

Did YouTube Attention Predict How Counties Voted in 2016?

Evidence from the Trump and Clinton Campaigns

Tristan Hotham*

2026

Preprint, not peer reviewed

Abstract

Do the political videos people watch online tell us how their area will vote? I test this with a record, by county and by day, of YouTube views of Donald Trump and Hillary Clinton content from 1 January to 15 October 2016, covering 2,184 US counties in 49 states. Counties where Trump took a larger share of candidate views voted more Republican, but nearly all of that is geography and past voting. Within a state, holding demographics and the 2012 vote equal, a 10 point higher Trump share goes with a 0.53 point larger Republican swing from 2012 to 2016 (95% CI 0.33 to 0.74). The 2008 to 2012 swing shows no such relation. Comparing neighbouring counties gives smaller but still positive estimates: 0.23 for neighbours in the same state, and 0.33 to 0.40 for neighbours across a state line once each border is held fixed. Allowing for noise in the view share raises the baseline to about 0.6. Attention measured before Trump was formally nominated predicts the swing about half as strongly as general election attention, which fits a stable county orientation better than exposure to the campaign. The swing comes mainly from lost Democratic votes (1.7% fewer per 10 points of share), with a smaller Republican gain and no change in total votes. Coarse Google Trends rankings from 2015 do not predict the swing, and the YouTube share does not improve prediction for counties in states the model has not seen (R^2 0.76 either way). In short, YouTube views are a small local sign of where a county was already heading. They do not predict votes beyond that, and they are not evidence that videos changed anyone's vote.

Keywords: 2016 US election; Trump; YouTube; digital trace data; county voting; electoral swing; ecological inference

1 Introduction

Search and video data have been used to measure things that surveys measure badly, from racial animus (Stephens-Davidowitz 2014) to the consumption of radical content (Hosseini et al. 2021). For electoral research the appeal is plain. Attention is recorded at a fine geography, in real time, for everyone with a connection. What it tells us that we did not already know is less clear. Counties that watch more Trump content are more Republican, but they were more Republican before he ran. A useful signal has to add information beyond a place's past. A claim about campaign effects faces a further hurdle, since

*Independent researcher. ORCID 0000-0003-4931-9349. Correspondence: contact@tristanhotham.com. The views expressed are the author's own and do not represent those of any employer.

the best field evidence finds that campaign contact persuades almost nobody in general elections (Kalla and Broockman 2018).

I use a county-by-day panel of YouTube views of Trump and Clinton content during the 2016 campaign to answer four questions. **RQ1.** Does the Trump share of candidate views predict the 2016 Republican vote, and how much of that survives geography and past voting? **RQ2.** Does it predict the swing from 2012 to 2016, and does that hold up against the usual threats? **RQ3.** Is the relationship timed like a campaign effect or like a stable county orientation? **RQ4.** Does it improve out-of-sample prediction of the vote?

The raw cross-section shows a large association, and most of it is geography. With state, demographics and the 2012 vote held fixed, a small, precisely estimated positive association with the swing remains, and it survives the robustness checks. A placebo on the 2008 to 2012 swing is not significant. The share measured in the primary season predicts the swing about half as strongly as the general election share, and still clearly, so I read the association as a marker of county orientation. The swing appears mostly as lost Democratic votes and is concentrated in counties that were less Republican in 2012. Across states, the measure does not help prediction.

The closest work is the county-level study of social media and the 2016 vote by Fujiwara et al. (2024), which uses variation in Twitter adoption to estimate an effect, and the work on media exposure and place-level voting that it builds on (Allcott and Gentzkow 2017). My design is far weaker on causation and cannot say whether attention moved votes. What it can do is describe how a fine-grained video signal sorts against the swing, which several accounts of 2016 predict should follow from economic and cultural conditions in particular places (Autor et al. 2020; Mutz 2018). Earlier electoral predictions from social media data have a poor record (Gayo-Avello 2013), which is one reason I test out-of-state prediction directly.

The paper extends two earlier preprints on Facebook data (Hotham 2026a,b), which looked at what audiences reward and how campaigns differed in tone, and my doctoral research on parties and social media (Hotham 2021). Here the outcome is an election result and the unit is a place rather than a post.

2 Data

2.1 YouTube attention

The attention data are a county-by-day panel of views of Trump and Clinton content on YouTube for 1 January to 15 October 2016 (288 days, as 29 February is missing; 2,899 counties and county equivalents), with a separate table of the most viewed videos by date. I hold the panel as an archive named “youtube-election-real-time” from my earlier research, and I have not been able to confirm its original source or collection method. I do not know which videos count as Trump or Clinton content, or whether counts are exact or sampled. The text below describes the data as I hold them. Views are attributed to the viewer’s county. I do not know whether counts are thresholded or noised. That matters for small counties, and it is why I restrict the main sample.

For each county I sum views of Trump content (T) and Clinton content (C) and define the *Trump share of candidate views* as $100T/(T + C)$. I also compute it for three phases of the campaign: before 8 June (the primaries), 8 June to 31 August (the end of the primaries, the conventions and summer), and 1 September to 15 October (the general election run-in). National attention is spiky (Figure 1), peaking around Obama’s endorsement of Clinton, the Republican convention and the debates. In one robustness




Figure 1. Daily views of Trump and Clinton content summed across the county panel, 1 January to 15 October 2016, log scale. Dotted lines mark selected events.

check I drop the 15 largest days.

The main sample is the 2,184 counties in 49 states with at least 200 Trump and Clinton views combined. Alaska drops out because its votes are not reported by county, and the District of Columbia is not in the merged data. The median county in the sample has about 2,500 views; a quarter have fewer than 800. I report results at thresholds from 50 to 1,000 views. The share is strongly clustered by state, with state fixed effects explaining 68% of its variance (Figure 2), so I always compare counties within the same state.

2.2 Votes and covariates

County presidential returns for 2008, 2012 and 2016 give the Republican two-party share (Republican votes over Republican plus Democratic votes). The outcome is the *swing*: the 2016 share minus the 2012 share, in percentage points. The placebo swing is 2012 minus 2008. The covariates are log population, log density, the percentages White, Black and Hispanic, the percentages aged 65 and over and aged 18 to 29, the percentage with a bachelor's degree, the poverty rate, average income, the 2016 unemployment rate and its change since 2012. They come from a public county compilation that merges data from the Census Bureau, the Bureau of Labor Statistics and state election returns. The compilation does not date every field, and some variables, notably education and income, may postdate 2016. Table 1 gives descriptives. Counties are unweighted unless stated, so the sample leans rural; the median population is 43,000.

3 Empirical strategy

Every model is a county-level least squares regression with standard errors clustered by state, and all but two (marked in Table 2) include state fixed effects. The baseline for RQ2 is

$$\text{Swing}_{cs} = \alpha_s + \beta \text{TrumpShare}_{cs}/10 + \gamma \text{GOP2012}_{cs} + X_{cs}\delta + \varepsilon_{cs},$$

Table 1. Descriptive statistics, 2,184 counties in the main sample.

	Mean	SD	25th	Median	75th
Trump share of candidate views (%)	46.8	13.3	37.9	45.6	53.1
GOP two-party share 2016 (%)	63.9	15.3	55.2	66.8	75.4
GOP two-party share 2012 (%)	58.6	14.0	49.7	59.5	68.8
GOP two-party share 2008 (%)	55.7	13.1	47.2	56.1	65.1
Swing 2012 to 2016 (points)	5.3	5.5	1.8	5.2	9.0
Swing 2008 to 2012 (points)	2.9	3.0	1.0	2.6	4.5
Population 2016 (thousands)	139.6	381.3	22.6	43.1	107.5
White (%)	76.2	18.9	65.2	82.5	91.4
Black (%)	9.1	12.9	1.1	3.3	11.2
Hispanic (%)	10.0	13.4	2.7	5.0	10.8
Aged 65+ (%)	19.0	4.6	16.1	18.6	21.2
Bachelor's degree (%)	23.9	10.1	16.6	21.2	29.1
Poverty rate (%)	13.6	5.2	9.8	12.9	16.4
Unemployment 2016 (%)	5.2	1.7	4.1	4.9	6.0

where s indexes states and X holds the covariates listed above. So β is the difference in swing, in points, between two counties in the same state whose Trump view shares differ by 10 points and whose observed demographics and 2012 vote are equal. Conditioning on 2012 is equivalent to modelling the 2016 share with the 2012 share on the right-hand side, and I build up the 2016 share model in stages to show where the association goes.

The baseline faces five checks. (1) A placebo: the same model on the 2008 to 2012 swing, which 2016 attention cannot have caused. (2) An Oster test of how strong selection on unobservables would need to be, relative to the observed controls, to push β to zero (Oster 2019). (3) Standard errors clustered by 3-degree spatial grid cell instead of by state, in the spirit of spatial correction (Conley 1999). (4) Leave-one-state-out re-estimation. (5) Re-estimation without the 15 national spike days, with votes as weights, at other view thresholds, and without extreme shares.

For RQ3 I compute the Trump share separately for each phase of the campaign and enter the phases first one at a time, then together. If campaign exposure drives the result, the general election share, closest to the vote, should matter most and the primary share little. If a county orientation drives it, all three should matter. To look at mechanism I model the change in log Republican votes, log Democratic votes and log total votes between 2012 and 2016.

To see how much of the within-state association is local, I compare neighbouring counties, in the manner of border-county designs (Dube et al. 2010; Huber and Arceneaux 2007; Spenkuch and Toniatti 2018). The first comparison pairs counties in the same state whose centroids lie within 50 km of each other. The second pairs each county with counties in other states within 50, 75 or 100 km. For each pair I difference the swing, the view share, the 2012 vote and the controls, then regress the swing difference on the share difference without a constant, clustering standard errors by state pair. A pair straddling a state line differs in the state-level component of the view share, which accounts for most of the variation in the share difference (70% at 75 km). That component may partly reflect how the data were built, so I report cross-border estimates both with no border effect and with a fixed effect for each state pair, which compares pairs along the same border segment and removes the state-level difference. Some counties appear in several pairs, so I also use only non-overlapping nearest-neighbour matches. The 2008 to 2012




Figure 2. Trump share of candidate YouTube views (left) and Republican two-party swing from 2012 to 2016 (right), by county centroid. The view share is strongly clustered by state; the swing is concentrated in the Upper Midwest and Appalachia.

placebo is repeated in this design.

To address measurement error in the share, I split each county’s views into two random halves, compute the share separately in each, and use the correlation between halves, stepped up by the Spearman-Brown formula, as a reliability. I correct the baseline for it and also instrument one half with the other. For sensitivity to unobserved confounding I report the bias-adjusted estimate of [Oster \(2019\)](#) under equal selection and the robustness value of [Cinelli and Hazlett \(2020\)](#), and for inference I use randomisation inference, permuting the share within states, and a baseline-minus-placebo contrast.

The second design uses Google Trends rankings from 2015 and early 2016 as coarse early attention flags: whether Donald Trump was the most searched Republican candidate in a county between 17 September and 23 October 2015, whether he was in the top two, whether he was first in the run-up to the 16 September debate and in 17 to 22 September, and whether Bernie Sanders was the most searched Democrat from June to October 2015 or from 19 December 2015 to mid-January 2016. These are rankings, not volumes ([Mellon 2013](#)). Trump was the most searched Republican in 89% of counties, so there is little variation to work with. For RQ4 I run ten-fold leave-state-out cross-validation, comparing a model with demographics and the 2012 vote against the same model plus the YouTube measures, using least squares and gradient boosting, and repeat it within states.

All results are correlations between aggregates. Nothing here identifies an individual-level relationship ([Robinson 1950](#)), and I do not identify a causal effect of watching.

4 Results

4.1 The cross-section is mostly geography

Table 2 builds the model up in stages. With no controls, a 10 point higher Trump share goes with a 3.4 point higher Republican share in 2016. Holding states fixed raises that to 6.4. Demographics cut it to 2.8, and the 2012 vote cuts it to 0.53. About 92% of the within-state association is gone once demographics and the 2012 vote are held equal. Most of what looks like predictive power is a place’s past.

Table 2. Trump share of views and the Republican vote. Coefficients are per 10 points of Trump view share. State fixed effects in all models except the first row of the upper panel and the second row of the middle panel; standard errors clustered by state.

Outcome and controls	β	95% CI	R^2	N
<i>Republican two-party share, 2016 (points)</i>				
No state effects, no controls	3.38	1.64 to 5.12	.09	2,184
State effects	6.40	4.43 to 8.38	.39	2,184
State effects + demographics	2.84	1.68 to 4.00	.84	2,184
State effects + demographics + 2012 share	0.53	0.33 to 0.74	.98	2,184
<i>Swing 2012 to 2016 (points)</i>				
State effects + 2012 share	1.23	0.79 to 1.66	.46	2,184
No state effects, full controls	0.23	−0.06 to 0.53	.79	2,184
Baseline: state effects, full controls	0.53	0.33 to 0.74	.87	2,184
<i>Placebo: swing 2008 to 2012 (points)</i>				
State effects + 2008 share	0.005	−0.30 to 0.31	.48	2,184
State effects, full controls	0.17	−0.11 to 0.46	.61	2,184

What remains is small, though clearly above zero. The baseline swing model gives 0.53 points per 10 points of share (CI 0.33 to 0.74, $R^2 = .87$). Within a state the standard deviation of the Trump share is 7.5 points, so a county one such deviation above its neighbours goes with about 0.4 points of extra swing. The swing has a standard deviation of 5.5 points, or 2.0 after state effects and controls, so one within-state deviation of the share moves the swing by about a fifth of its unexplained spread. The binned relationship in Figure 3 is close to linear across the range, and no small group of extreme counties drives it.

4.2 The swing result is robust, and the placebo is clean

Figure 4 and Table 3 show the checks, and the estimate barely moves. Weighting counties by 2016 votes gives 0.41 (CI 0.10 to 0.72): less precise, but enough to rule out a small-county artefact. Clustering on 126 spatial grid cells leaves the coefficient as it is and shrinks the standard error to 0.08. Dropping the 15 national spike days, when events rather than local interest drove attention, gives 0.40 (0.24 to 0.57). Across view thresholds the estimate runs from 0.35 (50 or more views) to 0.50 (500 or more), always significant. Dropping one state at a time gives 0.50 to 0.58, and restricting to counties with Trump shares between 10 and 90 gives 0.52.

Two results qualify this. Without state fixed effects the estimate falls to 0.23 and loses significance (CI −0.06 to 0.53). Much of the national relationship between view share and swing runs between states, where the share is a noisy proxy for regional differences, so this is a within-state finding. The placebo, too, is not a perfect zero. On the 2008 to 2012 swing the controlled estimate is 0.17 (CI −0.11 to 0.46): smaller than the baseline and not significant, but loosely bounded. It counts against a large pre-existing trend. It does not show there was none.

The Oster test asks about selection on unobservables. Going from the model with state effects and the 2012 share to the full model raises R^2 from 0.46 to 0.87 and cuts the coefficient from 1.23 to 0.53. Since 1.3 times the full-model R^2 would exceed 1, I set the maximum R^2 to 1. On that basis the unobservables would need to be about 2.4 times as important as the observed demographic controls to bring the coefficient to zero, above the benchmark of 1 that Oster (2019) proposes. That is reassuring up




Figure 3. Binned relationship between the Trump share of YouTube views and the 2012 to 2016 Republican swing, after removing state effects and the control variables from both. Twenty equal-sized bins; bars are 95% confidence intervals.

Table 3. Robustness of the baseline swing estimate (per 10 points of Trump view share).

Specification	β	95% CI	N	Note
Baseline	0.53	0.33 to 0.74	2,184	state clusters
Weighted by 2016 votes	0.41	0.10 to 0.72	2,184	
Spatial grid clusters (3 degrees)	0.53	0.38 to 0.69	2,184	126 clusters
Excluding 15 spike days	0.40	0.24 to 0.57	2,184	
50 or more views	0.35	0.20 to 0.49	2,595	
100 or more views	0.44	0.26 to 0.62	2,397	
500 or more views	0.50	0.28 to 0.72	1,829	
1,000 or more views	0.43	0.20 to 0.66	1,528	
Shares between 10 and 90	0.52	0.31 to 0.73	2,181	
Leave one state out	0.50 to 0.58			range over 49 fits

to a point. The comparison is with demographics only, because the 2012 vote, which already absorbs much of any latent county orientation, sits in both models. Assuming selection on unobservables equals selection on observables, the bias-adjusted estimate is 0.31. The robustness value of [Cinelli and Hazlett \(2020\)](#) is 16%: an unobserved confounder would need to explain 16% of the residual variance of both the share and the swing to bring the estimate to zero. Neither is a proof, but both say the estimate is not fragile to modest omitted variables, and adding the manufacturing share, life expectancy and the share in fair or poor health leaves it at 0.51 (CI 0.31 to 0.71, 2,165 counties with all three). Spatially correlated errors do not change the picture: standard errors that allow correlation up to 50, 200 and 1,000 km are 0.07, 0.10 and 0.11, and permuting the share within states 2,000 times never produces an estimate as large as 0.53. The baseline is 0.36 points larger than the placebo (SE 0.15, $p = .017$), so the 2016 relationship is stronger than the 2012 one, though by a margin that is not large.




Figure 4. Estimates of the association between a 10 point higher Trump view share and the 2012 to 2016 Republican swing, across specifications. Orange: swing 2012 to 2016. Grey: placebo on the 2008 to 2012 swing. Bars are 95% confidence intervals.

4.3 It looks like a county orientation, not campaign exposure

Figure 5 (left) splits the share by campaign phase. One at a time, the primary-season share (before 8 June) goes with a swing of 0.27 points (CI 0.16 to 0.37), the summer share with 0.28 (0.16 to 0.41) and the general election share with 0.52 (0.30 to 0.74). Entered together they give 0.20, 0.19 and 0.38, all significant. The general election share carries the most weight, but the primary share, measured before Trump was formally nominated, carries a good part of it. That is the pattern a stable county characteristic would produce: places that took an early interest in Trump were the places that swung. A pure exposure account would need the early share to matter little. Two things limit this reading. Trump was a constant presence in the primaries and the presumptive nominee from early May, so early attention may itself have been a first dose of exposure. And the general election share is the one most shaped by the debates and other events. I do not push the gradient far.

A latent orientation might seem at odds with the null placebo, since a stable trait should have predicted earlier swings too. It need not be. An orientation that only mattered once a candidate like Trump was on the ballot would predict the 2016 swing and not the 2012 one, and the data cannot say what that orientation is.

How much a county watched does not seem to matter, as against what it watched. Log views per capita has no relationship with the swing (coefficient -0.06 , $p = .53$) and leaves the share coefficient at 0.53. Entered together in one model, log Trump views per capita goes with a larger swing (1.23 per log unit, $p < .001$) and log Clinton views per capita with a smaller one (-1.26 , $p < .001$). That is the same relative-interest story in another form.

4.4 The swing is lost Democratic votes

The right panel of Figure 5 decomposes the swing. A 10 point higher Trump share goes with a 0.8% larger increase in Republican votes ($p = .002$), a 1.7% larger fall in Democratic votes ($p < .001$) and no relationship with total votes (-0.2% , not significant). The Republican gain is real but half the size of the Democratic loss, and total votes did not rise. In these counties the swing came mostly from Democratic




Figure 5. Left: swing coefficient by campaign phase of the Trump view share, estimated separately. Right: change in votes 2012 to 2016 on the Trump view share, in per cent per 10 points of share. Bars are 95% confidence intervals.

vote totals falling, not from a surge in turnout. That fits accounts in which the 2016 swing reflected a weakening Democratic vote in places already receptive to Trump. It fits poorly with the idea that online attention helped the campaign bring out large numbers of new Republican voters. County totals cannot tell us whether the missing Democratic votes went to Trump, to third parties or to abstention.

4.5 Where the association is strongest

The coefficient depends on how Republican a county already was. Its interaction with the 2012 Republican share is -0.28 per standard deviation ($p < .001$): the association is largest in less Republican counties and smaller in the most Republican ones. One possible reason is that the most Republican counties had less room to move, though I do not test this. The interaction with the White share is $+0.12$ per standard deviation ($p = .045$), small and only just significant. Interactions with education and density are not significant. By tercile the coefficient is 0.41, 0.64 and 0.50 for education and 0.53, 0.49 and 0.33 for density, so it appears in every group. These are descriptive. I have not corrected them for multiple comparisons, and the education and density terciles do not differ reliably.

4.6 It does not help prediction

If the YouTube share captured something that demographics and past votes miss, it should improve forecasts for counties in states the model has not seen. It does not. In ten-fold cross-validation holding out whole states, demographics and the 2012 vote predict the swing with $R^2 = 0.762$ and RMSE 2.66 points. Adding the YouTube measures gives 0.759 and 2.67. A gradient boosting model goes from 0.801 to 0.798. The signal is local to states and does not travel between them. For anyone forecasting counties in a state the model has not seen, the YouTube share adds nothing to what we already knew. This test is demanding by design: the out-of-state model cannot use the within-state variation where the association lives. Within states, where it can, the share does help a little. Repeated random ten-fold splits that keep each state in the training data raise R^2 from 0.856 to 0.860 and cut the RMSE from 2.07 to 2.04 points. The gain is real and small, in line with an effect of a fifth of the unexplained spread.




Figure 6. Association between a 10 point higher Trump view share and the 2012 to 2016 swing, by scale of comparison. Orange: counties in the same state. Blue: counties across a state line. Grey: placebo. All with the 2012 vote and demographics held equal; bars are 95% confidence intervals with standard errors clustered by state or state pair.

4.7 Neighbouring counties: the association is local, and smaller

Figure 6 and Table 4 report the neighbour comparisons. Among counties in the same state within 50 km of each other (2,930 pairs), a 10 point difference in view share goes with a swing difference of 1.46 points with no controls and 0.23 (CI 0.09 to 0.36) once the 2012 vote and demographics are held equal. That is under half the baseline of 0.53, so about half of the within-state association comes from broad regional differences within states, and about half survives a comparison of close neighbours.

Across state lines the answer depends on how the pairs are compared. For counties within 75 km of each other in different states (1,828 pairs, 91 state pairs), the full-controls estimate is 0.12 (CI -0.08 to 0.33) with no border effect, and the estimates at 50 and 100 km are 0.22 and 0.17. That design leans on the state-level difference in the view share, which explains 70% of the variation in the share difference at 75 km and which may partly be an artefact of how the data were built. With a fixed effect for each state pair, the estimates are 0.40 (CI 0.21 to 0.59) at 50 km, 0.33 (0.13 to 0.53) at 75 km and 0.39 (0.24 to 0.53) at 100 km, all clearly positive. Using only non-overlapping nearest-neighbour matches (456 pairs, median distance 49 km) gives 0.24 (0.02 to 0.46). The 2008 to 2012 placebo with border effects is -0.07, 0.09 and 0.14 at the three distances, none distinguishable from zero.

I read the pattern like this. The association is present at every scale, from whole states down to neighbouring counties, and it is smaller the more local the comparison: 0.53 within states, 0.23 between same-state neighbours and 0.33 to 0.40 along borders. Some of the decline is probably noise. The share is measured with error, and local comparisons throw away most of the true variation while keeping the noise. Split-half reliability of the share after controls is 0.87, which lifts the baseline to 0.61 (and to 0.64 with one half instrumenting the other), and the attenuation is worse for the within-pair differences. I do not claim that this recovers a true local effect. The reliable conclusion is that the association is not a state-level artefact, and that it is smaller, probably between a half and three quarters of the baseline, among close neighbours.

Table 4. Neighbouring-county comparisons. Coefficients are the swing difference per 10 points of difference in Trump view share. Differences are taken within pairs of counties; standard errors clustered by state or state pair.

Specification	β	95% CI	Pairs	State pairs
<i>Same state, within 50 km</i>				
No controls	1.46	1.00 to 1.92	2,930	
+ 2012 share and demographics	0.23	0.09 to 0.36	2,930	
<i>Across a state line, within 50 km</i>				
No controls	0.83	0.17 to 1.49	559	72
+ 2012 share	0.82	0.15 to 1.50	559	72
+ demographics	0.22	0.01 to 0.43	559	72
+ border-pair effects	0.40	0.21 to 0.59	559	72
Placebo 2008 to 2012, border-pair effects	−0.07	−0.50 to 0.35	559	72
<i>Across a state line, within 75 km</i>				
No controls	0.86	0.33 to 1.40	1,828	91
+ 2012 share	0.83	0.25 to 1.41	1,828	91
+ demographics	0.12	−0.08 to 0.33	1,828	91
+ border-pair effects	0.33	0.13 to 0.53	1,828	91
Placebo 2008 to 2012, border-pair effects	0.09	−0.17 to 0.35	1,828	91
<i>Across a state line, within 100 km</i>				
No controls	0.97	0.49 to 1.46	4,137	107
+ 2012 share	0.93	0.40 to 1.46	4,137	107
+ demographics	0.17	−0.03 to 0.36	4,137	107
+ border-pair effects	0.39	0.24 to 0.53	4,137	107
Placebo 2008 to 2012, border-pair effects	0.14	−0.06 to 0.34	4,137	107
Non-overlapping matches, + demographics	0.24	0.02 to 0.46	456	

4.8 Google Trends interest in 2015 does not predict the swing

Table 5 adds the Google Trends flags to the baseline swing model, with state effects and full controls. None is significant. Trump being the most searched Republican in autumn 2015 goes with a swing 0.21 points larger (CI −0.08 to 0.50), being in the top two with −0.34 (−1.03 to 0.35), and being first before the September debate with 0.17 (−0.12 to 0.46). The Sanders flags give −0.13 (−0.28 to 0.03) and −0.12 (−0.35 to 0.11). With the Trump flag in the model the YouTube coefficient stays at 0.53 (CI 0.33 to 0.73), so one measure does not stand in for the other. Counties where Sanders was the most searched Democrat in December 2015 and January 2016 have a Trump view share 0.9 points lower ($p=.009$); the other flags are unrelated to the share. These are weak tests. Each flag is a single ranking, Trump topped the Republican list in nearly nine counties in ten, and each window covers only a few weeks or months. What they show is that the YouTube share is not just a copy of early search interest. They cannot tell us whether early attention predicts the vote.

5 Discussion

The within-state estimate is positive, small and robust, and it does not come from state-level confounding: it shrinks but holds among close neighbours, including across a state line, and the placebo stays null. That does not make it evidence that YouTube moved votes. Three features argue against that reading.

Table 5. 2015 Google Trends flags and the 2012 to 2016 swing. Baseline specification (state effects, full controls, 2012 share). Coefficients are the swing difference in points between flagged and unflagged counties.

Flag	β	95% CI	N	Share flagged
Trump most searched GOP, 17 Sep to 23 Oct 2015	0.21	−0.08 to 0.50	2,164	89%
Trump in top two, same period	−0.34	−1.03 to 0.35	2,164	98%
Trump most searched before September debate	0.17	−0.12 to 0.46	2,164	90%
Trump most searched, 17 to 22 Sep 2015	−0.06	−0.28 to 0.16	2,150	78%
Sanders most searched Dem, Jun to Oct 2015	−0.13	−0.28 to 0.03	2,163	54%
Sanders most searched Dem, Dec 2015 to Jan 2016	−0.12	−0.35 to 0.11	2,162	28%

The association already appears in primary-season attention. It runs mainly through lost Democratic votes, with a smaller Republican gain and no rise in total votes. And it does not help predict counties in other states. The simplest account is that the Trump share of attention marks places that were already moving, for the economic and cultural reasons that shaped the Upper Midwest and Appalachia in 2016, and that the swing and the attention share common causes (Autor et al. 2020; Mutz 2018). That is in line with field evidence that campaign contact has close to no persuasive effect in general elections (Kalla and Broockman 2018). A causal effect of exposure of the order of the estimate cannot be excluded, but nothing here separates it from the common-cause account.

What the measure does offer is a legible local signal of relative interest in the two candidates, at a fine geography and in real time. A swing of 0.5 points per 10 points of share is not nothing, though it is small next to the swing itself, and it is smaller among close neighbours. Part of what makes a signal like this useful is knowing how little it adds, and how much of it is measurement noise.

5.1 Limitations

Ecological inference. Counties are not voters (Robinson 1950). **Provenance.** I do not know how the YouTube data were sampled or thresholded, and the strong clustering of the share by state suggests that part of the signal is regional, in how the data were built or in what content was on offer. State fixed effects remove this, at the price of using only within-state differences. **Selection of videos.** I do not know which videos count as Trump or Clinton content, and a share of views is a share of that set. Views of Trump content need not be sympathetic; some will be critical or mocking, so the share measures attention, not support. **Covariates.** Some covariates may postdate 2016. **No identification.** Nothing here is quasi-experimental. **Measurement error.** The share is noisy, especially in small counties, and the noise attenuates local comparisons more than state-level ones, so the neighbour estimates are probably too low; my reliability correction assumes the noise is classical. **Neighbour designs.** Pairs overlap, standard errors are clustered by state pair, and the border-pair fixed effects compare pairs along the same border, so counties near the largest borders count for more. **Trends flags.** The Google Trends data are rankings, not volumes. **Single election.** A result for 2016 may not hold in other cycles, so replication on 2020 is the obvious next step.

5.2 What would strengthen the design

A stronger study would (i) use attention volumes measured well before the campaign, not rankings, (ii) add a second platform with volumes, such as search, news or social media, (iii) use a source whose sampling and thresholds are documented, so that state-level differences in the share can be interpreted,

(iv) add individual-level survey or voter-file data to move from places to people, and (v) replicate on 2020 and 2024. None of this is possible with the data I hold. The Google Trends rankings I have are too coarse to do the first two.

6 Conclusion

County-level YouTube attention to Trump over Clinton is a small correlate of the 2016 Republican swing. It holds within states, shrinks but survives among close neighbours, and passes a placebo and a range of robustness checks. It does not help predict across states, and it reads better as a marker of county orientation than as a measure of persuasion. Its modesty is part of its value. It shows how little a fine-grained attention signal adds once we know where a place already stood.

Provenance and AI use

The YouTube panel is an archive I hold from research during my doctorate, described in Section 2. This paper builds on my doctoral research ([Hotham 2021](#)). The analysis and writing were completed in 2026. The analysis code and drafts of the text were prepared with the assistance of an AI system (Claude, by Anthropic), which was also used to review drafts. I checked the analysis and the text and take responsibility for both.

Data and code availability

I do not redistribute the underlying YouTube panel, because I cannot confirm the terms under which it was collected. The code, and the county-level estimates and figures reported here, are available on request.

References

- Allcott, H. and Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2):211–236.
- Autor, D., Dorn, D., Hanson, G., and Majlesi, K. (2020). Importing political polarization? The electoral consequences of rising trade exposure. *American Economic Review*, 110(10):3139–3183.
- Cinelli, C. and Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B*, 82(1):39–67.
- Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1):1–45.
- Dube, A., Lester, T. W., and Reich, M. (2010). Minimum wage effects across state borders: Estimates using contiguous counties. *Review of Economics and Statistics*, 92(4):945–964.
- Fujiwara, T., Müller, K., and Schwarz, C. (2024). The effect of social media on elections: Evidence from the United States. *Journal of the European Economic Association*, 22(3):1495–1539.
- Gayo-Avello, D. (2013). A meta-analysis of state-of-the-art electoral prediction from Twitter data. *Social Science Computer Review*, 31(6):649–679.

- Hosseinmardi, H., Ghasemian, A., Clauset, A., Mobius, M., Rothschild, D. M., and Watts, D. J. (2021). Examining the consumption of radical content on YouTube. *Proceedings of the National Academy of Sciences*, 118(32):e2101967118.
- Hotham, T. (2021). *How Do Political Parties Use Facebook and What Does It Offer to Their Campaigns*. PhD thesis, University of Bath.
- Hotham, T. (2026a). What went viral on Britain First's Facebook page? Outrage, questions and the cost of mobilisation, 2013–2018. Preprint, Zenodo, October 2026.
- Hotham, T. (2026b). Who won the Facebook referendum? Engagement, tone and nationalism on UK party and campaign pages, February to June 2016. Preprint, Zenodo, October 2026.
- Huber, G. A. and Arceneaux, K. (2007). Identifying the persuasive effects of presidential advertising. *American Journal of Political Science*, 51(4):957–977.
- Kalla, J. L. and Broockman, D. E. (2018). The minimal persuasive effects of campaign contact in general elections: Evidence from 49 field experiments. *American Political Science Review*, 112(1):148–166.
- Mellon, J. (2013). Where and when can we use Google Trends to measure issue salience? *PS: Political Science & Politics*, 46(2):280–290.
- Mutz, D. C. (2018). Status threat, not economic hardship, explains the 2016 presidential vote. *Proceedings of the National Academy of Sciences*, 115(19):E4330–E4339.
- Oster, E. (2019). Unobservable selection and coefficient stability: Theory and evidence. *Journal of Business & Economic Statistics*, 37(2):187–204.
- Robinson, W. S. (1950). Ecological correlations and the behavior of individuals. *American Sociological Review*, 15(3):351–357.
- Spenkuch, J. L. and Toniatti, D. (2018). Political advertising and election results. *Quarterly Journal of Economics*, 133(4):1981–2036.
- Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using Google search data. *Journal of Public Economics*, 118:26–40.