Tension Points: A Theory & Evidence on Migration in Brexit

Levi John Wolf
levi.john.wolf bristol.ac.uk

March 3, 2018

1 The Question

The GIS Research UK Conference Data Challenge asked: *was Brexit primarily driven by the rate of change in migration, rather than the total headcount?* To interrogate this, I used local regression methods, hierarchical models, migration data provided by the Office of National Statistics, and a novel method to extract population volatility from fine-grained Consumer Data Research Centre Data. Depending on the *implicit* hypothesis used to operationalize the contest question, I find Brexit voting at local authority level was driven in part by the rate of change in their population structure, but some types of change drove Leave voting and some drove Remain.[1]

2 An Answer

For this analysis, I examine Brexit results at the local authority level. I do this because the subset of ward-level data available for analysis is highly demographically conspicuous, whereas local authorities have all relevant data. Analyzing a demographically unrepresentative subset of wards would mislead rather than inform.[2]

Therefore, official Brexit referendum results at local authority level collated by the UK Electoral commission are used. Then, demographic data from the Annual Population Survey (APS) is combined with data from the Local Area Migration Indicators (LAMI) data put together by the Office of National Statistics (ONS). The relevant APS data involved covariates on education, employment, age, and national origin, described in Appendix 6. To examine the question, net flows were computed from the LAMI following the method in Appendix 4. The net flow measures new migrants as a fraction of the population in that local authority. Its rate-of-change would provide the acceleration (or deceleration) in the change. Thus, using the net directly provides an indication of change, since it provides a measure of *new people* arriving or leaving each local authority. Finally, a novel method outlined in Appendix 5 to incorporate the CDRC's Local Super Output Area data into the local authority-level analysis is applied. Flatly, the volatility measure characterizes how different each subsequent year's ethnic distribution is from the previous year. Computed at LSOA, it is then summarized at local authority.

A region-varying-intercepts specification of this model was fit using `Stan`, a Bayesian modeling framework.[3] This means that each region has a unique baseline level of propensity to vote Leave, as shown in the right side of Figure 5 in Appendix 5. This amounts to allowing for a model somewhere between a regional fixed effect model, where each region has its own baseline that provides no information about other areas, and a global intercept model. This flexibly accounts for any baseline heterogeneity in leave support over regions. [4]

The substantive effect estimates for this model are shown in Table 1. Notably, two out of three of the migration/change effects are associated with *reduced Leave sentiment*. This suggests local authorities with more citizens from outside the UK (regardless of race) *as well as* the local authorities that received more Britons than they lost were more likely to vote for Brexit than the areas being left by UK residents or areas with low non-British populations. This reinforces a "tension point" interpretation: the popular relocation destinations among Brits that also have many non-British residents are places that voted Leave. However, places with large net external migration flows, as well as places with larger volatility in their ethnic distributions, are actually more likely to vote Remain, after controlling for regional factors and other confounders.

	Median	2.5%	97.5%	Δ IQR
% No Qualifications	-0.0451	-0.2449	0.1601	-0.1802
% University Degree	-0.7830	-0.8575	-0.7069	-11.0009
Change % Manufacture	-0.0749	-0.2200	0.0715	-0.2621
White Unemployment %	0.1734	-0.0915	0.4204	0.5030
% Aged 16-20	-0.4187	-0.7674	-0.0694	-0.7639
% Aged 20-24	-0.3762	-0.6111	-0.1211	-1.1081
% Aged 50 and Up	0.0074	-0.0990	0.1061	0.0804
Votes Cast	-0.0162	-0.0269	-0.0053	-0.7894
% Ethnic non-UK Born	0.1443	0.0113	0.2732	0.8223
% White non-UK Born	0.1666	0.0196	0.3144	0.8247
Mean Population Volatility ($\overline{d_w}$)	-0.9391	-1.7660	-0.0877	-0.8995
Mean Net External Migration	-1.7658	-3.2541	-0.2560	-0.5335
Mean Net Internal Migration	2.0045	0.4037	3.5937	1.2078

Table 1: Substantive effect estimates for the varying-intercept model. To give a sense of the relative importance of these various factors on the propensity to vote for Brexit, the last column shows the change in predicted Leave % when moving from the 25th percentile to the 75th percentile of the data on that attribute alone, holding the remainder at their means. When a row is lined in grey, its 95% HPDI overlaps with zero.

3 Why not other answers?

Many other plausible models were fit. One clearly relevant model also considers the change in flows from year-to-year. However, these predictors were either not significant (in the case of the net internal migration rate-of-change) or had no distinct substantive impact. Mean change, as a direct second-order measure, works too well to capture this effect for these to be distinct. Thus, these derived second-order migration effects are dominated by their first-order behavior and a better direct second-order measure.

Another model that deserves discussion is the new multiple-bandwidth GWR model (MGWR) (Fotheringham, Yang, and Kang 2017). An analogue to the model in Table 1 was fit using this approach, allowing the migration-relevant parameters to be locally-varying. There is little evidence that the patterns of the net migration vary significantly, although each covariate has a distinct spatial scale and pattern. Three surfaces for these local effects (and the local baseline Brexit sentiment) are reported in Appendix 7 for illustration.[5]

Other models, such as region-varying slope models and spatially-lagged X models, were fit. Varying slope models muddied identification of which effects mattered where: letting all migration-relevant effects vary over regions, there were no nonzero migration effects outside London. Since this region-scale heterogeneity is strongly at odds with the structure of heterogeneity demonstrated in the GWR approaches, it seems likely that assuming slopes vary by region simply does not match the underlying spatial structure of nonstationarity, if any.[6] The lagged-X models provide no substantively-different effects when filtering by statistical significance as long as the mean change term is retained.

In conclusion, migration, population volatility, and voting Leave have geographies that are not simply quantified by regional distinctions. Net migration flows *do* drive Leave sentiment in local authorities, although whether the migrants are coming from within (vs. outside) the UK matters. Ethnic volatility measures accurately characterize rapid change over the entire ethnic distribution; volatility also consistently drives Remain voting. In total, this supports the idea that Brexit-aligned areas are “tension points:” places (possibly with many non-UK born residents) to which Britons are moving. It’s not the size of change, it’s the type of change.

References

Chodrow, Philip S. 2017. “Structure and information in spatial segregation”. *Proceedings of the National Academy of Sciences*: 11591–11596.

Fotheringham, A. Stewart, Wenbai Yang, and Wei Kang. 2017. “Multiscale Geographically Weighted Regression (MGWR)”. *Annals of the American Association of Geographers*, no. 6: 1247–1265. doi:10.1080/24694452.2017.1352480. eprint: https://doi.org/10.1080/24694452.2017.1352480.

Giles, Micheal W, and Kaenan Hertz. 1994. “Racial threat and partisan identification”. *American Political Science Review* 88 (2): 317–326.

Harris, Richard, and Martin Charlton. 2016. “Voting out of the European Union: Exploring the geography of Leave”. *Environment and Planning A*: 2116–2128.

Kullback, Solomon, and Richard A Leibler. 1951. “On information and sufficiency”. *The Annals of Mathematical Statistics* 22 (1): 79–86.

Linzer, Drew A. 2012. “The Relationship between Seats and Votes in Multiparty Systems”. *Political Analysis* 20 (3).

Olkin, I., and F. Pukelsheim. 1982. “The distance between two random vectors with given dispersion matrices”. *Linear Algebra and its Applications*: 257–263. doi:10.1016/0024-3795(82)90112-4.

Rizzo, Maria L, and Gábor J Székely. 2016. “Energy distance”. *Wiley Interdisciplinary Reviews: Computational Statistics* 8 (1): 27–38.

Rocha, R. R., and R. Espino. 2009. “Racial Threat, Residential Segregation, and the Policy Attitudes of Anglos”. 00058, *Political Research Quarterly* 62, no. 2 (): 415–426. ISSN: 1065-9129, visited on 05/09/2014. doi:10.1177/1065912908320931. http://prq.sagepub.com/cgi/doi/10.1177/1065912908320931.

Wolf, Levi John, Taylor M. Oshan, and A. Stewart Fotheringham. 2017. “Single and Multiscale Models of Process Spatial Heterogeneity”. *Geographical Analysis*. ISSN: 1538-4632. doi:10.1111/gean.12147. http://dx.doi.org/10.1111/gean.12147.

4 Appendix: Net Flows from LAMI

The LAMI provides a set of inflows and outflows at each local authority, broken down by whether the flows are coming/going from within the UK or outside the UK. With this, it is simple to define a net flow as a percentage of population:

$$net = \frac{inflow - outflow}{total\ population} * 100 \tag{1}$$

Thus, net is the difference in migrants as a percent of the population in a local authority; if everyone up-and-left, it would be -100, and if the population doubled due to migration, it would be 100. In practice, the average $net_{internal}$ from 2011 to 2016 is between -2 and 2, meaning that some places in the UK (e.g. Newham in London) saw around 2% of their population emigrate to other areas of the UK and others (e.g. East Devon in the Southwest) see an additional 2% on top of their population being added from other places in the UK. The five-year average $net_{external}$ is much higher & usually nonnegative, since many places do not send more people abroad than come to them from abroad.

Using this net flow for both internal and external migration, we can define the rate-of-change in flow by taking the average difference of the year-on-year net percent flow. This average difference (and the direct net percent flow) are used to assess the competition question. Further, the average gradient of this 5-year time series of net migration flows can also be used to identify an alternative to the distribution volatility measures presented in Appendix 5. This gradient is a measure of "rate of change in migration flow," in contrast to the more general distributional volatility presented there.

5 Appendix: Magnitude of Population Mix Change

To incorporate the CDRC dataset on ethnic mix, I used novel distribution dynamics methods (Chodrow 2017). The competition dataset contains a distribution vector $q_i^t = [p_1, p_2, \ldots, p_K]$ for area i, measuring the fraction of population in K ethnic categories in year t.[7]

Using measures of distributional difference, such as the *Wasserstein distance* (Olkin and Pukelsheim 1982), *energy distance* (Rizzo and Székely 2016), or *Kullback-Leibler divergence* (Kullback and Leibler 1951), the total change in a distribution vector between years can be computed.[8] Here, the Wasserstein distance ($d_w(.)$) between population fraction vector in year t, q_t, and next year, q_{t+1}, is:

$$d_w(q_i^t, q_i^{t+1}) = \sum_j^K |q_{ij}^t - q_{ij}^{t+1}| \tag{2}$$

Thus, the Wasserstein distance is the absolute difference in ethnic composition vectors, summed. Like many measures of distance, these tend to experience a power-law decay. A map of the average distance between subsequent years' population distributions in each LSOA is shown on the left of Figure 5

This reduces the $N \times K \times T$ information about population composition in each LSOA over time to a single $N \times 1$ statistic expressing the typical change in population composition from year to year. Large values indicate the distribution has changed radically, and small values indicate the distributions have not changed significantly. To use these directly in an analysis at the local authority level, the mean of distances within each local authority divided by their standard deviation is used.[9]

6 Appendix: Control Variates

% of working-age population (WAP) with no higher education
% of WAP who have completed university

Education often factors into political leanings. Here, I anticipated local authorities with more educated populations to be more likely to vote Remain.

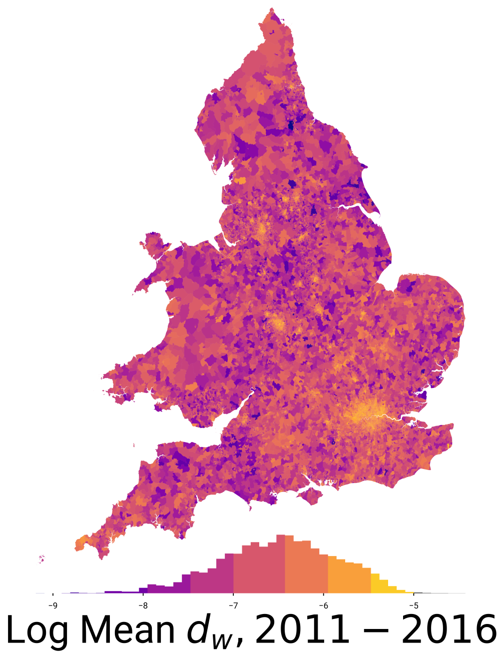

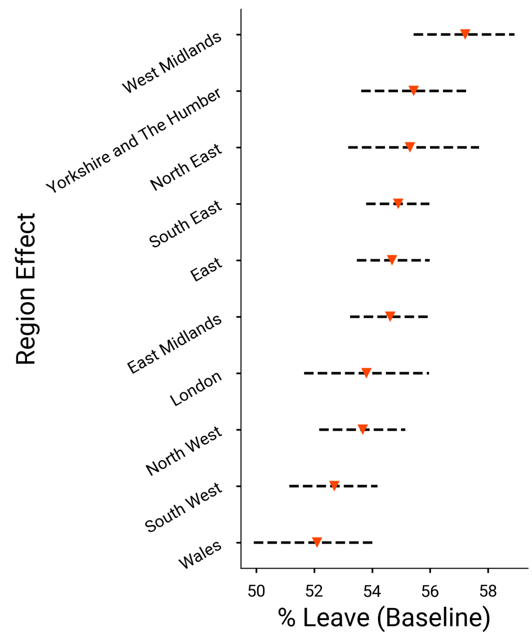

Change in % manufacturing employment, 2011 to 2016
Unemployment rate among WAP whites

We anticipated that deindustrialization and unemployment may play a role in whether a place voted Leave. Considering the issues raised in the Brexit campaign about employment, regulation, and *what* Britain does at work, I anticipated places with high unemployment or places where manufacturing has left in the last 5 years to vote Leave.

Percent of WAP who are white and not born in the UK
Percent of WAP who are ethnic and not born in the UK

For these raw population origination components, I did not anticipate them being relevant if the hypothesis of rate-of-change in migration (over state-of-being) were correct.

Working Age Population Age Structure

We expected places where the elderly live to be more likely to support Brexit, since they may be more emotionally connected to a period in time where life outside of the EU is now recalled on favorably. These covariates included the percentage of working-age population between 16-20, 20-24, & those 50+.

Votes Cast

Acknowledging Harris and Charlton (2016), the analysis of vote shares directly does not acknowledge the possibility that differently-sized constituencies may have different properties. This argument also has precedent in the analysis of heteroskedasticity and omitted-variate bias between vote shares depending on the structure of election size (Linzer 2012, e.g.). Therefore, to attempt to control for *only* the ommission of a constituency-size effect, I include the raw number of votes cast in a constituency (in thousands) as a predictor.

7 Appendix: Multiple-Bandwidth GWR results

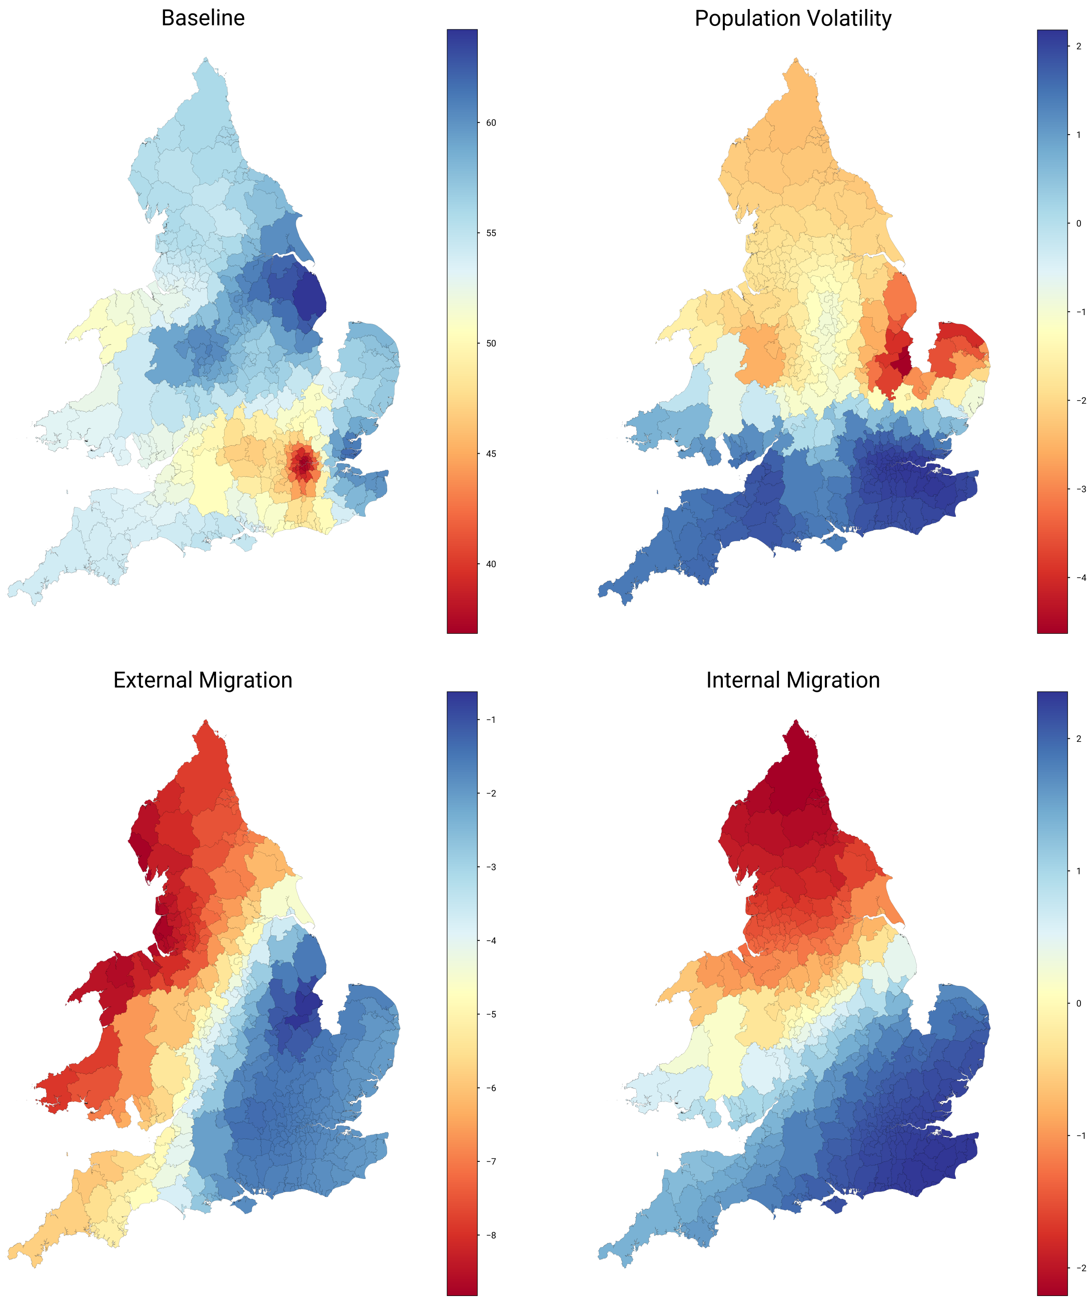

Figure 1: Coefficient surfaces for the multiple-bandwidth GWR models are shown above. In these models, each coefficient process is allowed to have its own spatial scale. Thus, each of these processes ostensibly operates at a different scale. Above, they are ordered from “most local” to “least local,” starting in the top left. An adaptive nearest-neighbor bandwidth is used, this means that both the increased density of elements in London and the Liverpool-Manchester conurbation can be accounted for without oversmoothing. With this, we see that internal/external migration has an (essentially-global) gradient although each has a slightly different tilt. Population volatility and the local intercept have distinct spatial patterns.

Notes

[1] By *implicit* hypothesis, I refer to the ways in which the rate of change of migrants in an area might plausibly drive individuals' choices in that area vote to Leave the European Union. If the claim is one suggesting a theory of racial (or national) threat (Giles and Hertz 1994; Rocha and Espino 2009), then we might would expect that the ethnicity or nationality of people, migrants or not, may play a role in determining whether some areas were more likely to vote to Leave. Second, if the claim is that rapid changes in a place's population mix drives revanchist sentiment, we might expect that the racial, ethnic, or national composition of these changes may not matter. Finally, I provide a preliminary examination of whether population dynamics affect regions differently. It may be the case that population dynamism drives Remain voting in some areas. Elsewhere, though, fast changes in population mix may drive a communal feeling of *anomie* and placelessness, emboldening a desire to make Britain (or that corner of it) great again.

[2] I modeled this demographic difference, but omit it here for brevity.

[3] The code for the analysis is published on Github.

[4] In a way, this punts the question, since fixed effects *in and of themselves* are not necessarily generative theories of why these fixed effects are salient.

[5] While MGWR does not have analytically-defined standard errors yet (due to its novelty), the patterns were largely robust when bootstrapped (Wolf, Oshan, and Fotheringham 2017)

[6] Indeed, (M)GWR presumes that such an underlying structure may exist and that it is smooth and detected conditionally-independently across covariates, which is not demonstrated in this brief note.

[7] One vector is available for all 2011 LSOAs since 1998. The ethnicities tracked are White British, White Irish, White Other, Asian British Bangladeshi, Asian British Chinese, Asian British Indian, Asian British Pakistani, Asian British Other, Black British African, Black British Caribbean, & all others, meaning $K = 11$. Often, however, the white-relevant categories account for more than 90% of the population. Further, sometimes q_i^t must be renormalized to account for rounding. This data does not contain information about the national origin of these percentages, nor information about the total count in the LSOA, so the data cannot be analyzed through reaggregation directly to the local authority and used alongside the LAMI directly.

[8] Wasserstein distance is sometimes called the *earth movers distance*, and KL divergence sometimes the *relative entropy*.

[9] The deviation standardization is not crucial, as the effect ends up being significant either way. However, it is the case that areas with local authorities with more population change tend to also have more heterogeneity in which LSOAs are experiencing that change, so accounting for that seems reasonable.