causal effects of social learning.

We exploit these appealing features to test for local social learning leading to increased polarisation between constituencies. In the following subsection, we present a simple model capturing the mechanism.

3 Theory and Hypotheses

We outline the mechanism by which local information signals contribute to polarisation: essentially, we predict belief convergence concerning a policy (i.e., Brexit) through social learning *within* a constituency that leads to polarisation *between* constituencies. This occurs as voters update their beliefs about the policy's net utility after observing its popularity among those in their vicinity following an information signal (i.e., the referendum). We first present a baseline model that describes how local social learning works generally. We then extend this framework by introducing a signal discontinuity, thus formalizing the "media label" from Figure 2 in our model. It is precisely this discontinuous label that we exploit in our empirical analysis for causal inference.


3.1 Baseline Model

We assume that voters in a constituency choose whether to support policy X according to their belief of the policy's net utility. We also assume that voter i 's prior belief of the policy's net utility x_i is drawn from a uniform distribution between $a - 1$ and a , where $0 < a < 1$. Voters' beliefs about the policy's net utility are distributed evenly between

¹²Other examples: The Daily Telegraph <https://tinyurl.com/y66h6526>, and The Daily Express <https://tinyurl.com/yxcqbqpc>.

¹³Intuitively, we assume that, on average, constituencies where Leave won 51% of the vote are otherwise similar to constituencies where Leave won 49% of the vote. We discuss our empirical strategy in-depth below.

Figure 2: Media Labels: Geographical Portrayal of the Brexit Referendum's Results



Note: Modified from the BBC - <https://tinyurl.com/BBCbrexit2016>. In blue (yellow) areas, Leave (Remain) won a majority of votes in the 2016 referendum. In our empirical analysis, we study whether being labelled 'Leave' or 'Remain' causally affects post-referendum local-level voting patterns. In particular, we ask whether the Brexit-opposing Liberal Democrats lost ground in favour of the Conservatives in Leave-labelled areas relative to Remain-labelled areas following the referendum.

$a - 1$ (which is the belief of the voter who gets the lowest net utility from the policy) and a (which is the belief of voter who gets the largest net utility).

The first vote provides the information signal. A voter decision rule is simple: they vote in favour of the policy if their prior belief on the policy's net utility $x_i > 0$, in opposition if $x_i < 0$, and indifferent if $x_i = 0$. The first vote on policy X is held (i.e., the referendum), and a proportion $\pi = a$ vote in favour of the policy.

Voters then update their beliefs about the policy's net utility according to the information signal they receive from the first vote's result - a mechanism consistent with contextual effects theories that emphasize the role of local learning (Cutler 2007). Essentially, voters update their beliefs if they have imperfect information about the policy's net utility to the community. Observing the policy's popularity may allow a voter to learn more about the impact the policy has on them directly or indirectly (Burbank 1997). More informally, the voter realizes that there likely is a reason why the others voted the way they did.

In the case of leaving the EU, the popularity of Brexit in a constituency may be influenced by the impact of Brexit on the constituency's economy (for example through its impact on trade - Graziano et al. 2021). Therefore a voter may learn more about the potential economic impact of Brexit (as well as a myriad of other outcomes) from the views of those nearby, and update their belief about the net utility they receive from Brexit.¹⁴ We also assume that voters update their beliefs after observing the average belief in the population. Their updated beliefs are x'_i .

It is this social learning mechanism that is critical in explaining how the proportion of voters in the constituency in favour of the policy changes from the first to the second vote. A second vote is held with individuals voting according to their updated beliefs, resulting in a proportion π' voting in favour of the policy.¹⁵ We now explain how the proportion of voters in favour of the policy changes from the first to second vote, conditional on the initial distribution of beliefs.

Voters observe the average belief across the population after the first vote, and update their beliefs accordingly before the second vote. Using linear opinion pooling according to Stone (1961), one can approximate the Bayesian updating for a given informed voter as follows:

$$x'_i = \lambda s + (1 - \lambda)x_i$$

¹⁴For example, upon observing that those nearby largely support Brexit, the voter is *ceteris paribus* more likely to conclude that Brexit is a good policy.

¹⁵As further discussed below, we do not have a second referendum in our setting, but instead use the performance of Liberal Democrats in general elections to proxy for post-referendum Brexit views.

Where $s = a - \frac{1}{2}$ is the signal that voters observe (the average belief in the population), and λ is the weight that they place on the signal (where $0 < \lambda < 1$). The more importance they place on the signal (whether because they feel connected to the community or do not hold particularly strong prior beliefs), the more their beliefs converge to the signal. In our context, λ represents how strongly voters update their beliefs about Brexit in response to observing the referendum's local results.

The distribution of beliefs among voters after the first vote is therefore a uniform distribution between $(a - 1) + \frac{\lambda}{2}$ and $a - \frac{\lambda}{2}$. Relative to the initial distribution, the distribution has moved inwards towards the average belief $a - \frac{1}{2}$, and the size of this contraction is determined by how much informed voters update their beliefs - i.e., λ .

Figure 3 illustrates how the distribution of beliefs changes from the first vote (red) to the second vote (blue). In the first vote, Constituency 1 is largely against the policy while the opposite is true in Constituency 2. Then, after observing the signal, we see beliefs converging towards the average belief in both constituencies. Importantly, note there is divergence, or polarisation, between the two constituencies: in Constituency 1 the proportion of voters who are in favour of the policy falls, whereas this rises in Constituency 2 *because of the signal leading to belief updating*. This exacerbates the pre-existing differences between constituencies, which is precisely the type of polarisation between constituencies that we explore empirically.

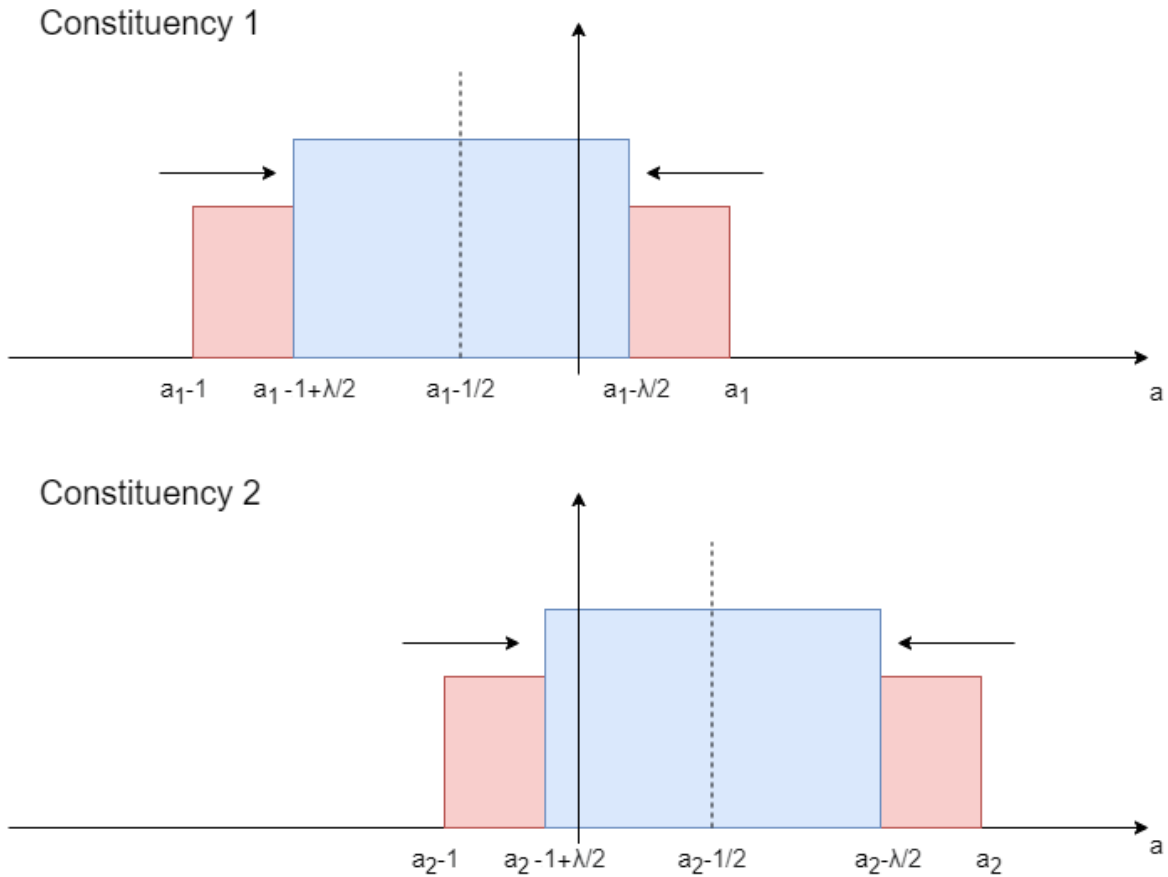
Formally, the proportion of the voters who vote in favour of the policy in the first vote is a . The proportion of voters who vote in favour of the policy in the second vote is the proportion of the new distribution whose posterior beliefs are above 0 (the blue distribution in Figure 3):

$$\pi' = \begin{cases} 0 & \text{if } a \leq \frac{\lambda}{2} \\ 1 & \text{if } a \geq 1 - \frac{\lambda}{2} \\ \frac{a - \frac{\lambda}{2}}{1 - \lambda} & \text{otherwise} \end{cases} \quad (1)$$

To conclude, voters update their beliefs and converge towards the average belief in their constituency, which causes polarisation *between* constituencies when measuring the proportion of voters in favour or against the policy (ultimately, what we care about is a binary variable: whether one is pro or anti-Brexit). Practically, having observed the results of the referendum within their constituency, voters update their beliefs about Brexit's net utility. As voters update their beliefs based on the average belief of Brexit within constituencies, this results in polarisation between constituencies.

Ceteris paribus, the proportion of voters in favour of Brexit in a constituency that voted largely for Brexit at the referendum increases *after the referendum*, while the opposite

Figure 3: The Local Social Learning Mechanism: Distribution of Original and Updated Beliefs



Note: We illustrate the impact of local social learning on voters' policy preferences. The initial distribution of beliefs for each constituency is depicted in red: uniformly between $a_1 - 1$ and a_1 for Constituency 1, and $a_2 - 1$ and a_2 for Constituency 2. After the voters observe the local information signal, their preferences converge towards the mean belief in their constituency. The resulting new distribution is depicted in blue, and is narrower than the original one. Now, beliefs are distributed between $a_1 - 1 + \lambda/2$ and $a_1 - \lambda/2$, and $a_2 - 1 + \lambda/2$ and $a_2 - \lambda/2$ in Constituencies 1 and 2, respectively, where $0 < \lambda < 1$ measures how informative the signal is perceived to be. Intuitively, the proportion of voters in favour of the policy falls (rises) in Constituency 1 (2). This illustrates how increased polarisation between two geographical units can arise as a result of social learning and belief convergence.

holds for a constituency that largely opposed Brexit initially. Voters learning more about local preferences exacerbates geographical polarisation in the proportion of voters in support of Brexit between constituencies.

3.2 Signal Discontinuity: The Label Effect

Voters may observe the popularity of a policy such as leaving the EU in different ways. One group of voters may take great interest in learning about the popularity of Brexit within their neighbourhood and take detailed notice of discussions in local news media. A second group of voters may take less interest, and only pay attention to the headlines in their local news media. As discussed in the previous section, one signal they may observe is the labelling of their geographic reference group as a "Leave" or "Remain" as in Figure 2.

To reflect this, we extend the baseline model by assuming that there are two types of voters: *Sophisticated* voters who update their beliefs after observing the average belief (as in the baseline model) and uninformed voters. *Unsophisticated* voters receive a binary signal: one value when policy X is supported by the majority of voters (the "Leave"-label), another value when the policy opposed by a majority (the "Remain"-label).

The population shares of the two groups are n_I , and n_U , respectively ($n_I + n_U = 1$), and the two groups have the same prior distributions of beliefs. As before, there are two votes on the policy. The proportion of voters in each group in favour of the policy is π_I and π_U for sophisticated and unsophisticated groups in the first vote, respectively (π'_I and π'_U for the second vote), so the proportion of all voters in favour of the policy at the first vote is $\pi = n_I\pi_I + n_U\pi_U$.

Unsophisticated voters receive a binary signal after the first vote: whether the first vote saw the majority of the population vote in favour of the policy, or not. That is, they observe the label.

Therefore, their updating can be summarised as follows (where the majority of voters vote in favour of the policy in the first vote if $a > \frac{1}{2}$):

$$x'_{iU} = \lambda s_U + (1 - \lambda)x_{iU}$$

Where the signal s_U is defined as:

$$s_U = \begin{cases} \underline{a} = -\frac{1-\lambda}{2} & \text{if } a < \frac{1}{2} \\ \frac{1}{2} & \text{if } a = \frac{1}{2} \\ \bar{a} = \frac{1-\lambda}{2} & \text{if } a > \frac{1}{2} \end{cases} \quad (2)$$

Therefore the uniform distribution of beliefs before the first vote (between $a - 1$ and a) is transformed to a uniform distribution between $\lambda s_U + (1 - \lambda)(a - 1)$ and $\lambda s_U + (1 - \lambda)a$. The proportion of unsophisticated voters in favour of the policy is, as a function of a :

$$\pi'_U = \begin{cases} 0 & \text{if } a \leq \frac{\lambda}{2} \\ a - \frac{\lambda}{2} & \text{if } \frac{\lambda}{2} < a < \frac{1}{2} \\ \frac{1}{2} & \text{if } a = \frac{1}{2} \\ a + \frac{\lambda}{2} & \text{if } \frac{1}{2} < a < 1 - \frac{\lambda}{2} \\ 1 & \text{if } a \geq 1 - \frac{\lambda}{2} \end{cases} \quad (3)$$

Sophisticated voters act in the same way as voters in the baseline model - for example the proportion of informed voters who are in favour of the policy in the second vote (π'_I) is described in (1). Combining the two groups, the total proportion of voters in favour of the policy in the second vote is:

$$\pi' = n_I \pi'_I + n_U \pi'_U = \begin{cases} 0 & \text{if } a \leq \frac{\lambda}{2} \\ n_I \left(\frac{a - \frac{\lambda}{2}}{1 - \lambda} \right) + n_U \left(a - \frac{\lambda}{2} \right) & \text{if } \frac{\lambda}{2} < a < \frac{1}{2} \\ \frac{1}{2} & \text{if } a = \frac{1}{2} \\ n_I \left(\frac{a - \frac{\lambda}{2}}{1 - \lambda} \right) + n_U \left(a + \frac{\lambda}{2} \right) & \text{if } \frac{1}{2} < a < 1 - \frac{\lambda}{2} \\ 1 & \text{if } a \geq 1 - \frac{\lambda}{2} \end{cases} \quad (4)$$


This is illustrated in Figure 4. There are two notable features of the updating. The first is evidence of polarisation *between* constituencies for different values of a . For a given value of a , π' lies further away from half than π , which reflects the conclusions of the baseline model.

The second is a discontinuity at $a = \frac{1}{2}$. This arises from the binary signal that unsophisticated voters receive, which depends on whether the policy won a majority in their constituency or not. The size of the discontinuity is $n_u \lambda$, which depends on the size of the unsophisticated group, and the strength of belief updating.

3.3 Empirical Implications and Testable Hypotheses

As we discuss in more detail in the following sections, this discontinuity - a result of voters observing a Brexit "media label" rather than a more nuanced signal - provides an opportunity to empirically identify the local social learning mechanism that leads to polarisation. While a correlation between π' and π may suffer from selection bias, the discontinuity allows for comparison between two groups that are similar in their prior

Figure 4: Proportion of Voters in Favour of the Policy Before and After the First Vote



Note: We illustrate the effect of local social learning on the proportion of voters supporting a given policy within a constituency. Here, the π schedule depicts the proportion of voters in favour of the policy, and is a function of a . The π' schedule illustrates the proportion of voters in support of the policy following the information signal received in a first vote on the matter. We note increased polarisation, as - for any value of a - the distance between π' and one half is greater compared to that corresponding to π .

preference for the policy (similar values of a), but differ solely in whether the policy enjoyed majority support in the constituency in the first vote, the referendum.

The baseline model and the extension presented above provides testable predictions that can be empirical analysed using the context of the EU referendum and following general elections. We would expect to observe polarisation *between* constituencies because of convergence (from social learning) *within* constituencies. Given the sharp contrast in policy between the Liberal Democrats and the Conservatives over Brexit, we would expect constituencies where leaving the EU received a greater share of the vote in the 2016 referendum to be associated with decreases in the share of the vote for the Liberal Democrats (and a corresponding increase in the vote share for the Conservatives) in the elections following the referendum compared to the elections that took place before the referendum.

Another prediction would be a discontinuity at the 50% threshold - where the "Leave" or "Remain" label changes - which allows for clean identification of polarisation *caused* by local convergence. This leads to two testable hypotheses:

Hypothesis 1: The vote share obtained by the anti-Brexit Liberal Democrats decreases following the referendum in Leave-labelled constituencies, relative to Remain-labelled constituencies.

Hypothesis 2: The vote share obtained by the Conservatives increases following the referendum in Leave-labelled constituencies, relative to Remain-labelled constituencies.

We now turn to the empirics.

4 Empirics

4.1 Data

Our unit of analysis is the constituency. Using information from the House of Commons Library, we construct a panel-dataset tracking the vote shares obtained by the LDs and the Conservatives over time. Since the referendum took place in 2016, we use information from the pre and post-referendum 2015 and 2017 general elections in our main analysis.¹⁶

In line with Hypothesis 1, because of their stark anti-Brexit stance (Hobolt 2018), we chiefly operationalize local anti-Brexit sentiments (our dependent variable) by the LD vote share. We assess Hypothesis 2 by evaluating the performance of the Conservatives.

Of course, we note that using vote shares, whose determinants are multifaceted (Mellon et al. 2018), to operationalize support for one policy is not ideal. Nevertheless, given that Brexit was at the forefront of political campaigns following the referendum (Prescott-

¹⁶We analyse the 2010 and 2019 elections in Section 5.3.

Smith 2019), we argue that vote shares represent a contextually appropriate proxy. In Section 5.4, we provide supporting evidence.

Retrieving constituency-level referendum results is not as straightforward since official tallies were not communicated, for the most part, at this level of aggregation, but rather at the larger local authority [LA] level (Becker et al. 2017). To overcome this difficulty, we simply assign to every constituency the 2016 referendum result communicated for its corresponding LA.

We prefer this matching approach to the alternative of using constituency-level vote share *estimates* constructed, for example, in Hanretty (2017) for two reasons. First, since the theorized mechanism linking local information and polarisation consists of social learning, it is not the actual results that will influence local preferences, but rather the numbers communicated to and in turn internalized by the local electorate. Because official figures were largely available for LAs only, it is these results that were widely circulated in the media (Figure 2). Then, if the label matters, we expect the electorate to be responsive precisely to these official figures - i.e., the relevant signal.

Second, as we explain below, we compare vote shares between constituencies where the referendum Leave figure was just above and just below fifty percent in order to identify a causal effect. If we relied on local estimates rather than official figures, the classification of constituencies as Leave-labelled or Remain-labelled would become noisy, because different sets of estimates would place some constituencies on different sides of the threshold purely due to random variation (Hanretty 2017). Hence, we believe the official LA results best suit our purposes, both theoretically and methodologically.

That said, one final difficulty remains, as roughly 32 percent of constituencies cannot be mapped to one particular LA. This happens because several wards within these constituencies are assigned to different LAs.¹⁷ Here, constructing suitable measures for the referendum's official outcome is not feasible at the constituency-level - different parts of the electorate likely received a different signal. Thus, we drop constituencies spanning multiple LAs.

Following these restrictions, 434 constituencies (out of 650) remain in our sample.

4.2 Empirical Strategy

We seek to assess if local social learning can exacerbate geographical polarisation by investigating whether the local results of the Brexit referendum affected subsequent voting patterns at the. However, a simple regression linking the Leave-share to the score obtained

¹⁷For example, the Arundel and South Downs constituency is comprised of sixteen wards - seven of which are assigned to Horsham as their LA, four to Arun, three to Chichester, and two to Horsham and Mid Sussex each.

by a particular party would likely produce biased estimates of the true causal effect of local information (Gallego et al. 2016). For example, a negative association between the Leave-share and the performance of the LDs may arise not due to the hypothesized learning channel, but rather because the LDs' anti-Brexit campaign prior to the referendum was more effective in LD strongholds. In cases such as this, cause and effect may be reversed, and any statistical associations would be uninformative.

We overcome this difficulty by first constructing a binary treatment variable $LEAVE LABEL_i$ equal to one (zero) in constituencies where the Leave-share recorded in 2016 was above (below) fifty percent. As discussed above, we assume that at this fifty-percent threshold a discontinuous change took place in the manner in which the referendum's results were framed by the media and, in turn, disseminated to local voters.¹⁸

When crossing fifty percent Leave-support, we posit that the narrative switches from belonging to a local reference group whose majority opposes Brexit to one where the policy has majority appeal. This 'label shift' suits our purposes, as it allows us to precisely isolate the impact of an information signal concerning local preferences on future policy support from the influence of other possibly relevant confounders, constituting a novel testing method for the local social learning mechanism.

We exploit the longitudinal structure of our data to assess how party vote shares evolved between 2015 and 2017 in Leave-labelled (treated) constituencies, *relative* to those labelled Remain (control) by running the following regression¹⁹:

$$Y_{it} = \alpha_i + \omega_{rt} + \beta LEAVE LABEL_i * POST_t + \epsilon_{it} \quad (5)$$

In equation 5, Y_{it} gives the constituency-level vote share obtained by a party competing in the general elections organized in year t , and $POST_t$ is a dummy variable equal to one in 2017 (zero in 2015). Our specification includes two sets of fixed effects. First, we add a vector of constituency dummies α_i , whose purpose is to correct for the influence of any time-invariant confounders that may be correlated with both a constituency's leave result as well as the electoral performance of the different parties. Similar to Gallego et al. (2016), including these terms ensures that the impact of the Leave majority label is quantified by only exploiting within-constituency variation in vote shares. Second, we include a set of region-by-electoral year dummies ω_{rt} in order to correct for any non-

¹⁸Corroborating this, we show that no significant effects can be retrieved by fitting our empirical model for "placebo" thresholds where the narrative switch from minority to majority Leave support did not take place (Section 5.2).

¹⁹Of course, implementing a regression discontinuity design (RD) would generally be a compelling alternative research method in this context. Unfortunately, the limited sample size prevents the use of an RD - see Freier et al. (2015) who, tackling a different question, rely on a difference-in-differences strategy in a setting where an RD design is inhibited by the small sample size. As long as the parallel trends assumption holds (discussed below), the method employed here yields consistent estimates.

linear time trends specific to the UK’s eleven regions that may otherwise bias our results. Including these regional terms is standard practice in studies working with longitudinal UK electoral data, since local and regional organizations are not fully independent (see [Fetzer 2019](#), p. 3859 for the model our own specification closely follows). Our results are however robust to instead using region-specific linear time-trends. Finally, ϵ_{it} is an error term. Throughout, we apply constituency-level clustering to account for the fact that the same geographical units are observed repeatedly ([Bertrand et al. 2004](#)).

β is the coefficient of interest. A significant non-zero β estimate implies that Leave-labelled constituencies experienced notable changes in party vote shares over time, relative to those labelled Remain. In line with our hypotheses, we expect this coefficient estimate to be negative (positive) when looking at LD (CON) vote shares

Nonetheless, because of the observational nature of our data, whether β ’s coefficient estimate retrieved via our specification may be seen as a *causal* treatment effect crucially hinges on the following "common trends" ([Angrist and Pischke 2008](#)) assumption holding: *Had the Brexit referendum not taken place, party vote shares would have evolved similarly in Leave-labelled and Remain-labelled constituencies*. Under this postulate, any divergent changes in party performance between treated and control units will arise *because of* the referendum’s recorded outcome.

We argue that this postulate is unlikely to hold when considering the full sample of constituencies. Indeed, it would be unwise to assume that no preference-impacting time-varying shock heterogeneously affected Leave-labelled constituencies relative to their Remain-labelled counterparts following the 2015 general election - other than the referendum itself - given large social, economic and political differences existing between these constituencies ([Becker et al. 2017](#)).

To overcome this difficulty, we proceed by fitting specification 5 subsamples of constituencies clustered progressively tighter around the fifty percent Leave threshold. That is, by imposing a margin of $h \in (0, 50)$ percentage-points, β is estimated by contrasting the evolution of party vote shares in the restricted treatment group - consisting of constituencies where the reported Leave-share lies between 50 and $50 + h$ - with that retrieved in the restricted control group - where $50 - h$ to 50 percent of voters expressed Leave support. We argue that the causal interpretation of the estimated label effect becomes more plausible as the margin narrows $h \rightarrow 0$, since the likelihood of heterogeneous time-varying shocks confounding the treatment effects is diminished. We bring evidence corroborating this argument in Section 5.2.

Before proceeding, we note that restricting the sample to increase the credibility of our estimates necessarily comes at a cost in terms of precision: because we include fewer constituencies, the power of our specification falls. Although this downside cannot be

directly addressed, since restricting the sample is at the core of our identification strategy, we present results from different margins to highlight the robustness of our findings by fitting our specification at different points on the "precision-credibility" spectrum.

5 Results

In Section 5.1, we show that Brexit support rose *after* the 2016 referendum in Leave-labelled relative to Remain-labelled constituencies, consistent with social learning exacerbating geographical polarisation. In Sections 5.2 and 5.3, we assess whether defining our treatment based on the fifty percent Leave-threshold is sensible, and check the validity of the common-trends assumption required for causal inference. We find corroborating evidence. In Section 5.4, we show that the label effects are larger in constituencies where the LDs performed strongly historically, bringing credence to our operationalization of Brexit views by party vote shares.

5.1 Main Results: The Polarising Effects of Local Brexit Labels

Results are reported in Table 1, where we show the coefficient estimates derived by fitting specification 5 on the vote shares obtained by the LDs and the Conservatives in rows one and two, respectively. Column (1) presents the β coefficient estimates from analysing the full final sample of constituencies. This figure tells us how vote shares evolved between the 2015 and 2017 elections in constituencies where a majority of the electorate supported Leave, relative to constituencies where Remain received majority support. In columns (2) through (5), we restrict our evaluation to observations in tighter bandwidths around the fifty-percent Leave threshold. Figure 5 graphically illustrates our results.

Overall, our findings constitute evidence for Hypothesis 1, whereby Brexit support became increasingly polarized due to the referendum's local label. Looking at the first row, the estimates suggest that the LDs' performance worsened in Leave-labelled constituencies relative to Remain-labelled ones after the referendum. In our preferred specification (column 4²⁰), the point estimate indicates a statistically-significant negative treatment effect of roughly two percentage points, a non-negligible change considering the formation's overall score of 7.4 percent obtained in 2017. Reassuringly the estimate is qualitatively insensitive to altering the margin.

²⁰We chose a 2.5 percentage-points margin for our preferred specification following the sensitivity checks performed in Sections 5.2 and 5.3, whose results indicate that wider margins may lead to spurious coefficient estimates due to the common trends assumption becoming unreasonable. Conditional on design validity, we prefer working with a wider margin to reduce low power issues arising from small restricted samples. The 2.5 percentage points choice satisfies this requirement. Nevertheless, our results are qualitatively insensitive to the margin choice.

Table 1: Main Results: The Electoral Effects of Local Referendum Labels

	Outcome: Party Vote Shares (% Turnout)				
	(1)	(2)	(3)	(4)	(5)
	Full Sample	7.5pp Margin	5pp Margin	2.5pp Margin	1.5pp Margin
LD (Anti-Brexit)	-0.878 (0.536) [0.103] 5.79***	-1.52** (0.684) [0.028] 2.59***	-1.48* (0.802) [0.067] 2.34***	-2.05*** (0.669) [0.003] 1.85*	-2.05*** (0.759) [0.009] 3.41**
CON	(0.492) [0.000]	(0.648) [0.000]	(0.769) [0.003]	(1.053) [0.082]	(1.69) [0.048]
Observations	868	416	302	172	116
Constituencies	434	208	151	86	58


Note: The treatment variable is the interaction of two dummy terms - LEAVE LABEL, equal to one in constituencies where the Leave vote share recorded in the 2016 Brexit Referendum is above fifty percent (zero otherwise), and POST, equal to one for observations recorded in 2017 (zero in 2015). All specifications include constituency and region-by-electoral year fixed effects. Standard errors, given in (round brackets), are clustered at the constituency-level; p-values are given in [square brackets]; ***p<0.01, **p<0.05, *p<0.1

Moreover, we find the opposite for the Conservatives, the party perceived as (largely - [Hobolt and Rodon 2017](#)) committed to and able to carry out Brexit. This party’s relative performance improved following the referendum in Leave-labelled constituencies, with our preferred specification retrieving a statistically significant positive coefficient estimate of just under two percentage-points, mirroring the loss suffered by the LDs. These findings corroborate Hypothesis 2.

Taken together, the label effects retrieved here suggest that learning about the referendum’s local results had a notable impact on the electorate’s subsequent voting behaviour, in a pattern consistent with voters becoming increasingly polarized on the matter as a result of contextual social learning ([Bala and Goyal 1998](#); [Bliuc et al. 2007](#)). To this point, the data suggests that the greatest relative electoral losses (gains) were experience by the party largely opposing (supporting) Brexit, precisely in areas where the Leave sentiment turned out to constitute a majority.

This novel finding helps us better understand the importance of one’s local political surroundings in causally shaping policy preferences ([McGarty et al. 2009](#)). In line with the scholarship on contextual effects ([Burbank 1997](#); [Cutler 2007](#)), our results bring new causal evidence for the hypothesis wherein a local learning mechanism may exacerbate

Figure 5: The Electoral Effects of Local Referendum Labels



Note: The treatment variable is the interaction of two dummy terms: LEAVE LABEL, equal to one in constituencies where the Leave vote share recorded in the 2016 Brexit Referendum is above fifty percent (zero otherwise), and POST, equal to one for observations recorded in 2017 (zero in 2015).

spatial-clustering, thus adding to the broader literature investigating the causes of increasing geographical polarisation (Kaplan et al. 2020).²¹

Of course, this interpretation hinges on the validity of our empirical method, which we assess in the following subsections.

5.2 Is Polarisation Truly Driven by the Label? Investigating Placebo Thresholds

We consider our treatment construction whereby constituencies were assigned the Leave-labelled treated or Remained-labelled control status based on their referendum Leave-share being above or below fifty percent, respectively. We argued that this definition is appropriate under the assumption that a discontinuous shift in the way the referendum's

²¹For completeness, in Appendix Figure A1, we investigate the label's effects on the performance of the Labour Party using analogous empirical methods. There is insufficient evidence to reject the zero effect null hypothesis. Hence, we cannot conclude that the performance of the Labour party was meaningfully affected by the local label. Speculatively, this may be the result of the Labour Party's perceived lack of clarity on the Brexit matter (see e.g., <https://bbc.in/2OoXAOD>, and Hobolt and Rodon 2017), which would make the vote share obtained by this party a poor operationalization choice for Brexit sentiments.

results are disseminated to and internalized by the electorate takes place here (Section 2).

We corroborate this assumption by testing the following: *if* a discontinuous change does exist at fifty percent, then similar "label effects" to those shown above should not be systematically retrieved at other thresholds, where the "minority to majority Leave-support" label switch does not occur.

We run placebo-specifications wherein we first assign a constituency to a new "treatment" ("control") group if the 2016 reported Leave-share falls above (below) an arbitrary threshold k unequal to fifty percent:

$$\text{PLACEBO LEAVE LABEL}_i^k = 1 \text{ if LEAVE SHARE}_i \geq k \text{ (0 otherwise)}$$

Where $k \in (0, 100)$. Then, we run the following for different k values:

$$Y_{it} = \alpha_i + \omega_{rt} + \beta_k \text{PLACEBO LEAVE LABEL}_i^k * \text{POST}_t + \epsilon_{it} \quad (6)$$

Equation 6 mirrors specification 5, adjusted for the placebo treatment definition. For a given placebo threshold value k , β_k 's coefficient estimate measures the 2015-2017 party vote share changes in placebo-treated ("placebo Leave-labelled") relative to placebo-control ("placebo Remain-labelled") constituencies.

Intuitively, if the Brexit label constitutes an information shock at fifty percent, we do not expect β_k 's coefficient estimate to be systematically statistically different from zero for $k \neq 50$ when zooming in on tighter margins.

As before, the number of observations included in the sample depends on the margin. For consistency with the main analysis, we use the 2.5 percentage-points margin in our preferred specification. Nevertheless, we again allow this margin to vary for robustness.

Tables 2 and 3 present the estimated coefficients retrieved for placebo thresholds chosen both above and below fifty percent, respectively.²²

We note that we are unable to reject the zero effect null hypothesis for any of the placebo thresholds considered in our preferred specification (column 4). Also, unlike the estimates presented in Table 1, the coefficients here are highly sensitive to the margin- choice both in magnitude and statistical-significance, suggesting that no robust associations can be retrieved at the different placebo thresholds.

These results corroborate our key postulate whereby a signal discontinuity occurs at the fifty percent Leave-share threshold, as it is solely here where we retrieve robust estimates

²²Since our empirical strategy involves restricting our sample tightly around the threshold, we do not consider k -values much larger or smaller in magnitude than fifty-percent. We do so because the restricted sample becomes too small for reliable statistical inference - e.g., very few constituencies exist in a 2.5 percentage points margin around 40 percent Leave-share.

Table 2: Assessing Label Effects at Placebo Thresholds Above Fifty Percent

Outcomes: Party Vote Shares (% Turnout)					
	(1)	(2)	(3)	(4)	(5)
	Full Sample	7.5pp Margin	5pp Margin	2.5pp Margin	1.5pp Margin
Panel A - Placebo Leave Threshold: 52.5%					
LD (Anti-Brexit)	-0.340 (0.377) [0.367] 5.36***	-0.284 (0.470) [0.546] 2.24***	-0.436 (0.506) [0.389] 1.42***	-0.063 (0.697) [0.928] -0.272	-0.381 (1.03) [0.713] 0.656
CON	(0.443) [0.000]	(0.476) [0.000]	(0.504) [0.005]	(0.717) [0.705]	(0.960) [0.498]
Observations	868	446	350	202	96
Constituencies	434	223	175	101	48
Panel B - Placebo Leave Threshold: 53.75%					
LD (Anti-Brexit)	-0.199 (0.352) [0.572] 5.11***	-0.214 (0.430) [0.620] 2.46***	-0.002 (0.473) [0.997] 1.48***	0.197 (0.566) [0.729] 0.398	0.294 (0.902) [0.746] -0.405
CON	(0.441) [0.000]	(0.477) [0.000]	(0.473) [0.002]	(0.734) [0.590]	(1.05) [0.701]
Observations	868	464	372	162	112
Constituencies	434	232	186	81	56
Panel C - Placebo Leave Threshold: 55%					
LD (Anti-Brexit)	0.015 (0.349) [0.965] 5.51***	0.163 (0.356) [0.647] 3.00***	0.482 (0.437) [0.272] 1.86***	0.433 (0.712) [0.544] 0.317	0.458 (0.911) [0.617] -0.290
CON	(0.434) [0.000]	(0.459) [0.000]	(0.509) [0.000]	(0.696) [0.650]	(0.917) [0.753]
Observations	868	480	346	178	120
Constituencies	434	240	173	89	60

Note: Standard errors, given in (round brackets), are clustered at the constituency-level; p-values are given in [square brackets]; ***p<0.01, **p<0.05, *p<0.1

Table 3: Assessing Label Effects at Placebo Thresholds Below Fifty Percent

Outcomes: Party Vote Shares (% Turnout)					
	(1)	(2)	(3)	(4)	(5)
	Full Sample	7.5pp Margin	5pp Margin	2.5pp Margin	1.5pp Margin
Panel A - Placebo Leave Threshold: 47.5%					
LD (Anti-Brexit)	-0.520 (0.665) [0.435]	-1.42* (0.810) [0.082]	-0.472 (1.20) [0.693]	0.917 (1.22) [0.457]	0.269 (2.46) [0.914]
CON	5.88*** (0.567) [0.000]	3.16*** (0.661) [0.000]	1.82 (1.18) [0.125]	0.851 (1.68) [0.614]	-1.66 (2.46) [0.506]
Observations	868	388	238	100	50
Constituencies	434	194	119	50	25
Panel B - Placebo Leave Threshold: 46.25%					
LD (Anti-Brexit)	-0.295 (0.624) [0.637]	-0.405 (0.728) [0.579]	0.290 (0.505) [0.567]	2.48 (1.81) [0.179]	2.89 (1.99) [0.161]
CON	5.86*** (0.536) [0.000]	3.42*** (0.568) [0.000]	2.16*** (0.659) [0.001]	0.100 (1.48) [0.947]	0.637 (1.61) [0.696]
Observations	868	352	244	86	46
Constituencies	434	176	122	43	23
Panel C - Placebo Leave Threshold: 45%					
LD (Anti-Brexit)	-0.523 (0.639) [0.414]	-0.443 (0.672) [0.511]	0.073 (0.645) [0.910]	0.384 (1.28) [0.766]	-0.462 (1.15) [0.692]
CON	5.29*** (0.565) [0.000]	2.33*** (0.548) [0.000]	1.07 (0.682) [0.121]	-1.23 (1.33) [0.362]	-1.51 (1.52) [0.330]
Observations	868	342	186	66	52
Constituencies	434	171	93	33	26

Note: Standard errors, given in (round brackets), are clustered at the constituency-level; p-values are given in [square brackets]; ***p<0.01, **p<0.05, *p<0.1

consistent with the predictions of the local social learning mechanism. At other thresholds - where the "minority to majority" label switch does occur - no systematic significant patterns are supported retrieved.

Hence, the evidence suggests that the label mattered.

5.3 Dynamic Analysis: Validating the Empirical Design

We assess whether the estimates in Table 1 may be *causally* interpreted. We do so by appraising the "common trends" assumption. We use our extended dataset (2010-2019 elections), and compare pre-referendum trends in party vote shares between Leave-labelled and Remain-labelled constituencies (see [Huet-Vaughn 2019](#) for a similar analysis investigating US presidential results).²³ If the local referendum label does trigger a change in Brexit preferences, then we should not observe systematic trend divergences between treated and control constituencies *prior* to 2016 when considering tight margins around the fifty percent Leave-share threshold.

We run the following:

$$Y_{it} = \alpha_i + \omega_{ct} + \sum_{\substack{j=2010 \\ j \neq 2015}}^{2019} (\beta_j \text{TREAT}_{it}^j) + \epsilon_{it} \quad (7)$$


Equation 7 closely resembles equation 5. The treatment variable here, TREAT_{it}^j , is the product of the LEAVE LABEL_i and a binary variable equal to one in election year j (zero otherwise). Note that an indicator corresponding to 2015 (the final pre-referendum general election year) is excluded in order to avoid the dummy-variable trap. Therefore, all the estimates should be interpreted relative to 2015 results.²⁴

The coefficients β_j capture any relative changes in vote shares between treated and control constituencies in year j . The coefficient estimate for the lead regressor β_{2010} allows us to test whether trend divergences occurred between the 2010 and 2015 pre-referendum elections. If Remain-labelled constituencies represent an appropriate counterfactual for Leave-labelled ones, we do not expect this estimate to be statistically distinguishable from zero when using narrow bandwidths. Finally, we also estimate the lag coefficient β_{2019} . While not crucially pertinent, doing so allows us to assess whether the label effects were persistent or transitory, by checking whether they are still notable three years after the referendum.

²³While data for pre-2010 elections are available, they do not correspond satisfactorily to current parliamentary constituencies because of the Fifth Periodic Review of Westminster Constituencies, which significantly altered electoral boundaries. Since our focus is geographical polarisation, we only consider the 2010 election in our pre-trends analysis, for which the correspondence is one-to-one.

²⁴Removing final pre-intervention periods is standard practice - e.g., [Fetzer \(2019\)](#).

Figure 6: Dynamic Analysis: The Electoral Effects of Local Referendum Labels



Note: The vertical line marks the referendum month, June 2016. The dependent variables are the vote shares obtained by the Liberal Democrats and the Conservative Party in orange and blue, respectively. The margin used is 2.5 percentage-points. Statistically significant coefficient estimates are marked as follows: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

In Figure 6, we plot the estimated β_j coefficients against their corresponding year. The dependent variables are the LD and CON vote shares in orange and blue, respectively.²⁵ As before, the estimates are retrieved via our preferred specification which restricts the sample of constituencies to those lying within a 2.5 percentage points margin around the fifty-percent Leave-share threshold. We illustrate the robustness of our findings to different margins in Table 4.

We argue that the results corroborate the contextual validity of the common trends assumption. In our preferred specification, we are unable to reject the zero effect null hypothesis when looking at β_{2010} 's estimate for any of the outcomes. Moreover, this coefficient's estimates - unlike those corresponding to the post-referendum coefficient β_{2017} - are once more highly-sensitive in both sign and magnitude to our margin choice,

²⁵In Figure A2, we conduct an analogous analysis using the Labour Party's vote share as our dependent variable.

Table 4: Dynamic Analysis: The Electoral Effects of Local Referendum Labels

		Outcomes: Party Vote Shares (% Turnout)				
		(1)	(2)	(3)	(4)	(5)
		Full Sample	7.5pp Margin	5pp Margin	2.5pp Margin	1.5pp Margin
LD (Anti-Brexit)	$TREAT_{2010}$	-1.74***	-0.513	-0.613	-0.163	0.716
		(0.613)	(0.748)	(0.867)	(1.28)	(2.07)
		[0.005]	[0.493]	[0.481]	[0.898]	[0.730]
	$TREAT_{2017}$	-0.878	-1.52**	-1.48*	-2.05***	-2.05***
		(0.538)	(0.690)	(0.811)	(0.682)	(0.782)
		[0.104]	[0.029]	[0.070]	[0.004]	[0.010]
	$TREAT_{2019}$	-2.44**	-1.98	-1.78	-2.30	-3.21
		(1.02)	(1.27)	(1.45)	(1.52)	(1.96)
		[0.017]	[0.120]	[0.219]	[0.134]	[0.106]
CON	$TREAT_{2010}$	0.435	0.060	0.166	-0.044	2.21
		(0.530)	(0.530)	(0.869)	(1.17)	(1.53)
		[0.412]	[0.935]	[0.849]	[0.970]	[0.156]
	$TREAT_{2017}$	5.79***	2.60***	2.34***	1.85*	3.41*
		(0.493)	(0.653)	(0.777)	(1.07)	(1.74)
		[0.000]	[0.000]	[0.003]	[0.088]	[0.055]
	$TREAT_{2019}$	9.07***	3.31***	3.34***	2.59	3.26
		(0.856)	(1.59)	(1.19)	(1.79)	(2.82)
		[0.000]	[0.000]	[0.006]	[0.152]	[0.252]
Observations		1,736	832	604	344	232
Constituencies		434	208	151	86	58

Note: Standard errors, given in (round brackets), are clustered at the constituency-level; p-values are given in [square brackets]; ***p<0.01, **p<0.05, *p<0.1

suggesting that no statistically robust trend divergences exist between Leave and Remain-labelled constituencies prior to 2016, as expected.²⁶

Summarising, the results from Sections 5.2 and 5.3 suggest that *i*) no systematic *post-referendum* label effects exist at placebo thresholds where the "minority to majority" switch did not take place, and that *ii*) party vote shares evolved in a parallel manner in Leave-labelled and Remain-labelled constituencies between the two elections taking place *pre-referendum*. These findings bring credence to the causal reading of our main results wherein local Brexit preferences changed as a result of the discontinuous label shift taking place at fifty percent.

²⁶Finally, we prefer to not draw any conclusions concerning persistence. As seen here, β_{2019} coefficient estimates retain the sign of their β_{2017} counterparts and, for the most part, are larger in magnitude. Nevertheless, standard errors are also notably larger. While speculatively this happens because more time has elapsed since the information signal, we cannot confidently make this assertion given our evidence.

5.4 Is Using Party Vote Shares Sensible? Heterogeneity by Historical LD Strength

Thus far, we found evidence for the LDs losing ground in favour of the Conservatives in Leave-labelled relative to Remain-labelled constituencies following the referendum, corroborating a mechanism wherein geographical polarisation is exacerbated by local learning. As discussed, however, the evidence is indirect since we use party votes shares to proxy Brexit views. Our justification for doing so has been that one’s perceived benefits of voting for an anti-Brexit party such as the LDs fall when learning that a majority of constituents in one’s geographical reference-group supports the policy.

We bring further evidence supporting this interpretation by testing whether our estimated coefficients are larger in magnitude in constituencies where the LDs performed strongly historically, relative to areas where they underperformed. Our reasoning is as follows: We argue that if indeed the label effect operates by influencing *Brexit preferences*, then this should only lead to individuals changing *which party they vote for if* they believe that their new choice can notably impact how Brexit is carried out.

What does this mean for the LDs? Intuitively, voters will update their choice only if they think that their new decision may influence the LD candidate’s chances of entering parliament and (potentially) influence the Brexit process. Since, in most constituencies, the LD party did not have a realistic chance of winning (Hobolt and Rodon 2017), we argue that the treatment effects should be stronger in constituencies where the LDs were perceived as competitive, and could therefore plausibly enter parliament.

We use 2010 constituency-level LD vote shares to proxy for historical LD performance. Thus, our assumption here is that, on average, the party’s vote share in 2010 constitutes a reasonable proxy for the LDs’ *expected* performance in 2017 - a reasonable assumption since the 2010 election was the last one to occur prior to the investigated period.

Empirically, we assess whether the referendum’s local label effects are stronger in constituencies where the LDs’ 2010 vote share was relatively large by fitting our preferred specification²⁷ 5 separately on the subsamples of constituencies belonging to the bottom and top quartiles of the 2010 LD vote share variable, respectively.²⁸

²⁷For consistency, we continue to use a 2.5 percentage-points margin around the fifty-percent leave-share threshold in our preferred specification. Nevertheless, we check robustness to employing different margins, and our findings remain qualitatively insensitive. See Table A1.

²⁸For completeness, we also run our regression on the subsample of constituencies belonging to the middle two quartiles. Note that, since we are interested in comparing coefficient estimates retrieved at the two extremes of our moderating variable’s distribution, we prefer our approach to the alternative whereby an interaction term between our treatment variable and the LD’s 2010 vote share is introduced into the model (see Hainmueller et al. 2018 for further details on why a multiplicative interaction model may be inappropriate in settings such as ours, where one may not reasonably assume the existence of linear interaction effects).

Table 5: Heterogeneity of the Label Effect by Historical LD Performance

	Outcome: LD (Anti-Brexit) Vote Share (% of Turnout)				
	(1)	(2)	(3)	(4)	(5)
	Full	Weak	Medium	Strong	Difference
	Sample	Historical LD	Historical LD	Historical LD	p-value
LEAVE LABEL	-2.05*** (0.669) [0.003]	-0.157 (0.248) [0.534]	-0.833 (0.821) [0.316]	-6.23*** (1.35) [0.000]	0.000
Observations	172	44	84	44	-
Constituencies	86	22	42	22	-
R-squared	0.411	0.669	0.594	0.677	-

Note: The p-value reported in column (5) is associated with the one sided null hypothesis that treatment effect estimates are larger in magnitude in column (4) relative to column (2). Standard errors, given in (round brackets), are clustered at the constituency-level; p-values are given in [square brackets]; ***p<0.01, **p<0.05, *p<0.1


The findings are given in Table 5. In column (1), we replicate the analysis from Table 1, column (4). Then, we present results derived from fitting our preferred specification on the subsample of constituencies in the bottom quartile, the two middle quartiles, and the upper quartile of the 2010 LD vote share in columns (2), (3) and (4), respectively. We graphically illustrate these estimates in Figure 7.

Overall, we find evidence suggesting that the (relative) losses suffered by the LDs in Leave-labelled constituencies were only statistically-significant in areas where they were perceived as competitive (we estimate a fall of 6.2 percentage points in column 4, compared to a negligible decrease of 0.1 in column 2). While these estimates should be cautiously interpreted given the small sample size, they bring credence to our operationalization choice wherein party vote shares were used to proxy for post-referendum local Brexit views.

6 Related Literature

Following [Martin and Webster \(2020\)](#), we posit that disentangling the determinants of geographical polarisation is not solely an academic curiosity. First, such an analysis helps us better understand changes in the intensity of electoral competition. To this point, [Bishop \(2008\)](#) argues that increased spatial-clustering has led decreased electoral pressures, negatively impacting the ability of voters to hold representatives accountable (see also [Sussell 2013](#), [Hopkins 2017](#)). Second, polarisation has been documented to affect both the allocation of local public-goods ([Gerber and Gibson 2009](#)), and the way in which

Figure 7: Heterogeneity of the Label Effect by Historical LD Performance



Note: We illustrate the estimates from Table 5, with rows one through four corresponding to columns (1) through (4).

votes translate to legislative seats ([Chen and Rodden 2013](#)). Finally, one might question what factors influence locally polarized support for key policies such as Brexit ([Becker et al. 2017](#)).

Our contribution is twofold. First, our simple model linking geographical political polarisation to local learning sheds light on an understudied channel driving spatial-clustering. Building on insights from an emerging literature in social psychology ([Bliuc et al. 2007](#); [McGarty et al. 2009](#)), we provide a framework wherein information shocks contribute to the exacerbation of group-based identities, and political preferences. This effort complements existing work emphasizing the role of economic factors ([Carreras 2019](#)), political trust ([Mitsch et al. 2021](#)), and that of the rural-urban divide ([Scala and Johnson 2017](#); [Rohla et al. 2018](#)) in explaining polarisation.

Second, we test our arguments by exploring geographical polarisation in the UK concerning Brexit following the 2016 referendum, a setting which allows us to estimate the causal effects of local information shocks in the form of media labels. Related to our analysis, [Hobolt et al. \(2020\)](#) find evidence that strong "affective" polarisation is associated with the UK's withdrawal, with Brexit identities becoming increasingly prevalent after the referendum, ultimately resulting in prejudice between those in the "Leave" and "Remain" camps. We rely precisely on the Leave-Remain division (applied to constituencies rather than individuals) introduced here to assess whether the local information contained in such labels may influence subsequent voting patterns.

Given the local learning mechanism investigated above, our article also adds to a literature on contextual effects, according to which one's socio-economic and political environment causally influences one's policy preferences. As explained by [Gallego et al.](#)

(2016), existing work shows that contextual effects operate through numerous mechanisms including interpersonal contact (Butler and Stokes 1974; Huckfeldt et al. 2005), party mobilization (Denver and Hands 1997), neighbourhood cues (Cho and Rudolph 2008), and strategic sorting.

More pertinent for our investigation, several studies highlight the importance of contextual learning in shaping political behaviours (Conover 1985; Burbank 1997; Mutz 2002; Cantoni and Pons 2021). Books and Prysby (1988) argue that it is precisely contextual information that shapes neighbourhood effects. The authors also explicitly emphasize that causal inference in contextual analyses is difficult, and suggest using longitudinal data for more accurate inference. We proceed as such by investigating the impact of referendum’s local results using a panel design.

Focusing on local effects explicitly, Cutler (2007) asserts that citizens weigh local interests when forming political views. Using Canadian data, he finds evidence that indicators of local interests are associated with individual opinions and group feelings, suggesting that local information may influence preferences. We expand upon the insights in Cutler (2007) and contend that one’s exposure to information signals on local group interests is precisely a mechanism contributing to increased geographical clustering. We, however, move beyond simple statistical associations. By employing a difference-in-differences design, we make progress towards better understanding the precise *causal* role local information plays in shaping policy preferences.

7 Conclusion

We seek to explain the recently-documented rise in geographical political polarisation (Jennings and Stoker 2019; Kaplan et al. 2020). We propose that one catalyst is a local social learning mechanism, triggered by voters becoming better informed about the views of those nearby. This results in preferences becoming more homogeneous within local electoral clusters, but increasingly heterogeneous between clusters. Exploiting the Brexit referendum and its aftermath, we find evidence of voters learning from contextual information signals in the form of media labels, causally contributing to polarisation between constituencies.

Future work may build upon our findings. In particular, it would be insightful to further investigate the transmission channel underlying local learning. First, one might ask what factors influence how fast voters learn. Do, say, the strength of local media outlets or the pervasiveness of social media matter? Second, a related question is whether the parties themselves also learned from and then strategically reacted to the local label. For instance, were more campaigning resources targeted towards constituencies where Leave

won? Third, we note again that the vote shares obtained by the Liberal Democrats and the Conservatives served as proxies for local Brexit sentiments, as we are interested in geographical polarisation concerning *policy* choices. Therefore, one cannot conclude solely based on our analysis that the effects we document spill-over to persistent changes in the electorate's taste for different *parties* - an interesting question that may help us better understand how the appeal of two-party versus multi-party systems changes as local reference groups become increasingly relevant (Prosser 2018).

Moreover, while our focus is geographical, our theoretical framework has broader appeal. As information signals on reference groups²⁹ becomes more widespread (for instance, due to increased accessibility to social media - Zhuravskaya et al. 2020) so does social learning. Thus, future work may check whether the effects documented here can be generalized.³⁰

Finally, our findings suggest that the manner in which reference groups are "labelled" can play a role in exacerbating group-based identities. We posit that further studying this labelling mechanism, perhaps experimentally, is worthwhile.

²⁹Be them organized geographically or otherwise.

³⁰For instance, a cross-disciplinary literature studies whether group-based identities explain Covid-19 vaccination uptake (e.g. Wakefield and Khauser 2021; Chu et al. 2021; Sinclair and Agerström 2021).

References

- Abramowitz, A. I. and K. L. Saunders (2008). Is Polarization a Myth? *Journal of Politics* 70(2), 542–555.
- Aidt, T., F. Grey, and A. Savu (2021). The Meaningful Votes: Voting on Brexit in the British House of Commons. *Public Choice* 186, 587–617.
- Allison, P. D. (1994). Using Panel Data to Estimate the Effects of Events. *Sociological Methods & Research* 23, 174–199.
- Angrist, J. and J. Pischke (2008). *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press.
- Autor, D. (2003). Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of Labor Economics* 1(21), 1–42.
- Axe-Browne, A. and M. E. Hansen (2020). Still dividing the electorate? Brexit and voter evaluation of candidates. *Journal of Elections, Public Opinion and Parties*, 1–23 <doi: 10.1080/17457289.2019.1711097>.
- Bala, V. and S. Goyal (1998). Learning from Neighbours. *Review of Economic Studies* 65, 595–621.
- Becker, S. O., T. Fetzner, and D. Novy (2017). Who voted for Brexit? A comprehensive district-level analysis. *Economic Policy* 32(92), 601–650.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How Much Should We Trust Differences-in-Differences Estimates? *Quarterly Journal of Economics* 119(1), 249–285.
- Bishop, B. (2008). *The Big Sort: Why the Clustering of Like-Minded America is Tearing Us Apart*. New York: Houghton Mifflin Harcourt.
- Bliuc, A., C. McGarty, K. Reynolds, and D. Muntele (2007). Opinion-based group membership as a predictor of commitment to political action. *European Journal of Social Psychology* 37, 19–32.
- Books, J. and C. Prysby (1988). Studying Contextual Effects on Political Behavior. *American Politics Quarterly* 16(2), 211–238.
- Branton, R. P. and B. S. Jones (2005). Reexamining Racial Attitudes: The Conditional Relationship between Diversity and Socioeconomic Environment. *American Journal of Political Science* 49(2), 359–372.

- Brown, J. and R. Enos (2018). *Partisan Segregation*. Harvard University <<https://scholar.harvard.edu/files/renos/files/brownenos.pdf>>.
- Brown, J. R. and R. D. Enos (2021). The measurement of partisan sorting for 180 million voters. *Nature Human Behaviour* 5, 998–1008.
- Burbank, M. J. (1997). Explaining contextual effects on vote choice. *Political Behavior* 19(2), 113–132.
- Butler, D. and D. E. Stokes (1974). *Political Change in Britain: The Evolution of Electoral Choice*. London: Macmillan.
- Cantoni, E. and V. Pons (2021). Does Context Outweigh Individual Characteristics in Driving Voting Behavior? Evidence from Relocations within the U.S. *American Economic Review* (forthcoming).
- Carreras, M. (2019). 'What do we have to lose?': Local economic decline, prospect theory, and support for Brexit. *Electoral Studies* 62, 102094.
- Chen, J. and J. Rodden (2013). Unintentional gerrymandering: Political geography and electoral bias in legislatures. *Quarterly Journal of Political Science* 8(3), 239–269.
- Cho, W. K. T. and T. J. Rudolph (2008). Emanating Political Participation: Untangling the Spatial Structure Behind Participation. *British Journal of Political Science* 38(2), 273–289.
- Chu, J., S. Pink, and R. Willer (2021). Religious identity cues increase vaccination intentions and trust in medical experts among American Christians. *PNAS* doi.org/10.1073/pnas.2106481118.
- Conover, P. J. (1985). The impact of group economic interests on political evaluations. *American Politics Quarterly* 13, 139–166.
- Converse, P. (1964). "The Nature of Belief Systems in Mass Publics" In D.E. Apter (Ed.) *Ideology and Discontent*. New York: Free Press.
- Cox, K. R. (1969). The voting decision in a spatial context. *Progress in Geography* 1, 81–117.
- Cutler, F. (2007). Context and attitude formation: Social interaction, default information, or local interests? *Political Geography* 26(5), 575–600.
- Denver, D. T. and G. Hands (1997). *Modern Constituency Electioneering: Local Campaigning in the 1992 General Election*. Abingdon, Oxon: Routledge.

- Evans, G. and A. Menon (2017). *Brexit and British Politics*. Cambridge: Polity Press.
- Fetzer, T. (2019). Did Austerity Cause Brexit? *American Economic Review* 11(109), 3849–3886.
- Ford, R. and M. Goodwin (2017). Britain After Brexit: A Nation Divided. *Journal of Democracy* 28(1), 17–30.
- Freier, R., M. Schumann, and T. Siedler (2015). The earnings returns to graduating with honors - Evidence from law graduates. *Labour Economics* 34, 39–50.
- Gallego, A., F. Buscha, P. Sturgis, and D. Oberski (2016). Places and Preferences: A Longitudinal Analysis of Self-Selection and Contextual Effects. *British Journal of Political Science* 46(3), 529–550.
- Gerber, E. R. and C. C. Gibson (2009). Balancing Regionalism and Localism: How Institutions and Incentives Shape American Transportation Policy. *American Journal of Political Science* 53(3), 633–648.
- Gimpel, J. G. and I. S. Hui (2015). Seeking politically compatible neighbors? The role of neighborhood partisan composition in residential sorting. *Political Geography* 48, 130–142.
- Graziano, A. G., K. Handley, and Limao (2021). Brexit Uncertainty and Trade Disintegration. *Economic Journal* 131(635), 1150–1195.
- Hainmueller, J., J. Mummolo, and Y. Xu (2018). How Much Should We Trust Estimates from Multiplicative Interaction Models? Simple Tools to Improve Empirical Practice. *Political Analysis* 27(2), 163–192.
- Hanretty, C. (2017). Areal interpolation and the UK’s referendum on EU membership. . *Journal of Elections, Public Opinion and Parties* 27(4), 466–483.
- Hobolt, S. B. (2018). Brexit and the 2017 UK General Election. *Journal of Common Market Studies* 56(S1), 39–50.
- Hobolt, S. B., T. J. Leeper, and J. Tilley (2020). Divided by the vote: affective polarization in the wake of the Brexit referendum. *British Journal of Political Science*, 1–37 <doi: 10.1017/S0007123420000125>.
- Hobolt, S. B. and T. Rodon (2017). Cross-cutting issues and electoral choice. EU issue voting in the aftermath of the Brexit referendum. *Journal of European Public Policy* 27(2), 227–245.


- Hopkins, D. (2017). *Red Fighting Blue: How Geography and Electoral Rules Polarize American Politics*. New York: Cambridge University Press.
- Huckfeldt, R., I. Ken'ichi, and F. U. Pappi (2005). Patterns of Disagreement in Democratic Politics: Comparing Germany, Japan, and the United States. *American Journal of Political Science* 49(3), 497–514.
- Huckfeldt, R. and J. Sprague (1987). Networks in Context: The Social Flow of Political Information. *American Political Science Review* 81(4), 1197–1216.
- Huet-Vaughn, E. (2019). Stimulating the vote: ARRA road spending and vote share. *American Economic Journal: Economic Policy* 11(1), 292–316.
- Jennings, W. and G. Stoker (2019). The Divergent Dynamics of Cities and Towns: Geographical Polarisation and Brexit. *The Political Quarterly* 90, 155–166.
- Johnston, R., D. Manley, K. Jones, and R. Rohla (2020). The geographical polarization of the American electorate: a country of increasing electoral landslides? *GeoJournal* 85, 187–204.
- Kaplan, E., J. L. Spenkuch, and R. Sullivan (2020). Partisan Spatial Sorting in the United States: A Theoretical and Empirical Overview. Technical report, <https://www.semanticscholar.org/paper/Partisan-Spatial-Sorting-in-the-United-States%3A-A-Kaplan-Spenkuch/0ca6460bbe72960cd838670d1bf36cd273cca68a>.
- Levendusky, M. S. (2013). Why Do Partisan Media Polarize Viewers? *American Journal of Political Science* 57(3), 611–623.
- Martin, G. J. and S. W. Webster (2020). Does residential sorting explain geographic polarization? *Political Science Research and Methods* 8(2), 215–231.
- McGarty, C., A. Bliuc, E. F. Thomas, and R. Bongiorno (2009). Collective Action as the Material Expression of Opinion-Based Group Membership. *Journal of Social Issues* 65(4), 839–857.
- Mellon, J., G. Evans, E. Fieldhouse, J. Green, and C. Prosser (2018). Brexit or Corbyn? Campaign and Inter-Election Vote Switching in the 2017 UK General Election. *Parliamentary Affairs* 71(4), 719–737.
- Mitsch, F., N. Lee, and E. R. Marrow (2021). Faith no more? The divergence of political trust between urban and rural Europe. *Political Geography*, doi.org/10.1111/1467-923X.12612.

- Moore, M. and G. Ramsay (2017). UK media coverage of the 2016 EU Referendum campaign.
- Mutz, D. C. (2002). The Consequences of Cross-Cutting Networks for Political Participation. *American Journal of Political Science* 46(4), 838–855.
- Pickard, H., V. Bove, and G. Efthymoulou (2022). You (Br)exit, I stay: The effect of the Brexit vote on internal migration. *Political Geography* doi.org/10.1016/j.polgeo.2021.102576.
- Prescott-Smith, S. (2019). Which issues will decide the general election? <<https://yougov.co.uk/topics/politics/articles-reports/2019/11/07/which-issues-will-decide-general-election>>.
- Prosser, C. (2018). The strange death of multi-party Britain: the UK General Election of 2017. *West European Politics* 41(5), 1226–1236.
- Rohla, R., R. Johnston, K. Jones, and D. Manley (2018). Spatial scale and the geographical polarization of the American electorate. *Political Geography* 65, 117–122.
- Scala, D. J. and K. M. Johnson (2017). Political polarization along the rural-urban continuum? The geography of the presidential vote, 2000-2016. *The ANNALS of the American Academy of Political and Social Science* 672(1), 162–184.
- Sinclair, S. and J. Agerström (2021). Do Social Norms Influence Young People’s Willingness to Take the COVID-19 Vaccine? *Health Communication* doi.org/10.1080/10410236.2021.1937832.
- Stone, M. (1961). The Opinion Pool. *The Annals of Mathematical Statistics* 32(4), 1339–1342.
- Sussell, J. (2013). New Support for the Big Sort Hypothesis: An Assessment of Partisan Geographic Sorting in California, 1992-2010. *PS: Political Science and Politics* 46(4), 768–773.
- Wakefield, J. and A. Khauser (2021). Doing it for us: Community identification predicts willingness to receive a COVID-19 vaccination via perceived sense of duty to the community. *Journal of Community & Applied Social Psychology* 31(5), 603–614.
- Weatherford, M. S. (1983). Evaluating economic policy: a contextual model of the opinion formation process. *Journal of Politics* 45(4), 158–174.
- Zhuravskaya, E., M. Petrova, and R. Enikolopov (2020). Political Effects of the Internet and Social Media. *Annual Review of Economics* 12, 415–438.

Appendix

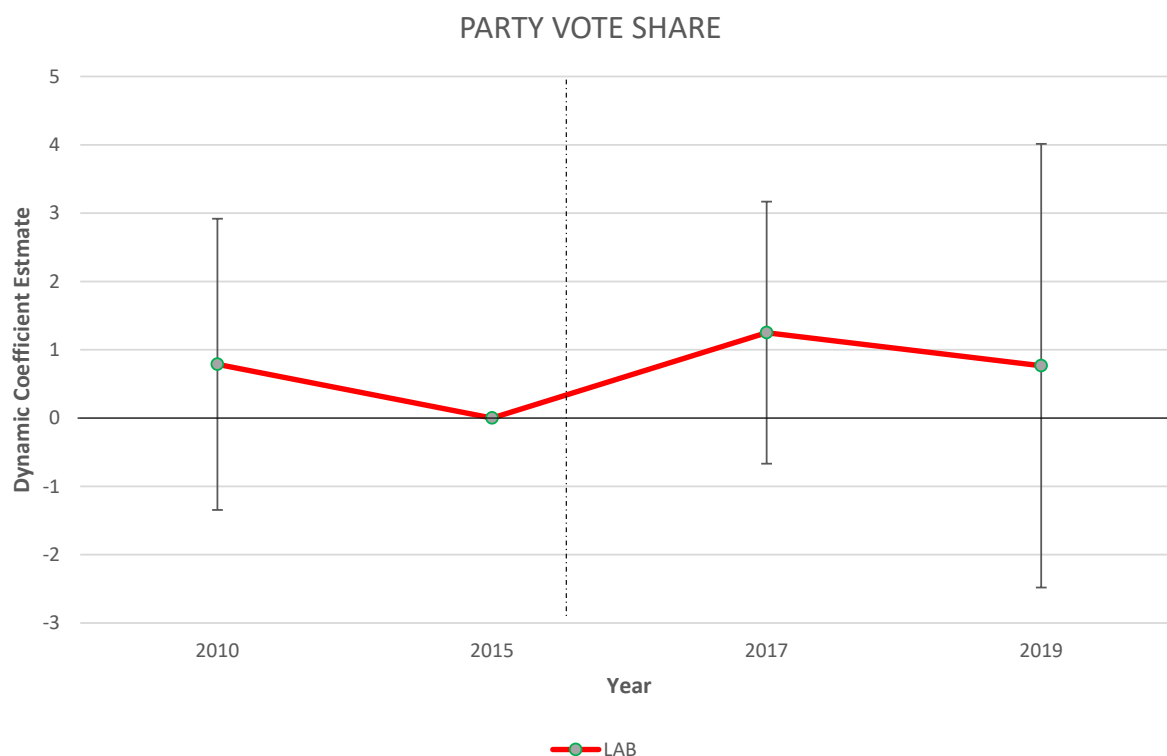
Appendix A - Supplementary Analyses

Figure A1: The Effect of Local Referendum Labels on the Labour Party's Vote Share



Note: We replicate the baseline analysis performed in Figure 3 using the vote share obtained by the Labour Party as our dependent variable. We illustrate the coefficient estimates and corresponding 90 percent confidence intervals retrieved by fitting our difference-in-differences specification 5 on the full final sample of constituencies, and on subsamples defined by progressively tighter margins around the fifty percent Leave share threshold. Our treatment variable captures the interaction of two dummy terms - LEAVE LABEL, equal to one in constituencies where the Leave vote share recorded in the 2016 Brexit Referendum is above fifty percent (zero otherwise), and POST, equal to one for observations recorded in 2017 (zero in 2015).

Figure A2: Dynamic Analysis: The Effect of Local Referendum Labels on the Labour Party's Vote Share



Note: We replicate the dynamic analysis from Figure 6 using the vote share obtained by the Labour Party as our dependent variable. To do so, we plot the coefficient estimates, alongside their 95 percent confidence intervals, retrieved from fitting specification 7 on the full final sample, against the corresponding year. The dotted black line marks June 2016, when the referendum took place. The margin considered is 2.5 percentage points around the fifty percent leave share threshold.

Table A1: Heterogeneity of the Label Effect by Historical LD Performance: Robustness to the Bandwidth Choice

Dependent Variable: Vote Share obtained by the LDs (% of Turnout)					
	(1)	(2)	(3)	(4)	(5)
	Full	7.5pp	5pp	2.5pp	1.5pp
	Sample	Margin	Margin	Margin	Margin
All Historical LD	-0.878 (0.536) [0.103]	-1.52** (0.684) [0.028]	-1.48* (0.802) [0.067]	-2.05*** (0.669) [0.003]	-2.05*** (0.759) [0.009]
Weak Historical LD	-0.135 (0.251) [0.591]	-0.392* (0.223) [0.085]	-0.196 (0.225) [0.389]	-0.157 (0.821) [0.534]	-0.425*** (0.131) [0.006]
Medium Historical LD	-0.764 (0.534) [0.154]	-1.08 (0.751) [0.154]	-1.05 (1.10) [0.343]	-0.833 (0.821) [0.316]	0.373 (0.536) [0.492]
Strong Historical LD	-2.91* (1.66) [0.082]	-6.23*** (1.97) [0.003]	-5.75** (2.34) [0.019]	-6.23*** (1.35) [0.000]	-5.04*** (1.36) [0.002]
Difference p-value	0.058	0.002	0.009	0.000	0.000

Note: We assess the robustness of the estimates presented in Table 5 to altering the margin employed when restricting our sample of constituencies to those tightly clustered around the fifty percent Leave vote share recorded in the 2016 referendum. To do so, we report the difference-in-differences coefficient estimates obtained by fitting specification 5 on the full final sample of constituencies in column (1), and on subsamples defined by progressively tighter margins around the fifty percent Leave share threshold in columns (2) to (5). In row one, we do not exclude any constituencies based on the historical performance of the LDs. In rows two, three and four we only include constituencies belonging to the bottom, middle and top quartiles of the 2010 LD vote share variable, respectively. The dependent variables is the vote share obtained by the Liberal Democrats. Our treatment variable of interest captures the interaction of two dummy terms - LEAVE LABEL, equal to one in constituencies where the Leave vote share recorded in the 2016 Brexit Referendum is above fifty percent (zero otherwise), and POST, equal to one for observations recorded in 2017 (zero in 2015). The p-values reported in the final row are associated with the one sided null hypothesis that treatment effect estimates are larger in magnitude in row four relative to row two. Standard errors, given in (round brackets), are clustered at the constituency-level; p-values are given in [square brackets]; Significance levels: ***p<0.01, **p<0.05, *p<0.1