



BIROn - Birkbeck Institutional Research Online

Ahlstrom-Vij, Kristoffer and Allen, W. (2021) As we like it: did the UK's 2016 EU referendum reveal the "Will of the People?". *PS: Political Science and Politics*, ISSN 1049-0965. (In Press)

Downloaded from: <https://eprints.bbk.ac.uk/id/eprint/46877/>

Usage Guidelines:

Please refer to usage guidelines at <https://eprints.bbk.ac.uk/policies.html>
contact lib-eprints@bbk.ac.uk.

or alternatively

As We Like It: Did the UK's 2016 EU Referendum Reveal the "Will of the People?"

Kristoffer Ahlstrom-Vij
Birkbeck, University of London
k.ahlstrom-vij@bbk.ac.uk

William Allen
Oxford University
william.allen@politics.ox.ac.uk

Abstract: Since the UK's 2016 referendum on EU membership, politicians on both sides of the debate have repeatedly invoked the "will of the people," either to claim that the matter is settled or to justify the need for subsequent confirmatory action. Following a long tradition of information effects research in political science, we argue that the extent to which voters are informed about politically relevant issues is a key factor in evaluating whether any given result accurately reflects an electorate's collective "will." By using counterfactual modeling applied to British Election Study data, we estimate that support for leaving the EU would have dropped by 7-10 percentage points, had voters been more fully informed. More generally, in the context of relatively infrequent referenda, we suggest that such models can usefully indicate the sensitivity of outcomes to different levels of information, and allow legislators to gauge the extent to which a given outcome likely reflects voters' collective will.

Rhetoric surrounding the UK's 2016 referendum on continued EU membership has frequently invoked the "will of the people." Addressing the House of Commons in March 2019, then-Prime Minister Theresa May stated that "my sense of responsibility and duty has meant that I have kept working to ensure that we deliver on the result and the will of the people" (March 27, 2019).¹ May's successor and current Prime Minister, Boris Johnson, appealed to the same notion when suggesting in the *Daily Telegraph* (September 15, 2019) that opposition parties were "united in wanting to cancel the referendum result...and overturn the will of the people." On the other side of the debate, Caroline Lucas (currently the sole MP for the UK's Green Party) stated that "[e]very recent opinion poll shows that the will of the people has changed since [the referendum]" (December 4, 2018).

These statements raise a key question: what precisely is "the will of the people" in electoral contexts? We argue that voters' true preferences—and, in the aggregate, the "will of the people"—comprise the preferences that *would* have been reported had people been fully informed on relevant matters (e.g., Ahlstrom-Vij 2020; Harsanyi 1997). As such, given widespread public ignorance on politically relevant matters (Delli Carpini and Keeter 1996; Achen and Bartels 2016), there are likely to be gaps between the choices made by the electorate at the polls, and what is in their collective "will." Such gaps have been investigated in a variety of electoral contexts under the banner of "information effects," including the US (Althaus 2003; Delli Carpini and Keeter 1996; Bartels 1996), Canada (Blais et al. 2008), Denmark (Hansen 2009), Sweden (Oscarsson 2007), and in European Parliamentary elections (Bhatti 2010).

Building on this work, we examine whether the results of the UK's 2016 referendum on leaving the EU—51.9% for "Leave" and 48.1% for "Remain"—were potentially sensitive to differences in voters' levels of political information. We do so by developing and applying a counterfactual model of how information influences vote choice to the 2017 British Election Study (BES) face-to-face survey wave (N = 2,067)—the wave immediately following the referendum. In a more fully informed electorate, we find that support for leaving the EU would likely have dropped by

¹ For this and all subsequent quotes from Members of Parliament, see <https://hansard.parliament.uk/>, unless otherwise stated.

up to 9.5 percentage points, producing a 57.6%-42.4% split in favor of Remain by 15.2 percentage points. While our results are a simulation, we view them as not only good reasons for taking the extent to which voters are informed more seriously, but also as a cautionary lesson for politicians and policymakers as they interpret electoral results and confidently make claims about how such results reflect their constituents' collective will.²

What is the “Will of the People?”

The “will of the people” consists of the preferences we would have reported, had we been informed.³ How so? Consider the preferences of an individual first. While the most straightforward way of finding out what someone wants is to ask them, interpreting their answers is difficult when they are mistaken about the nature or implications of the options in question. To illustrate, say you are presented with two environmental policies. Policy *A* is disastrous for the environment, while Policy *B* is good for the environment—but in all other ways, they are identical in line with their outcomes. As it happens, you have it the wrong way around: although you care deeply about the environment, you pick policy *A*. But this does not actually mean you prefer *A* over *B*. Indeed, imagine you realize your mistake after choosing *A*. At that point, you would likely deny that you had ever wanted *A*: the only reason you chose it was because you thought it had the properties actually possessed by *B*.

This is what underlies Robert Goodin’s observation that, under circumstances of ignorance, “we can serve a person’s ‘real’ preferences only by censoring the misleading indication of his preferences that is revealed in his choices” (1995: 137). John Harsanyi makes a similar point when distinguishing “those choices of a person that really express his true preferences [...] from those choices of his that fail to do so because they are based on incorrect information” (1997: 133). Generalizing this point to the collective, Delli Carpini and Keeter (1996: 5) suggest that “the real interest of an individual—and by extension of a group and of the polity as a whole—are reflected in the choices one would make if

² We are not suggesting that the “will of the people” is a myth (Weale 2018), but that, at least in this case, the referendum likely failed to uncover it.

³ While we provide a summary of our rationale here, see Ahlstrom-Vij (2020) for a more detailed defense.

he or she were fully informed about the consequences.” Similarly, we hold that the “will of the people” comprises the preferences its constituent individuals would hold, were relevant yet false beliefs corrected. Assuming a majoritarian electoral system, the will of a population is thereby indicated by what the majority prefers (if anything) under conditions of having full information.

When examining that will, we are squarely in the political context—a context where mistaken beliefs are in fact rife (e.g., Achen and Bartels 2016; Friedman 1998).⁴ To be sure, if the public reports wanting something, there usually is a strong case for government to act accordingly. However, in settings where what voters say they want might be due to their not having had the privilege of taking into account relevant information, directly reading the “will of the people” from electoral results becomes difficult. It would be handy, then, if there were a way to identify such gaps between what people *say* they want, and what they *would have said* if they had possessed full information. In the next section, we argue that work on “information effects” offer just such a way.

Measuring the Sensitivity of Electoral Outcomes to Information

The intuition behind counterfactually modeling information effects

There is a long tradition of political scientists using statistical models to examine “information effects,” or gaps between actual outcomes or distributions of attitudes and the counterfactual outcomes or attitudes that we likely would have seen under full information. In an electoral context, the procedure follows four steps. First, you construct a knowledge scale and place respondents on that scale. Such scales (see e.g., Delli Carpini and Keeter 1996) are typically built from questions one might see on civics exams—the roles that named people hold, or the platforms of political parties—

⁴ It has been suggested that the public is able to rely on cues and shortcuts to act *as if* informed (e.g., Popkin 1991), or that voters are able to vote retrospectively on minimal amounts of information (e.g., Key 1966; Fiorina 1981). See Bartels (1996) and Achen and Bartels (2016) for critical discussions, as well as Ahlstrom-Vij (2019) for a critical survey of suggestions for how democracies might cope with low levels of voter knowledge.

and are meant to indicate generalist political knowledge.⁵ Second, using data on reported votes, you build and fit a model that estimates the probability of someone reporting having voted some particular way as a function of that person’s demographics and level of knowledge. Third, you increase each person’s placement on the knowledge scale to the maximum level while leaving everything else as is: this simulates how that person *would* have voted had they been fully informed. (If you want to relax this assumption of being “fully” informed, you can move individuals’ knowledge scores to a different threshold, such as the 90% percentile.) Finally, by noting the difference between the actual electoral outcome and the model’s estimated outcome, you obtain a measure of the aggregate information effect. A substantial difference indicates high sensitivity to voters’ level of information, and a greater risk that the result might to a non-trivial degree be an artifact of mistaken beliefs.

What counts as “substantial” in the context of elections? Larry Bartels (1996: 220) finds information effects in the range of 2-5 percentage points in US presidential elections. André Blais and colleagues (2008) find an average information effect of 2.3 percentage points across parties in six Canadian elections, while Henrik Oscarsson (2007) finds an average net gain of 2.7 percentage points for right parties in six Swedish elections. Yosef Bhatti (2010) models three European Parliament elections (in Denmark, Finland, and Sweden) with an average information effect of 3.5 percentage points across parties and elections. Collectively, these studies suggest finding information effects of more than about 3 percentage points would be substantial—and, crucially, these effects do not consistently benefit any particular party or ideology.

⁵ Some object to this by arguing competent voting does not require knowing answers to these kinds of survey items (e.g., Lupia 2006). We do not contest this point. Rather, we view these questions as *diagnostic* of whether people know things that *are* necessary for competent voting. Why? These scales have high internal consistency while also correlating well with both interviewer ratings of respondents’ political knowledge and behaviors associated with knowledge such as participation (Delli Carpini and Keeter 1996). Moreover, information effects based on these scales tend to be consistently higher for groups that we have independent reason to believe will be less politically informed, such as high-school dropouts, people who are poorer, and people in rural areas; and consistently lower for groups we have reason to believe will be more politically informed, such as those who are college educated, affluent, and living in cities or suburbs (Althaus 2003).

Modeling a more informed EU Referendum

In contrast to previous scholars' focus on regularly scheduled elections, we consider whether, to what extent, and in which direction the UK's EU referendum result may have displayed an information effect. This is motivated not only by observing politicians' invocations of the "will of the people", as described in the introduction, but also by referenda representing a relatively rare event in British politics—and one in which information conveyed through campaign efforts has been experimentally shown to have sizeable effects on attitudes towards EU membership (Goodwin et al. 2018).

To this end, we used the British Election Study's 2017 face-to-face survey wave (N=2,194)—the wave following the referendum—which contains demographic information, responses to six questions measuring political knowledge, and respondents' reported vote in the referendum (Fieldhouse *et al.* 2018). Dropping 127 observations that lacked survey weights left 2,067 observations, of which 1.4% had missing values which were imputed with multiple imputation using *aregImpute* in R's *Hmisc* package (R Core Team 2017; Harrell et. al 2019). Meanwhile, we fitted an item response theory (IRT) model using *ltm* (Rizopoulos 2006) to estimate the underlying knowledge of respondents. Diagnostic tests revealed that four of the six questions could be used to construct a valid knowledge scale.⁶

Since we were looking to model a counterfactual—how the electorate likely *would* have voted, had we somehow been able to intervene on their level of knowledge by increasing its value—we are using an explicitly causal model.⁷ In terms of likely confounders, we controlled for gender (vanHeerde-Hudson 2020; Plutzer 2020), level of education (Hebbelstrup and Rasmussen 2016) and income (Vowles 2020; Plutzer 2020). To reduce noise in the models, we also controlled for ethnicity

⁶ The four (true/false) knowledge items were: "No-one may stand for parliament unless they pay a deposit," "The Liberal Democrats favour a system of proportional representation," "MPs from different parties are on parliamentary committees," and "The number of members of parliament is about 100." The test information function suggested good precision (with a peak around mean ability), and tests suggested unidimensionality, local independence (by Yen's Q3), and good model fit (evaluated through a plot of observed versus expected values).

⁷ Pearl (2000) is the central text here, but also see also Morgan and Winship (2015) on causal modelling in the social sciences, and Keele (2015) for an overview of causal inference in political science.

(e.g., Dawson 1994), religion and social class (Evans and Northmore-Ball 2020), marital status (Denver 2008), and age (Plutzer 2020). We viewed partisanship as a mediator because it is likely affected by political knowledge—specifically, knowledge of party and candidate positioning (Brader and Tucker 2018). Controlling for a mediator, or a causal node located on a direct or indirect pathway between political knowledge and political preference, would misestimate the relevant causal effect.⁸

Our dependent variable was an individual’s reported vote (“Leave,” “Remain,” or “Did not vote.”) We fitted two multinomial logistic models: a *demographic* model, containing all variables mentioned above except partisanship, and a *partisanship* model—containing all variables *including* partisanship—as a robustness check. The models were fitted with R’s *nnet* package (Venables and Ripley 2002) using “doubly robust” estimation of causal effects (e.g., Morgan and Winship 2015). For purposes of such estimation, we recoded the knowledge variable as a binary variable, with all observations representing someone with the maximum ability score being coded as 1 (a full 28% of respondents, suggesting that our definition of a “more informed” voter does not set the bar at an unrealistically high level), and everyone else as 0. Then, we estimated two sets of propensity scores—one for each model—using boosted logistic regression as implemented in the *twang* package (Ridgeway *et al.* 2020). Diagnostics confirmed improved balance between the two groups (those with a maximum score and those not) across measured covariates. These propensity scores were then used as weights in fitting the models.⁹

⁸ Even if partisanship is not a mediator, controlling for it in this context is likely unnecessary. Socialization is centered around group-identity considerations relating to religion, ethnicity, gender, social class, and the like. All of these factors help shape individuals’ conceptions of who they are, and consequently also about what positions “people like us” take in politics (Green *et al.* 2002). Consequently, controlling for such group-level variables, in the manner done here, would already account for partisanship.

⁹ The McFadden value for the partisanship model was 0.222, and the coefficients for the knowledge variable 0.169 for “Leave” ($p = 0.155$) and 0.616 for “Remain” ($p = 0.000$), with “Did not vote” as the reference category. The McFadden value for the demographic model was 0.160, and the coefficients 0.293 for “Leave” ($p = 0.011$) and 0.760 for “Remain” ($p = 0.000$), again with “Did not vote” as the reference category. The variance inflation factor for each predictor did not diverge substantially from 1 on either model, suggesting an absence of multicollinearity.

Finally, to estimate what distribution of support we would likely see for “Leave” and “Remain” in this counterfactually informed electorate, we compared the actual outcome with the distribution estimated by the models after setting the (binary) knowledge variable for each respondent to 1 (representing being maximally informed), with each observation being weighted using the survey weights included with the dataset, to approximate representativeness.¹⁰ Figure 1 gives the actual referendum outcome alongside the estimated outcomes on the two models.

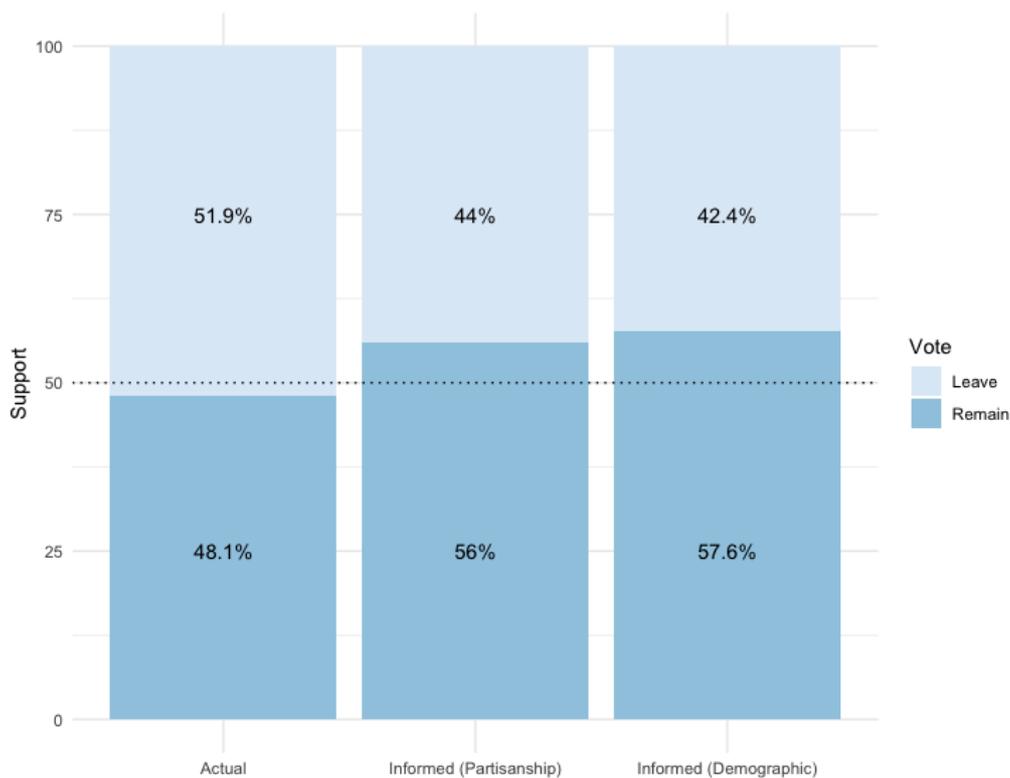


Figure 1. Actual and estimated referendum result under full information, including survey weights.

¹⁰ For each observation, the models ascribed a probability of voting “Remain,” voting “Leave,” or not voting, respectively, given full information. Each probability was multiplied by the observation’s survey weight. To calculate the amount of informed support for “Remain,” the sum of weighted probabilities for “Remain” was then divided by the combined sum of the weighted probabilities for “Remain” and for “Leave” (thereby excluding informed non-voters).

On the partisanship model, the proportion in support for “Leave” drops by 7.9 percentage points, from 51.9% to 44%. Using the demographic model, which assumes partisanship acting as a mediator, the “Leave” vote drops by 9.5 percentage points to 42.4%. Reading across the models, not only does the aggregate result switch from “Leave” to “Remain,” the information effect also exceeds the standard set for a substantial information effect (3 percentage points) by a fairly wide margin. This suggests the results were highly sensitive to voters’ levels of information.

Discussion

It may be well and good to simulate electoral results as we have done, but what does this exercise meaningfully contribute to understandings of politics and political behavior? First, when supported by a robust causal model, this approach provides an empirical and replicable measure of the degree to which an outcome is sensitive to shifts in voters’ political knowledge, which should be of interest to any political scientist interested in the role of information in voter choice. Crucially, this measure does not presume *a priori* that additional information benefits any particular side—a helpful and welcome feature in a political context characterized by motivated reasoning and affective polarization (see Hobolt *et al.* 2020).

Second, as with causal modeling generally, this approach forces us as researchers to be explicit and transparent about the mechanisms we think are at work when people vote, behave, and make sense of political issues. Importantly, the framework also enables us to highlight what assumptions matter, as illustrated above by how robustness checks involving partisanship—clearly an important factor for political choices (e.g., De Vries *et al.* 2018)—enable gauging the sensitivity of effects to different model specifications.

Third, information effects highlight contexts in which changing voters’ levels of information might offer an important lever for politicians, media, and civil societies. That said, successfully isolating the effects of knowledge in a causal model should not be read as offering a recipe for attitude change. For example, it might be true that, had I known *X*, I would have changed my view on some specific topic. But it might simultaneously be the case that motivated reasoning would have prevented

me from ever believing (and therefore knowing) X in the first place. What models of informed choice can do is offer evidence about *where* information-based interventions might make a difference, while accompanying experimental methods can provide clues as to *what* and *for whom* specific interventions are likely to be effective.

Fourth, and perhaps more provocatively, counterfactual modeling itself generates new knowledge accessible for legislators, by providing a measure of the relevant outcome's sensitivity to information.¹¹ *Not* finding substantial information effects—and certainly any effects that would potentially change the aggregate result—could serve as one of several criteria that a government could require as a condition for implementing referenda results. And even if such a criterion were never actually imposed, having empirical evidence of substantial information effects should certainly give us pause in directly and conclusively inferring the “will of the people” from any result, especially in contexts of relatively rare and unusual electoral events such as referenda.

¹¹ Though probably via their friendly neighborhood political scientist.

Works Cited

- Achen, C. and L. Bartels (2016). *Democracy for Realists: Why Elections do not Produce Responsive Government*, Princeton, NJ: Princeton University Press.
- Ahlstrom-Vij, K. (2020), 'The Case for Modelled Democracy,' *Episteme* 1-22. <https://doi.org/10.1017/epi.2020.10>.
- Ahlstrom-Vij, K. (2019). 'The Epistemic Benefits of Democracy: A Critical Assessment,' *the Routledge Handbook of Social Epistemology*, ed. by M. Fricker, P. Graham, D. Henderson, and N. J. L. L. Pedersen, Routledge.
- Althaus, S. (2003). *Collective Preferences in Democratic Politics*, Cambridge: Cambridge University Press.
- Bartels, L. (1996). 'Uninformed Voters: Information Effects in Presidential Elections,' *American Journal of Political Science* 40(1): 194–230.
- Bhatti, Y. (2010). 'What Would Happen if We Were Better Informed? Simulating Increased Knowledge in European Parliament (EP) Elections?' *Representation* 46(4): 391-410.
- Blais, A., Gidengil, E., Fournier, P., and Nevitte, N. (2008). 'Information, Visibility and Elections: Why Electoral Outcomes Differ When Voters are Better Informed,' *European Journal of Political Research* 48: 256-280.
- Brader, T. and J. A. Tucker (2018). 'Unreflective Partisans? Policy Information and Evaluation in the Development of Partisanship,' *Advances in Political Psychology* 39, Suppl. 1: 137-157.
- Dawson, M. C. (1994). *Behind the Mule: Race and Class in African American Politics*, Princeton: Princeton University Press.
- Delli Carpini, M. X., and S. Keeter (1996). *What Americans Know about Politics and Why It Matters*, New Haven, CT: Yale University Press.
- Denver, D. (2008). 'Another Reason to Support Marriage? Turnout and the Decline of Marriage in Britain,' *British Journal of Politics and International Relations* 10(4): 666-680.
- De Vries, C. E., S. B. Hobolt, and J. Tilley (2018). 'Facing up to the Facts: What Causes Economic Perceptions?' *Electoral Studies* 51(February): 115–122.
- Evans, G. and K. Northmore-Ball (2020). 'Long-Term Factors: Class and Religious Cleavages,' in Fisher et. al. (2020): 123-135.

- Fieldhouse, E., J. Green, G. Evans, H. Schmitt, C. van der Eijk, J. Mellon, and C. Prosser (2018). *British Election Study, 2017: Face-to-Face Post-Election Survey* [data collection].
<http://doi.org/10.5255/UKDA-SN-8418-1>.
- Fiorina, M. (1981). *Retrospective Voting in American National Elections*, New Haven, CT: Yale University Press.
- Fisher, J., E. Fieldhouse, M. N. Franklin, R. Gibson, M. Cantijoch and C. Wlezen (2020) *The Routledge Handbook of Elections, Voting Behavior and Public Opinion*, Abingdon: Routledge.
- Friedman, J. (1998). 'Introduction: Public Ignorance and Democratic Theory,' *Critical Review* 12(4): 397-411.
- Goodin, R. E. (1995). *Utilitarianism as Public Policy*. Cambridge: Cambridge University Press.
- Goodwin, M., S. Hix, and M. Pickup (2018). 'For and Against Brexit: A Survey Experiment of the Impact of Campaign Effects on Public Attitudes toward EU Membership'. *British Journal of Political Science*, 1–15. <https://doi.org/10.1017/S0007123417000667>.
- Green, D. P., B. Palmquist, and E. Schickler (2002). *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*, New Haven, CT: Yale University Press.
- Harrell, F. E, C. Dupont, and collaborators (2019). 'Hmisc: Harrell Miscellaneous,' *R* package version 4.2-0; <https://CRAN.R-project.org/package=Hmisc>.
- Harsanyi, J. C. (1997). 'Utilities, Preferences, and Substantive Goods,' *Social Choice and Welfare* 14: 129-145.
- Hebbelstrup, S. and R. Rasmussen (2016). 'Education or Personality Traits and Intelligence as Determinants of Political Knowledge?' *Political Studies* 64(4): 1036-1054.
- 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*, 1–18.
<https://doi.org/10.1017/S0007123420000125>.
- Key, V.O. (1966). *The Responsible Electorate: Rationality in Presidential Voting 1936-1960*, Cambridge, MA: Belknap Press.
- Lupia, A. (2006). 'How Elitism Undermines the Study of Voter Competence,' *Critical Review* 18(3): 217-232.
- Oscarsson, H. (2007). 'A Matter of Fact? Knowledge Effects on the Vote in Swedish General Elections, 1985–2002,' *Scandinavian Political Studies* 30(3): 301-322.
- Plutzer, E. (2020). 'Demographics and the Social Bases of Voter Turnout,' in Fisher *et. al.* (2020): 69-82.
- Popkin, S. L. (1991). *The Reasoning Voter: Communication and Persuasion in Presidential Elections*, Chicago: University of Chicago Press.

- R Core Team (2017). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Ridgeway, G., D. McCaffrey, A. Morral, B. A. Griffin, L Burgette, and M. Cefalu (2020). ‘*twang*: Toolkit for Weighting and Analysis of Nonequivalent Groups,’ *R* package version 1.6; <https://CRAN.R-project.org/package=twang>.
- Rizopoulos, D. (2006). ‘ltn: An R package for Latent Variable Modelling and Item Response Theory Analyses,’ *Journal of Statistical Software* 17(5): 1-25.
- vanHeerde-Hudson, J. (2020). ‘Political Knowledge: Measurement, Misinformation and Turnout,’ in Fisher et. al. (2020): 369-382.
- Venables W.N. and Ripley B.D. (2002). *Modern Applied Statistics with S*, 4th edition, New York, NY: Springer.
- Vowles, J. (2020). ‘The Big Picture: Turnout at the Macro-Level,’ in Fisher et. al. (2020): 57-68.
- Weale, A. (2018). *The Will of the People: A Modern Myth*, Cambridge: Polity Press.